

Coarse Bayesian Updating

Alexander M. Jakobsen*

November 2, 2025

Abstract

Studies have shown that the standard law of belief updating—Bayes’ rule—is descriptively invalid in various settings. In this paper, I introduce and analyze a generalization of Bayes’ rule—*Coarse Bayesian updating*—accommodating much of the empirical evidence. I characterize the model axiomatically, show how it generates several well-known biases, and derive its main implications in static and dynamic settings. Each axiom expresses a property of Bayes’ rule but, conditional on the others, stops just short of making the agent fully Bayesian. The model employs standard primitives, making it suitable for applications; I demonstrate this by applying it to a standard setting of decision under risk, leading to a close relationship with the Blackwell information ordering and comparative measures of cognitive sophistication and bias.

1 Introduction

Bayesian updating plays a central role in economic theory. A number of studies, however, document behavior that cannot be reconciled with Bayes’ rule. For example, individuals may under-react to new information or even ignore it altogether; others may over-react by falsely extrapolating or, more generally, engaging in pattern-seeking behavior. “Motivated” reasoning, among other mechanisms, may lead individuals to under-react to some signals but over-react to others. Still others may be Bayesian except when information is too extreme or unexpected. Such heterogeneity, both within and between individuals, poses an interesting challenge to the Bayesian paradigm and calls not just for new models of behavior, but for

* Kellogg MEDS, Northwestern University; alexander.jakobsen@kellogg.northwestern.edu. I thank the editor and three anonymous referees for feedback that substantially improved the paper. I am also grateful to David Dillenberger, Yoram Halevy, Shaowei Ke, Dimitri Migrow, Rob Oxoby, Scott Taylor, Benjamin Young and seminar and conference audiences at UC Davis, Edinburgh, Kellogg MEDS, Michigan, Montreal, UNSW, Purdue, UTS, Toronto, UWO, BSE Summer Forum (2022), 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.

analysis of the features of Bayesianism that are compatible or incompatible with the wide range of documented behavior.

In this paper, I introduce and analyze a generalization of Bayesian updating—*Coarse Bayesian updating*—encompassing many documented phenomena. The main results characterize, axiomatically, both the model and the gap between it and standard Bayesian updating. There are two advantages to this approach. First, each axiom expresses a falsifiable property of Bayes’ rule. This provides a normative foundation for Coarse Bayesian behavior and assurance that, despite its flexibility, the model is not so general as to accommodate anything. Second, the characterizations identify not just which properties of Bayes’ rule are compatible with Coarse Bayesian behavior, but also which properties are necessarily violated by proper (non-Bayesian) Coarse Bayesians. This provides a transparent comparison of the model’s conceptual trade-offs and, as described below, a sense in which it is a small departure from standard Bayesianism. The end result is a framework capturing a variety of departures from Bayes’ rule while remaining tractable for economic applications; I illustrate this by deriving its main implications in a general setting of decision under risk.

Intuitively, a Coarse Bayesian simplifies the world by considering only a subset of the probability space. Given this restriction, the agent applies subjective criteria to switch among beliefs in that set. More precisely, a Coarse Bayesian agent is characterized by (i) a partition of the probability simplex into convex cells, and (ii) a representative distribution for each cell of the partition, one of which is the prior. After observing a signal, the agent determines which cell contains the Bayesian posterior and adopts the representative of that cell as posterior belief (see Figure 1). Importantly, the agent need not point-identify the Bayesian posterior; instead, he merely approximates it by determining which cell it belongs to. For example, the agent might analyze information in small steps, gradually eliminating candidates for the true distribution until only one cell remains in contention. Thus, the procedure is not “more difficult” than Bayesian updating to begin with.

The parameters of the model—cells and their representative points—are 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 (1954), then, Coarse Bayesians exhibit subjectivity not only in their prior beliefs but in their criteria for revising those beliefs. Consequently, different agents may exhibit over-reaction, under-reaction, or other biases depending on the signal, the partition, and the positions of representative points within their cells. There are several ways of interpreting this behavior, such as categorical thinking or signal distortion—I discuss these, and other, interpretations in section 2.

The first result provides a simple characterization of the updating procedure. I take as primitive a finite, exogenous state space and an updating rule specifying an individual’s be-

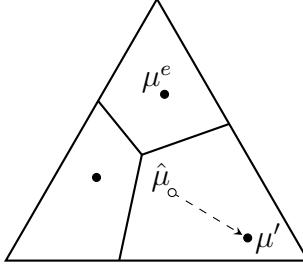


Figure 1: Coarse Bayesian updating. In this example, there are three feasible beliefs (solid dots). The point μ^e is the prior. After observing a signal, the agent determines which cell of the partition contains the Bayesian posterior $\hat{\mu}$, then adopts the representative of that cell (in this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

iefs at every possible signal. In my framework, signals represent messages that can be generated by stochastic information structures. Thus, a signal is a profile of numbers representing likelihoods of the associated message being generated in different states. By employing such primitives, the model is readily adaptable to standard economic or game-theoretic settings.

The characterization involves three testable axioms on the updating rule, each capturing a property of standard Bayesian behavior. First, *Homogeneity* states that beliefs are invariant to scalar transformations of signals: 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 axiom 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.

Next, Proposition 1 establishes that, under mild assumptions, strengthening any of the axioms to an if-and-only-if form forces the agent to be Bayesian. For example, *Homogeneity* states that if two signals have the same likelihood ratios, then they induce the same beliefs. The proposition implies that if one adds a fourth axiom, “two signals have the same likelihood ratios if they induce the same beliefs” (the converse to *Homogeneity*), then the agent must be Bayesian—the added responsiveness to information implied by this converse statement closes the gap between Bayesian and Coarse Bayesian behavior. The same property holds for *Cognizance* and *Confirmation*: adding the converse statement to either axiom makes the agent Bayesian. This is the sense in which Coarse Bayesian updating is, qualitatively,

a “small” departure from Bayes’ rule. Moreover, since Bayes’ rule satisfies all three axioms and their converse statements, the converse statements capture the features of Bayes’ rule that are violated by (proper) Coarse Bayesians.

Section 3 explores the main implications of the model and how it might be applied. In section 3.1, I discuss how the framework can be used as a tool for modeling various biases (section 3.1.1), for predicting or understanding real behavior (section 3.1.2), or for testing or identifying coarse cognition in experiments (section 3.1.3). For example, I show in section 3.1.1 that Coarse Bayesians may exhibit over/under-reaction, “motivated” belief updating, limited perception, or other biases and establish in section 3.1.2 that, generically, all non-Bayesian behavior in the model stems from the combination of three particular biases (an implication of Proposition 1); this fully characterizes the predictions of the model and informs part of the discussion in section 3.1.3 on the design of experiments. I also discuss, at an intuitive level, how the framework can be used in more applied settings to shed light on behavior like financial decision making, stereotyping, and discourse or non-informative persuasion. Finally,¹ section 3.2 explores some basic properties of the model in dynamic settings. I consider two categories of dynamic updating rules: *pooled* rules and *sequential* rules. Pooled rules incorporate, at every time period, the full history of signal realizations; consequently, pooled rules satisfy strong forms of path-independence and have simple convergence properties. Sequential rules, however, involve signal-by-signal updating, introducing various degrees of path dependence and more nuanced convergence properties.

Section 4 applies the model to a standard setting of decision under risk. 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, then establish 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. I characterize the menus (decision problems) in which a given Coarse Bayesian benefits from Blackwell improvements and show that the connection runs much deeper: two Coarse Bayesians are identical—same cells, same representative points—if and only if they benefit from the same Blackwell improvements. Thus, the parameters of the model can be uniquely identified from the agent’s menu-contingent rankings of Blackwell-comparable experiments.

In section 4.2, I examine how a Coarse Bayesian’s welfare changes as he becomes “more Bayesian.” I consider three such orderings. First, an agent is *more sophisticated* if he employs a finer partition. I show that more-sophisticated agents are characterized by height-

¹The Online Appendix contains additional results related to observational learning and the connection between my model and others involving maximum likelihood-style reasoning.

ened responsiveness to information, as captured by ex-ante value of information. Second, one agent is *more biased* than another if his updating rule exhibits larger distortions away from Bayesian posteriors. I show that greater bias is characterized by greater susceptibility to harmful exploitation in that worst-case losses, relative to a Bayesian, increase as bias increases. Importantly, neither greater sophistication nor lower bias imply the agent is better off at all menus or signal realizations. The final result shows that such welfare enhancements require the agent to be perfectly Bayesian on a larger set of signal realizations, giving rise to a third ordering that jointly refines the sophistication and bias orderings.

Throughout the paper, my focus is on the general class of Coarse Bayesian representations and their properties. I do not take a stance on where partitions or representative elements “come from,” viewing them instead as subjective (but identifiable) characteristics of an individual, much like subjective prior beliefs. There are several ways to restrict or endogenize the parameters by adding assumptions about the decision problem(s) agents expect to face, the signaling structure, and costs or constraints on the fineness of the updating rule (for example, a bound on the number of cells in the partition). The results of section 4.2 suggest a slightly different approach may be valuable: rather than solving for an optimal updating rule in the context of a specific environment, one may prefer a more robust objective—characterized by the bias ordering, for example—accommodating uncertainty about the environment. I discuss this at the end of section 4.2.

To summarize, the main contribution of the paper is a new model of belief updating accompanied by analysis of its essential properties and implications. Few, if any other models proposed in the literature can accommodate the range of behavior studied in section 3.1. Any model that can is necessarily quite flexible, but the Coarse Bayesian framework has some advantages. First, it is testable: not all behavior satisfies the axioms. Second, it involves a clear separation between properties of Bayes’ rule that are compatible with biases in belief updating and those that are not. Third, the model is portable and tractable for applications; it is one of just a few to take general stochastic signals as the starting point, allowing it to be directly imported to standard settings in economics and game theory. Section 4 demonstrates this by deriving the main implications of Coarse Bayesian updating in standard settings of decision under risk—a core component of any new model of updating behavior. Finally, it is simple: the key ideas can be captured by a single picture (Figure 1), and the axioms and characterizations are easy to state, prove, and interpret.

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 experimental protocols² and the patterns of behavior displayed by subjects. For example, under-reaction is quite common but by no means an established law of behavior—over-reaction occurs as well; there is mixed evidence for asymmetric processing of ego-relevant information—subjects may or may not respond differently to good news than they do to bad news; and numerous studies document individual heterogeneity—some subjects are more Bayesian than others (see Benjamin, 2019 for a survey and meta-analysis of the literature).

Motivated by this evidence, several authors have developed models to better understand the mechanisms behind, and consequences of, non-Bayesian updating. Models 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). My emphasis, particularly in sections 3 and 4, is on implications that are reasonably independent of any particular application. As such, I employ standard primitives (a finite state space; stochastic information structures; general decision problems) that can be adapted to any economic model.

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, over-reaction, and other biases; Epstein et al. (2008) extend this model to an infinite-horizon setting. Zhao (2022) axiomatizes an updating rule for signals indicating that one event is more likely than another. Like these authors, I take a general approach and characterize behavior axiomatically. My model is not geared toward a specific bias or application, but provides a general framework of coarse cognition that accommodates (and generates) a variety of non-Bayesian behavior.

Coarse Bayesian updating resembles, to a degree, the well-known representativeness heuristic of Kahneman and Tversky (1972), wherein an individual “evaluates the probability of an uncertain event, or a sample, by the degree to which it is: (i) similar in essential properties to its parent population; and (ii) reflects the salient features of the process by which it is generated” (Kahneman and Tversky, 1972 p. 431). One might interpret Coarse Bayesian representations—cells and their representative points—as a way of formalizing the representativeness heuristic by providing an agent’s subjective assessment of “similarity,” “essential

²For example, studies differ in whether subjects observe individual signals or larger samples/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.

properties,” or “salient features.” There are at least two problems with this. First, the representativeness heuristic requires agents to ignore base rate (prior) information, which is inconsistent with Coarse Bayesian behavior: the prior is directly relevant to a Coarse Bayesian because the agent adopts the representative of the cell containing the Bayesian posterior. Second, the Coarse Bayesian framework accommodates behavior that is at odds, intuitively, with the representativeness heuristic. For example, Coarse Bayesian updating permits agents to be perfectly Bayesian as long as they “notice” the signal (see section 3.1). Despite the freedom afforded by the definition above, it would be a stretch to categorize such behavior as an instance of the representativeness heuristic when other explanations, like limited attention, seem more appropriate. In section 2, I offer other interpretations of Coarse Bayesian behavior that avoid these difficulties.

Three studies are especially relevant to Coarse Bayesian updating. First, the *hypothesis testing* model introduced by Ortoleva (2012) posits that agents apply Bayes’ rule except when news is sufficiently “surprising,” in which case a maximum-likelihood criterion is applied using a second-order prior.³ Specifically, an agent applies Bayes’ rule if the prior probability of the signal 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. I show that Coarse Bayesian updating can accommodate similar behavior and (in the Online Appendix) compare Coarse Bayesian rules to a general class of Maximum-Likelihood updating rules. I show that Coarse Bayesian rules can be expressed as Maximum-Likelihood rules if there are only two states but that, in general, neither category subsumes the other. Notably, Maximum-Likelihood rules may violate the Confirmation property—perfect evidence of a candidate belief does not guarantee that that belief is selected.

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 can be represented as a (dynamic) Coarse Bayesian updating procedure. Naturally, the parameters of this representation—cells and their representative points—depend on features of the environment like the signal structure, the stakes of the decision problem, and the bound K . Like Bayesian updating, Coarse Bayesian updating procedures do not depend on any factors other than the informational content of realized signals. I do not require Coarse Bayesian representations to be optimal in any sense, nor do

³See also Dominiak et al. (2023), who study a behaviorally equivalent model.

I impose cognitive bounds such as a restriction on the number of cells. This allows my model to capture documented behavior (for example, Bayesian updating except when signals are too “extreme”—see section 3.1) that is inconsistent with Wilson’s model.

Third, in 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 Mullainathan’s 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 the Online Appendix for a concrete example.

2 Model

Let $\Omega = \{1, \dots, N\}$ 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 with N rows, finitely many columns, and entries in $[0, 1]$ such that each row is a probability distribution and each column has a nonzero entry. Columns represent messages that might be generated, and rows state-contingent probability distributions over messages. Let \mathcal{E} denote the set of all experiments, with generic element σ .

As in Jakobsen (2021), 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. A signal s represents a column (message) of some experiment, and s_ω the likelihood of the message being generated in state ω . The notation $s \in \sigma$ indicates that s is a column of σ . I reserve e to denote the signal $e \in S$ such that $e_\omega = 1$ for all $\omega \in \Omega$; note that e qualifies as an (uninformative) experiment.⁴ Using the notation of signals, an experiment can be viewed as a collection (matrix) of signals, of state-contingent distributions over signals, or of points in S that sum to e ; see Figure 2.

For 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 $\|v\|$.

For $\hat{\mu} \in \Delta$ and $s \in S$ where $s \cdot \hat{\mu} \neq 0$, let $B(\hat{\mu}|s) := \frac{s \hat{\mu}}{s \cdot \hat{\mu}} \in \Delta$ denote the **Bayesian**

⁴Any signal s such that $s_\omega = s_{\omega'}$ for all $\omega, \omega' \in \Omega$ is uninformative, as is any experiment composed of such signals; e is a convenient representative because it qualifies as both a signal and an experiment.

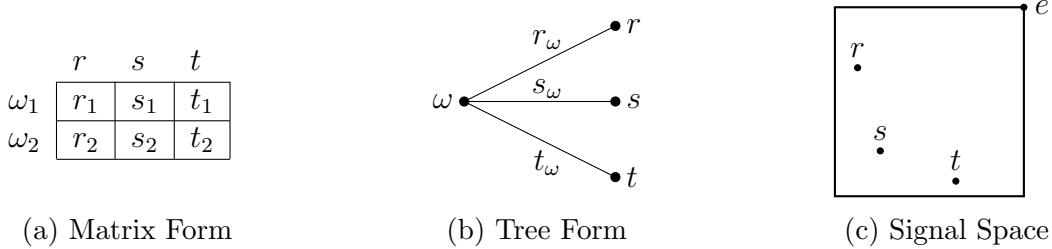


Figure 2: Three representations of an experiment $\sigma = [r, s, t]$.

posterior of $\hat{\mu}$ at s . To allow non-Bayesian behavior, the key primitive of the model is a function $\mu : S \rightarrow \Delta$ (an **updating rule**) assigning distributions $\mu^s := \mu(s) \in \Delta$ to signals $s \in S$. The interpretation is that μ^s is the agent's posterior belief conditional on observing signal s . Being a function on S , the updating rule specifies beliefs at all conceivable signals, not just those generated by a particular experiment. I assume μ^e , the **prior**, has full support.⁵

My notion of an updating rule implicitly makes two assumptions about behavior. First, updating rules condition beliefs on signal realizations s but not the experiment(s) generating them. In practice, one might record posterior beliefs as $\mu^{(\sigma, s)}$ where $s \in \sigma$, changing the domain of μ to (a subset of) $\mathcal{E} \times S$. Like Bayesian updating, however, Coarse Bayesian updating depends on s but not the other columns of σ ; therefore, I omit dependence on experiments σ . Second, the agent's prior coincides with his posterior belief after observing e . This, too, is a property of Bayes' rule that Coarse Bayesians satisfy.

2.1 Coarse Bayesian Representations

Consider an agent whose behavior is summarized by an updating rule $\mu : S \rightarrow \Delta$. In this section, I show that Coarse Bayesian updating is characterized by three axioms on μ . Each axiom expresses a property of Bayes' rule and is falsifiable with data in the form of an updating rule. The axioms also lead to a simple comparison with Bayes' rule (Proposition 1), capturing the sense in which the model is a “small” departure from Bayes' rule and the exact properties of Bayesian rationality that are violated by (proper) Coarse Bayesians.

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

Homogeneity requires the agent's analysis of a signal s to depend only on the likelihood

⁵Some authors define updating rules as functions $\varphi : \Delta \times S \rightarrow \Delta$, with the interpretation that $\varphi(\hat{\mu}, s)$ is the posterior belief given a prior $\hat{\mu}$ and signal s . That approach is not appropriate here because it assumes the agent can hold any belief $\hat{\mu}$. Coarse Bayesians entertain only a subset of Δ as possible beliefs, so one cannot freely vary the prior. In my model, the set of candidate beliefs (including the prior) is subjective and revealed via updating behavior. Dynamic extensions (section 3.2) that condition beliefs on histories of signals provide a way to study updating as the agent moves around the set of candidate beliefs.

ratios $s_\omega/s_{\omega'}$. This is a key feature of Bayesian updating and it implies 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, if necessary, 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 axioms.

Axiom 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 has realized. This interpretation stems from the fact that $s + t$ is a “garbled” signal indicating that either s or t was generated.⁶ 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 to the same posterior belief—that is, if he is cognizant of his own updating procedure—then he ought to adopt that belief.⁷

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 of updating behavior. For such exercises to make sense, an assumption like Cognizance is required.

Axiom 3 (Confirmation). If $t \approx \mu^s/\mu^e$, then $\mu^t = \mu^s$.

To understand Confirmation, observe that the set of attainable beliefs, $\{\mu^s : s \in S\}$, serves as a kind of consideration set—no other points in Δ are candidate posteriors. Confirmation states that if the Bayesian posterior is a point in that set, the agent adopts that point as posterior belief. In particular, $t \approx \mu^s/\mu^e$ satisfies $B(\mu^e|t) = \mu^s$ (that is, t *confirms* μ^s by providing “perfect” evidence of it for an agent with prior μ^e), so $\mu^t = \mu^s$. In contrapositive form, this means that if μ^s does not coincide with the Bayesian posterior at s , then the Bayesian posterior is not a candidate belief. Although quite intuitive and normatively appealing, Confirmation is not satisfied by some related models—see the Online Appendix.

⁶For example, if $s, t \in \sigma$, there is a garbling matrix M such that $\sigma' = \sigma M$ collapses columns s and t to a single column $s + t$ without altering any other columns of σ .

⁷Axioms 1 and 2 can be combined into one statement: if $\mu^s = \mu^t$, $\alpha, \beta \geq 0$ and $\alpha s + \beta t \in S$, then $\mu^{\alpha s + \beta t} = \mu^s$ (it would not suffice to replace the statement with convex combinations of s and t , ie, $\mu^{\alpha s + (1-\alpha)t} = \mu^s$, even though this property is implied by the axioms; conic combinations are needed). I have separated this statement into two axioms because they capture intuitively distinct features of behavior.

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) *$\mu^P \in P$ for all $P \in \mathcal{P}$, 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

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

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 Axioms 1–3 are satisfied (Axiom 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 $\mu^{\alpha s + (1-\alpha)t} = \mu^s$. Combined with Homogeneity, 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 equation (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 $C \in \mathcal{C}$. For each $P \in \mathcal{P}$, let μ^P denote the unique distribution $\hat{\mu}$ such that $\mu^s = \hat{\mu}$ for all $s \in C$, where $P = P^C$. Confirmation implies $\mu^s = \mu^P \in P$ whenever $B(\mu^e|s) \in P \in \mathcal{P}$. Uniqueness of $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ follows from uniqueness of \mathcal{C} . \square

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 axioms imposes a degree of Bayesian rationality on the agent by expressing a property of Bayes' rule—indeed, each axiom is satisfied by a standard Bayesian. As we shall see, Coarse Bayesian updating nonetheless accommodates a variety of behavioral biases and other violations of Bayes' rule.

Coarse Bayesians partition the probability simplex, assign a representative point to each cell, and adopt the representative of a cell as posterior if the Bayesian posterior belongs to that cell. Why might an agent behave this way? Below, I offer four interpretations of the model, some of which may be more appropriate than others depending on the application.

1. *Competing Theories.* Here, the agent simplifies the world by considering a set of candidate theories (representative points), sets criteria (the partition) for switching between them, and analyzes signals to the extent necessary to determine whether a change is justified. The agent is only interested in whether the evidence satisfies his “standard of proof” for a given theory, so he does not necessarily point-identify the Bayesian posterior. For example, he might process signals in small steps, gradually eliminating points in the simplex until he determines which cell applies.

2. *Limited Computation.* An agent might wish to compute the Bayesian posterior but be unable to point-identify it. Consequently, the agent lumps several posteriors together with a single point, making the representation a simplifying heuristic or approximation to Bayes' rule. Since different agents may employ different partitions or representative points, they may disagree on what constitutes a hard problem or a good approximation.

3. *Signal Distortions.* Here, to update beliefs, the agent mentally transforms signals before applying Bayes' rule. Thus, apparent deviations from Bayes' rule are the result of imperfect perception or attention—not necessarily computational constraints. Theorem 2 below formalizes the concept of Signal Distortion Representations and establishes their equivalence to Coarse Bayesian Representations in static settings.⁸ In dynamic settings, however, the distinction matters (see section 3.2).

4. *Categorical Thinking.* Here the agent reasons about categories of beliefs, each represented by a cell of the partition. This way, a cell represents distributions that share some properties of interest, and its representative point is a natural example (or “archetype”) of a distribution with those properties. When information arrives, the agent determines which category applies and adopts its archetype as posterior. The key difference between this and

⁸The term “signal distortion” is often associated with Grether (1980) updating, where distorted signal and/or prior likelihoods are used in the updating process. My notion of signal distortion employs a different functional form and does not involve distorted priors; see Theorem 2 below.

the competing-theories interpretation is that here the cells (not their representative points) are focal: the agent is primarily interested in whether the true distribution belongs to a given category, and uses representative points to envision the category.

In each case, the parameters of the representation are subjective *characteristics of the individual*: agents may differ in their priors, partitions, or representative points. In the same way that standard Bayesian theories are agnostic about the source of one's prior beliefs, my model does not take a stance on how partitions or representative points are formed. Rather, Theorem 1 characterizes Coarse Bayesian behavior in terms of observable primitives (the updating rule) and establishes that all parameters can be uniquely identified from those primitives—with or without additional assumptions about how they came to be.⁹

The next result provides a simple comparison between Bayesian and Coarse Bayesian behavior. To proceed, an additional definition is required; some subsequent results in the paper also utilize this definition.

Definition 1. Given $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$, a cell $P \in \mathcal{P}$ is **regular** if it has full dimension in Δ and its representative μ^P belongs to the relative interior of P . If every cell $P \in \mathcal{P}$ is regular, then $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ is **regular**.

Proposition 1. *Suppose μ is non-constant and has a Coarse Bayesian Representation $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ where every non-singleton cell of \mathcal{P} is regular. Then μ is Bayesian (that is, $\mu^s = B(\mu^e|s)$ for all $s \in S$) if and only if any of the following three conditions hold:*

- (i) $\mu^s = \mu^t$ implies $s \approx t$;
- (ii) $\mu^{s+t} = \mu^s$ implies $\mu^s = \mu^t$;
- (iii) $\mu^t = \mu^s$ implies $t \approx \mu^s/\mu^e$.

Proposition 1 states that, under mild regularity conditions, strengthening *any* of Axioms 1–3 to an if-and-only-if form forces a Coarse Bayesian agent to be perfectly Bayesian. Statement (i), the converse to Homogeneity, makes the agent highly responsive to changes to information: different likelihood ratios lead to different posterior beliefs. Statement (ii), the converse to Cognizance, requires that if the agent is unaffected by the knowledge that t may have been generated instead of s , then s and t must lead to the same beliefs. Finally, statement (iii), the converse to Confirmation, asserts that if t leads to the same posterior as s , then t must be perfect evidence of μ^s . These statements are themselves properties

⁹See also the discussion at the end of section 4.2 regarding approaches to endogenizing the parameters.

of Bayes' rule, and the proposition implies that if a Coarse Bayesian agent satisfies any of them, then the agent actually satisfies all three and behaves like a standard Bayesian.¹⁰

There are two key takeaways from Proposition 1. First, the “wedge” between Bayesian and Coarse Bayesian updating is, qualitatively, fairly small: Axioms 1–3 impose enough Bayesian rationality that slightly strengthening any of them eliminates non-Bayesian behavior. Nonetheless, the model permits many documented departures from Bayes' rule. Thus, one can accommodate a variety of non-Bayesian behavior without abandoning tenets of Bayesian rationality that, combined, almost make the agent perfectly Bayesian. Second, Proposition 1 identifies the properties of Bayes' rule that are necessarily violated by proper Coarse Bayesians. This provides a more comprehensive understanding of the model, enables direct comparison of its conceptual trade-offs and, as we shall see, leads to a full characterization of the non-Bayesian behavior predicted by the model.

I conclude this section by providing an alternative representation of Coarse Bayesian behavior and a brief discussion of some limitations of the model.

Theorem 2. *An updating rule μ has a Coarse Bayesian Representation if and only if there is a function $d : S \rightarrow S$ such that*

- (i) $s \approx t$ implies $d(s) \approx d(t)$,
- (ii) $d(s) \approx d(t)$ implies $d(\lambda s + (1 - \lambda)t) \approx d(s)$ for all $\lambda \in [0, 1]$,
- (iii) $d(d(s)) = d(s)$ for all s ,

and $\mu^s = B(\mu^e | d(s))$ for all s . The function d is a **Signal Distortion Representation** of μ . If d' is another such representation, then $d'(s) \approx d(s)$ for all $s \in S$.

Signal Distortion Representations formalize the signal distortion interpretation of Coarse Bayesian behavior, replacing the parameters $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ with a function d satisfying three properties analogous to Axioms 1–3; in particular, an agent who receives signal s applies Bayes' rule to a distorted signal $d(s)$. Theorem 2 establishes that Coarse Bayesian and Signal Distortion behavior is equivalent in static settings; however, as shown in section 3.2, this equivalence fails in dynamic settings.

Naturally, Coarse Bayesian updating is not without its limitations. Although Theorem 1 and Proposition 1 fully characterize and contrast Coarse Bayesian and standard Bayesian behavior, it is worth highlighting a few additional implications of the framework.

¹⁰The regularity assumptions of Proposition 1 are only needed to establish that statement (ii) forces the agent to be Bayesian—statements (i) and (iii) each make *any* Coarse Bayesian perfectly Bayesian.

1. *Only the realized signal matters.* More precisely, Homogeneity requires that only the likelihood ratios of s can affect posterior beliefs. This rules out sensitivity to the way information is framed, as well as the possibility that extraneous features of the environment might impact beliefs.
2. *Beliefs are represented by probability distributions.* For example, the conjunction fallacy (illustrated by the well-known “Linda” problem of Tversky and Kahneman, 1983) occurs when subjects declare an event E less likely than a conjunction $E \cap F$. Such beliefs cannot be represented by probability distributions and therefore fall outside the scope of the model.
3. *Discontinuities in s .* Jumps can occur when perturbations to a signal make the Bayesian posterior cross over a cell boundary. This is a feature of any model involving threshold-style behavior, including that of Wilson (2014), the Categorical-Thinking model of Mullainathan (2002), the Hypothesis-Testing model of Ortoleva (2012) and the related Maximum Likelihood models examined in the Online Appendix. If continuity is an essential conceptual feature of some pattern of behavior—rather than a convenient technical assumption—then Coarse Bayesian updating will, at most, provide an approximation to that behavior.
4. *Convex cells.* This convexity is driven by Cognizance and can be discarded by dropping that axiom. However, as explained above, Cognizance is potentially important in applications because it means agents correctly forecast their own updating behavior.

3 Implications, Applications, and Dynamics

This section explores the main implications and applications of the model as well as its relationship to empirical work on non-Bayesian updating. Sections 3.1 and 3.2 are independent of each other and can be read in any order; section 4 is also independent of this section.

3.1 Implications and Applications

To map out the main implications and applications of Coarse Bayesian updating, this section considers three ways one might employ the framework: (i) as a tool for modeling specific biases (section 3.1.1), (ii) as a model for predicting or understanding behavior (section 3.1.2), and (iii) as a guide for designing experiments and testing for coarse cognition (section 3.1.3).

3.1.1 Modeling Biased Updating

The Coarse Bayesian framework does not target a specific bias (or collection thereof) but instead provides a standalone model of coarse cognition that will, in various circumstances,

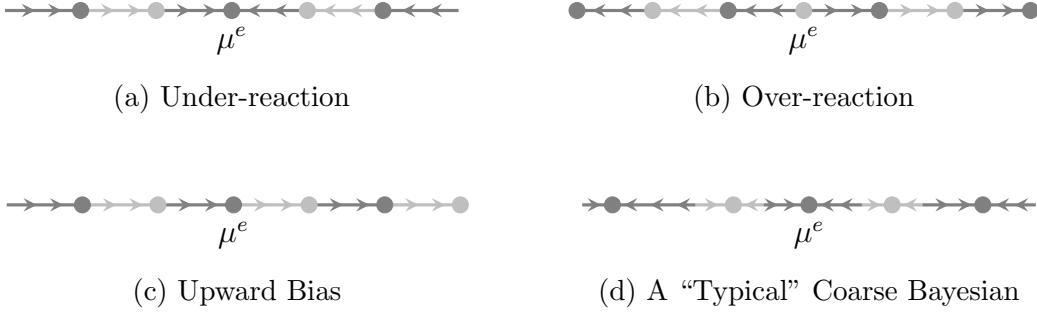


Figure 3: Four Coarse Bayesian Representations on $\Delta = [0, 1]$.

generate well-known biases. That said, one might be interested in modeling specific biases when bringing the framework to applications. Below, I illustrate how a variety of documented biases can be represented as Coarse Bayesian behavior. Since they are Coarse Bayesian Representations, Axioms 1–3 provide foundations while other findings in the paper, like those of sections 3.2 and 4, provide general tools and results for their analysis.

1. *Under-reaction, Over-reaction, and Asymmetric Updating.* Conservative updating, or under-reaction to information, is a well-documented behavior violating Bayes’ rule.¹¹ Benjamin (2019) conducts meta-analysis of the experimental literature and finds that under-reaction is the most common bias. On the other hand, individuals also over-react to information in various settings. De Bondt and Thaler (1985), for example, find evidence of over-reaction in financial markets (in particular, to unexpected news); more recently, Thaler (2021) finds evidence of over-reaction to weak signals and under-reaction to strong signals.

When information is “ego-relevant,” subjects may respond asymmetrically to information. Eil and Rao (2011) find that if information concerns personal attributes such as attractiveness, individuals under-react to negative signals but are approximately Bayesian when processing positive signals; see also Sharot and Garrett (2016) for a survey of related studies.

To represent such behavior in the Coarse Bayesian framework, I follow the literature by considering two-state settings; this way, Δ can be identified with the unit interval. Figures 3a and 3b illustrate under- and over-reaction. In 3a, the agent never over-reacts but typically under-reacts: his posterior belief (solid dot) is as close as possible to μ^e given the partition of Δ into sub-intervals (light/dark gray regions representing different cells). In 3b, the agent never under-reacts but typically over-reacts: his posterior is farthest away from μ^e given the partition. Figure 3c exhibits a biased agent who favors one state: posteriors typically assign higher probability to state 1 than the Bayesian posterior and never less. Thus, it is relatively easy for the agent to revise beliefs upward and more difficult to revise downward.

¹¹See Phillips and Edwards (1966) and Edwards (1968) for early experiments on conservative updating.

A common feature of Figures 3a-3c is that representative points of cells sit on cell boundaries; this is needed to model such biases in the Coarse Bayesian framework because there is either an ideal belief the agent aspires to or an unappealing belief they seek to avoid. Figure 3d depicts a more typical Coarse Bayesian: representative points do not necessarily sit on the boundaries of cells, so both over- and under-reaction occur, depending on the signal.

2. *Limited Perception, Extreme-Belief Aversion, and Reactions to Unexpected News.* The model also accommodates agents who behave like standard Bayesians except in particular circumstances. For example, consider Figure 4a. In this representation, the agent retains prior μ^e unless the Bayesian posterior is sufficiently far away from μ^e , in which case he applies Bayes' rule. An interpretation is that the agent only notices signals that are sufficiently strong or provocative to yield a large shift in the Bayesian posterior. This provides a way of capturing imperfect attention or perception.¹²

Figure 4b exhibits rather the opposite behavior: the agent is Bayesian unless posterior beliefs are too “extreme”—that is, close to degenerate distributions representing certainty about the state. Ducharme (1970) argues that such behavior may explain some of the experimental evidence for under-reaction (see also Benjamin et al., 2016, who introduce the term “extreme-belief aversion”). Indeed, a Coarse Bayesian employing the representation in Figure 4b would effectively under-react to signals that strongly support any particular state.

Figure 4c illustrates an updating rule that coincides with Bayes' rule unless the observed signal is sufficiently “surprising.” In this case, the prior strongly supports a particular state and the agent exhibits non-Bayesian behavior only if the signal has a low probability of occurrence in that state. Several studies, such as De Bondt and Thaler (1985), find that updating behavior at such unexpected signals may be inconsistent with Bayes' rule. See also Ortoleva (2012), who develops a model to accommodate this, and related, evidence.

3.1.2 Understanding & Predicting Behavior

What does the Coarse Bayesian model predict, and how might it help explain real behavior? Since the model accommodates a variety of biases, it does not necessarily predict whether one bias prevails over another—additional assumptions are needed to make such comparisons. Nonetheless, the model makes several concrete predictions about non-Bayesian behavior and provides meaningful implications (and explanations) in a variety of settings.

¹²This makes the most sense in the signal-distortion interpretation of the model, where the agent transforms signals before applying Bayes' rule. The underlying signal distortion function d represents the agent's attention, avoiding the “circularity” of having the agent compute the Bayesian posterior of ignored signals—a common criticism of rational inattention models.

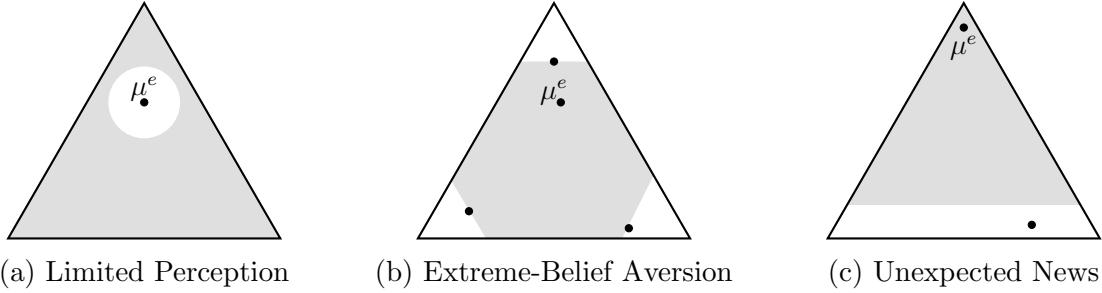


Figure 4: Limited Perception, Extreme-Belief Aversion, and Reactions to Unexpected News. Each point in the shaded regions represents a singleton cell.

General Predictions

1. *(Non)Genericity of “quasiconvex” biases.* While the model accommodates (for example) those who *always* under-react or *always* over-react, it also suggests that such extremes are unusual or non-generic. To see why, consider Figure 3. In Figure 3a, the agent always under-reacts: relative to Bayesian posteriors, realized beliefs are closer to the prior. For this to hold, representative points must sit on the boundaries of their cells. Therefore, global under-reaction is a hairline case; a more typical Coarse Bayesian, depicted in Figure 3d, under-reacts to some signals and over-reacts to others. More broadly, the biases depicted by Figures 3a–3c are “quasiconvex”—the set of beliefs that distort upward is convex, as is the set that distorts downward—and such quasiconvexity is non-generic because (outside of trivial cases) it requires representative points to sit on cell boundaries.

Note, however, that the model predicts *local* uniformity of directional biases: if an individual (say) over-reacts to a signal s , he likely over-reacts to signals near s as well. This holds for regular Coarse Bayesian Representations and is a special case of the stability property discussed next.

2. *Local stability of Bayesian and non-Bayesian behavior.* As illustrated by Figure 4, the model accommodates individuals who apply Bayes’ rule in many (even most) circumstances. In regions where the agent is Bayesian, cells are singletons and posterior beliefs vary smoothly (and non-trivially) with the signal. In contrast, signals yielding non-Bayesian reactions can typically be perturbed without affecting beliefs; this holds because in regular representations, non-singleton cells have full dimension, implying that if a Bayesian posterior belongs to the cell, the cell almost surely contains a neighborhood around that posterior—the only exception is if the posterior sits on the boundary of the cell. Thus, non-Bayesian behavior is locally stable: if a signal s evokes a non-Bayesian response, signals near s typically will, too. Moreover, such signals yield the same (non-Bayesian) posterior.

Interestingly, Bayesian updating also tends to hold locally: if behavior at s is consistent

with Bayes' rule, then (almost surely) so is behavior at nearby signals. This is so because consistency with Bayes' rule at s implies Bayes' rule is violated at nearby signals only if (i) μ^s is the representative point of a non-singleton cell, or (ii) μ^s is the representative of a singleton cell that sits on the boundary of a region of non-singleton cells (like those illustrated in Figure 4). Both scenarios are non-generic and amount to zero-probability events in signal space. Therefore, conditional on being consistent with Bayes' rule at s , the agent is very likely to be consistent with Bayes' rule at nearby signals.

3. Full characterization of predicted non-Bayesian behavior. Proposition 1 provides a complete characterization of the non-Bayesian behavior that Coarse Bayesians must exhibit. In particular, proper Coarse Bayesians satisfy the *negations* of properties (i)–(iii) in Proposition 1. This yields the following predictions:

- (i) There exist signals $s \not\approx t$ such that $\mu^s = \mu^t$. In other words, Coarse Bayesians consider some signals to be equivalent that a Bayesian would not.
- (ii) There exist signals s, t such that $\mu^s \neq \mu^t$ while $\mu^{s+t} = \mu^s$. That is, there exist signals that a Coarse Bayesian distinguishes unless he is uncertain about which one was generated, in which case one of the signals becomes the default.
- (iii) There exist s, t such that $\mu^s = \mu^t$ and $t \not\approx \mu^s/\mu^e$. That is, there are signals s, t such that t brings the agent to μ^s even though t is not perfect evidence of μ^s .¹³ This is a kind of false extrapolation.

Statements (i)–(iii) are, in effect, three different biases that must be exhibited by a Coarse Bayesian and they account for all non-Bayesian behavior generated by the model.

Applications

The following explores some of the ways the model can be used to understand behavior in more applied contexts. The goal is to illustrate various applications of the model without delving too deep into any of them, so the discussion is kept at a fairly high-level.

1. Categorization, Extrapolation and Financial Markets. There is a vast literature on coarse (or categorical) thinking in psychology and economics. Broadly speaking, coarseness involves lumping “similar” situations together and basing analysis (and subsequent decisions) on features of categories. My model provides a unifying framework for settings where categorization involves belief updating; naturally, this affects choice and consumption behavior.

¹³Note that these signals may satisfy $s \approx t$, so this is not equivalent to property (i).

For example, patterns in financial markets and investor behavior can be understood through the lens of Coarse Bayesian updating. Stocks are often classified into categories and “style investors” choose how to invest across categories rather than individual stocks. In a Coarse Bayesian framework, states represent characteristics of stocks, categories are cells and representative points reflect the average within a cell—here, the cells are plausibly exogenous. Crucially, information is processed on the level of categories and this is what drives style investing. Such belief updating naturally entails false extrapolation to all stocks within a category and explains, among other regularities, co-movement within categories (see Sharpe et al., 1992 or Barberis and Shleifer, 2003). More recently, Teeple (2022) employs a version of Coarse Bayesian updating to study behavior involving support and resistance levels; again, the cells (levels) are naturally exogenous.

Stereotyping also fits the model particularly well. Suppose states are multidimensional, with different dimensions capturing characteristics people or places might have. Cells of the representation group signals—that is, noisy information about attributes—the agent might receive about another person and representative points are the “stereotypes” associated with such information. For example, the representation might make the agent believe, erroneously, that an individual with attribute x is likely to have attribute y . The exposure an individual has to different people naturally affects how the individual groups or categorizes them; consequently, we should expect the partition to be finer (coarser) in regions the individual encounters more (less) frequently. Thus, a basic version of the model, with minimal assumptions on how cells and representative points are formed, can explain patterns in stereotyping; for example, stereotypes contain a “kernel of truth” and inferences about out-group members are more error-prone than those about in-group members—see Bordalo et al. (2016) and Bursztyn and Yang (2022).

These and many other phenomena involving coarseness (polarization, inertial behavior, etc) are widely studied but typically involve stylized models and specific assumptions about how agents assess “similarity” to form categories. My framework allows such assessments to be fully subjective and shows that they can be identified from behavior. The tools developed here are based on standard primitives and therefore can aid the analysis of such applications; for example, section 4 establishes that more information, or finer categories, need not be advantageous, and characterizes under what circumstances they actually are.

2. Discourse and Persuasion. While individual behavior is the focus of this paper, there are interesting ways the model can be used to shed light on interactive phenomena, particularly those involving communication or persuasion. Much of our discourse involves not just information provision but arguments about how information should be interpreted and whether

it justifies switching positions on some issue. Do recent data suggest a recession is coming? What did the policymaker really mean when they said they were exploring options? Can existing models adequately explain our observations or do we need a new theory?

A natural way to apply the model to such settings is to allow players (“persuaders”) to influence its parameters. For example, a juror may wish to persuade another that the evidence proves guilt beyond reasonable doubt. This is not a matter of acquiring new information (admissible evidence has already been presented in court) but of arguing about where the cutoff should be between two competing theories (innocence or guilt); effectively, the jurors argue about what “reasonable doubt” means and, thereby, how to evaluate available information. In other contexts, persuaders might wish to influence others’ actions by proposing new theories (representative points) to explain or “frame” available information; a carefully-constructed set of theories can make the receiver adopt beliefs and take actions that are beneficial to the persuader.¹⁴ Again, this kind of persuasion is not about providing new information but influencing the way available information is perceived.

Viewed this way, the model provides a lens through which patterns in discourse can be better understood. A given Coarse Bayesian Representation captures the heuristics and biases an agent employs when processing information. This affects a persuader’s incentives for information provision but also for influencing those heuristics—that is, challenging the set of theories under consideration or the standard of proof for switching between them. Tactics like false dichotomies (arguing one of only two possible theories must be correct) or straw-man arguments (misrepresenting a theory in order to more easily refute it) are but two common examples that fit the model well and are difficult to explain without a notion of coarse cognition.

3.1.3 Testing & Identifying Coarse Cognition

The results of this paper can help inform the design of experiments on non-Bayesian updating. Below, I discuss key criteria for testing the model in the lab and how one might separate Coarse Bayesian updating from competing explanations for non-Bayesian behavior.

1. *Theorem 1 and Proposition 1 as guides for experiments.* As a full axiomatic characterization of Coarse Bayesian updating, Theorem 1 provides a recipe for an ideal experiment: Axioms 1–3 describe patterns in updating behavior that must be satisfied by any Coarse Bayesian and thereby provide a guide for eliciting comparisons in the lab. *Homogeneity*, for example, indicates that updating behavior is invariant to scalar transformations of signals,

¹⁴See Schwartzstein and Sunderam (2021) for an approach where the theories proposed by the persuader are evaluated by the receiver according to a maximum-likelihood criterion.

so a good experimental test of the model would expose subjects to signals of the form s and λs ; if a subject responds differently to λs than to s , they are not Coarse Bayesian. Similarly, *Cognizance* and *Confirmation* provide testable properties of behavior that can falsify the model by eliciting the right kinds of comparisons. This is generally true of any axiomatic characterization provided the axioms are falsifiable, as is the case here.

As explained in section 3.1.2, Proposition 1 identifies patterns capturing *proper* Coarse Bayesian behavior. Specifically, the negations of properties (i)–(iii) in Proposition 1 amount to three different biases that must be exhibited by a proper Coarse Bayesian. Thus, documenting such behavior helps support (not refute) the hypothesis that coarse cognition is at work and can be used to distinguish Coarse Bayesian updating from alternative explanations for non-Bayesian updating.

2. Exploring the signal space. A key implication of the model, discussed above, is local stability of both Bayesian and non-Bayesian behavior. Intuitively, stability means that if there is a non-Bayesian response to some signal, then nearby signals tend to generate the same response. If instead behavior is consistent with Bayes' rule, then it is likely consistent at nearby signals as well. This yields two implications for experimental design. First, it is important to consider signals near an original signal that was tested. If, for example, a subject violates Bayes' rule at s , behavior at signals near s indicates whether Coarse Bayesian behavior is in play; if it is, then nearby signals yield the same posterior μ^s as s . Second, if behavior at s is inconsistent with Bayes' rule, one should also test signals that, under Bayes' rule, would lead the subject to beliefs μ^s ; this helps establish whether μ^s is, in fact, the representative of some cell in a Coarse Bayesian representation.

More generally, the implications of the model suggest it is important to explore the signal space. The fact that over/under-reaction are likely to hold locally but not globally in the model (see section 3.1.2) suggests, for example, that one should test a broad set of signals, not just those confined to a particular region. The same holds for Bayesian behavior: the model predicts that adherence to Bayes' rule, when it occurs, tends to hold at nearby signals as well. Consequently, a broader range of signals must be considered. The axioms, as well as Proposition 1, can help guide the selection of signals to test.

3. Qualitative experimental data. Experiments sometimes elicit “soft” data in addition to standard measurements like action choices. For example, experiments on level-k strategic reasoning might ask subjects to explain their action choices; such descriptions, while difficult to analyze quantitatively, can provide insight about subjects' reasoning processes and help establish whether they employ a level-k heuristic.

The Coarse Bayesian framework lends itself quite naturally to such data: its leading interpretations involve boundedly-rational thinking procedures or heuristics that subjects might knowingly engage in. For example, if asked to explain their assessments, Coarse Bayesians following the competing-theories heuristic might explicitly describe how they narrowed down the possibilities and decided which one best fits the data, while those adhering to a limited-computation heuristic might describe how they arrived at their approximation.

Going beyond such descriptions, an analyst might ask subjects more concrete questions targeting the various interpretations of, and mechanisms for, Coarse Bayesian behavior. Questions like “what would it take to change your mind?” probe a subject’s standard of proof for switching between competing theories, while eliciting ranges of beliefs or asking for a 90% confidence interval can help tease out limited-computation heuristics. Asking subjects to classify their beliefs (eg, “high/medium/low” likelihood of a given state) can capture categorical thinking, and asking whether they incorporated a piece of information into their analysis or how similar they consider different pieces of information to be can capture signal distortion. These approaches provide yet another way to test the model or weigh it against competing explanations.

3.2 Dynamics

This section examines some basic dynamic properties of the model. Suppose an agent observes a sequence of signals $\vec{s} = (s^1, \dots, s^n)$, where s^t is the signal generated in period t . How do properties of \vec{s} affect the agent’s final belief? Must beliefs converge to the truth?

For standard Bayesians, terminal beliefs do not depend on how signals are pooled or ordered. For example, consider a sequence $\vec{s} = (s^1, s^2, s^3)$. The terminal Bayesian belief is $B(\mu^e | s^1 s^2 s^3)$ regardless of whether the sequence is rearranged (eg. (s^2, s^1, s^3)), pooled differently (eg. $(s^1, s^2 s^3)$), or both.¹⁵ Another feature of Bayesian updating is that, for sufficiently informative structures σ , repeated draws of signals from σ make beliefs converge to the truth (a point mass δ_ω on the true state ω). More precisely, suppose the true state is ω and that for every n , s^n is an independent draw from σ (if $t \in \sigma$, then $s^n = t$ with probability t_ω). For a Bayesian, the sequence $(s^n)_{n=1}^\infty$ induces a sequence $(B^n)_{n=1}^\infty$ of beliefs $B^n = B(\mu^e | s^1 s^2 \dots s^n)$ such that $B^n \rightarrow \delta_\omega$ almost surely, provided σ is sufficiently informative.¹⁶

Under non-Bayesian updating, including Coarse Bayesian updating, dynamics are more nuanced. For example, the terminal belief of an agent who incorporates the full history of

¹⁵See Cripps (2018) for a general analysis of updating rules that are invariant to how an agent partitions histories of signals.

¹⁶For example, the uninformative structure $\sigma = e$ yields $B^n = \mu^e$ for all n , so that beliefs converge to μ^e instead of δ_ω . For beliefs to converge to the truth, the distribution over signals $s \in \sigma$ for state ω (that is, the row of matrix σ corresponding to state ω) must differ from that of other states.

signal realizations typically differs from that of one who performs signal-by-signal updating. Similarly, matters of belief convergence depend not only on σ , but on how the (static) non-Bayesian updating rule is extended to a dynamic updating rule. Fortunately, Coarse Bayesian updating yields fairly simple results.

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

Definition 2. 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 $\vec{s} = (s^1, \dots, s^n)$ of length $n \geq 2$ and all $k < n$, $\mu^{\vec{s}} = \mu^{(s^1, \dots, s^{k-1}, s^k s^{k+1}, s^{k+2}, \dots, s^n)}$.

Definition 2 formalizes two different notions of history independence. Under invariance to signal ordering, any history \vec{s} can be reordered without affecting the final belief.¹⁷ 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.

Consider first the following dynamic extension of a Coarse Bayesian updating rule:

Definition 3. A dynamic updating rule μ is a **Pooled Coarse Bayesian** updating rule if either of the following equivalent conditions hold:

- (i) There is a Coarse Bayesian Representation $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ such that, for all histories (s^1, \dots, s^n) , $\mu^{(s^1, \dots, s^n)} = \mu^P$ where $B(\mu^e | s^1 s^2 \dots s^n) \in P \in \mathcal{P}$.
- (ii) There is a Signal Distortion Representation $d : S \rightarrow S$ such that, for all histories (s^1, \dots, s^n) , $\mu^s = B(\mu^e | d(s^1 s^2 \dots s^n))$.

A Pooled Coarse Bayesian updating rule works by applying, at every n , the full history of signals up to that point. The pooled signal $s^1 s^2 \dots s^n$ represents the joint likelihood of having

¹⁷Rabin and Schrag (1999) analyze a model of history-dependent updating where, in each period, information is distorted to support the agent's current belief. Such procedures are not invariant to signal ordering.

observed the sequence, and these likelihoods are applied either to the Coarse Bayesian Representation $\langle \mathcal{P}, \mu^P \rangle$ or its associated Signal Distortion Representation d . Naturally, Pooled Coarse Bayesian updating rules are invariant to signal pooling and, hence, signal ordering.

To study belief convergence, an additional definition is needed. A Coarse Bayesian Representation $\langle \mathcal{P}, \mu^P \rangle$ is **stable at ω** if there exists $P \in \mathcal{P}$ and $\varepsilon > 0$ such that the ε -ball $\{\hat{\mu} \in \Delta : \|\hat{\mu} - \delta_\omega\| < \varepsilon\}$ around δ_ω is contained in P . The next result summarizes the dynamic properties of Pooled Coarse Bayesian updating rules.

Proposition 2. *Pooled Coarse Bayesian updating rules are invariant to signal ordering and pooling. If $\langle \mathcal{P}, \mu^P \rangle$ is stable at ω , $(s^n)_{n=1}^\infty$ is the stochastic sequence generated by σ in state ω (that is, $s^n = t \in \sigma$ with probability t_ω), and $B(\mu^e | s^1 \dots s^n) \xrightarrow{a.s.} \delta_\omega$, then $\mu^{(s^1, \dots, s^n)} \xrightarrow{a.s.} \mu^P$, where $\delta_\omega \in P \in \mathcal{P}$.*

Proposition 2 states that if $\langle \mathcal{P}, \mu^P \rangle$ is stable at ω and σ is sufficiently informative for Bayesian beliefs to converge to δ_ω , then Pooled Coarse Bayesian beliefs converge to the representative μ^P of the cell P containing δ_ω . Thus, Pooled Coarse Bayesian beliefs converge whenever Bayesian beliefs do, but not necessarily to the point δ_ω .

Next, consider the following two types of signal-by-signal updating:

Definition 4. A dynamic updating rule μ is:

- (i) A **Sequential Coarse Bayesian** updating rule if there is a Coarse Bayesian Representation $\langle \mathcal{P}, \mu^P \rangle$ 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 **Sequential Signal Distortion** updating rule if there is a Signal Distortion Representation $d : S^0 \rightarrow S^0$ for histories of length 1 such that, for every history (s^1, \dots, s^n) of length $n \geq 2$, $\mu^{(s^1, \dots, s^n)} = B(\mu^{(s^1, \dots, s^{n-1})} | d(s^n))$.¹⁸

A Sequential Coarse Bayesian updating rule employs a fixed Coarse Bayesian Representation to perform signal-by-signal updating. Starting at prior μ^e , the agent applies $\langle \mathcal{P}, \mu^P \rangle$ to reach posterior μ^{s^1} after observing s^1 . Then, treating μ^{s^1} as the prior, the agent applies the same representation $\langle \mathcal{P}, \mu^P \rangle$ to reach posterior $\mu^{(s^1, s^2)}$ after observing s^2 , and so on. A Sequential Signal Distortion rule follows a similar procedure, substituting d for $\langle \mathcal{P}, \mu^P \rangle$. Thus, sequential rules apply to agents who have imperfect memory and rely on current beliefs as summary statistics of the history.

¹⁸Restricting d to take values in S^0 ensures that $B(\mu^{(s^1, \dots, s^{n-1})} | d(s^n))$ is well defined at all possible histories.

Proposition 3. *Let μ^e have full support. Then:*

- (i) *Sequential Signal Distortion rules are invariant to signal ordering but not necessarily to signal pooling.*
- (ii) *Sequential Coarse Bayesian updating rules need not be invariant to signal ordering nor to signal pooling. If there are full-support representatives $\mu^P \neq \mu^{P'}$ and a signal s^* such that both $B(\mu^P|s^*) \in P$ and $B(\mu^{P'}|s^*) \in P'$, then the updating rule is not invariant to signal ordering.*

Proposition 3 establishes that the path-dependence properties of sequential updating rules depend on how the static rule is extended to a dynamic rule: Sequential Signal Distortion rules are invariant to signal ordering, but Sequential Coarse Bayesian rules need not satisfy either type of path-independence. The requirements specified by the second part of (ii) are satisfied by many Coarse Bayesian rules; such rules fail to be invariant to signal ordering and, therefore, fail to be invariant to signal pooling as well. Intuitively, these differences stem from the fact that fixing d while updating beliefs signal-by-signal effectively yields different Coarse Bayesian Representations at different histories, making Sequential Signal Distortion rules quite different from Sequential Coarse Bayesian rules.

The distinction between Sequential Coarse Bayesian and Signal Distortion rules also has implications for belief convergence. In general, sequences of beliefs induced by Sequential Coarse Bayesian rules need not converge to the true state, or even to converge at all. Sequential Signal Distortion rules, however, do induce belief convergence, though not necessarily to the true state:

Proposition 4. *Suppose μ is a Sequential Signal Distortion rule with distortion function d . Fix $\sigma = [t^1, \dots, t^J]$ and $\omega \in \Omega$. Let $(s^n)_{n=1}^\infty$ denote a sequence of random vectors $s^n \in \sigma$ independently and identically distributed by σ in state ω (for all n , $s^n = t^j \in \sigma$ with probability t_ω^j). Let*

$$t^* = d(t^1)^{t_\omega^1} d(t^2)^{t_\omega^2} \dots d(t^J)^{t_\omega^J} := \left(d(t^1)_{\omega'}^{t_\omega^1} d(t^2)_{\omega'}^{t_\omega^2} \dots d(t^J)_{\omega'}^{t_\omega^J} \right)_{\omega' \in \Omega}. \quad (2)$$

Then $\mu^{(s^1, \dots, s^n)} \rightarrow B(\mu^e|t_{E^})$ almost surely, where $E^* = \{\omega' \in \Omega : t_{\omega'}^* \geq t_{\omega''}^* \ \forall \omega'' \in \Omega\}$ and $t_{E^*} = 1_{[\omega' \in E^*]} \in S$ is the indicator vector for E^* .*

Proposition 4 states that, in the limit, Sequential Signal Distortion narrows the set of possible states down to $E^* = \arg\max_{\omega'} t_{\omega'}^*$, where t^* is the “average” distorted signal generated by σ in state ω . For standard Bayesians, $E^* = \{\omega\}$ provided σ is sufficiently informative. As the next example illustrates, however, E^* need not contain the true state; thus, although

beliefs converge, they need not converge to the true state.

Example 1. Consider a two-state setting. Let

$$d(s) = \begin{cases} e & \text{if } \frac{s_2}{s_1} \leq 3 \\ \left(\frac{1}{4}, 1\right) & \text{otherwise} \end{cases}$$

and $\sigma = [s, t]$ where $s = (\frac{1}{5}, \frac{4}{5})$ and $t = (\frac{4}{5}, \frac{1}{5})$. Then $d(s) = (\frac{1}{4}, 1)$ and $d(t) = e$, so that in state 1 we have $t^* := d(s)^{s_1} d(t)^{t_1} = ((\frac{1}{4})^{1/5}, 1)$. Since $t_1^* < t_2^*$, beliefs converge to state 2.

4 The Value of Information

Assessing the value of information is a fundamental part of decision making in many economic models. In this section, I study the Coarse Bayesian value of information, including its relationship to the Bayesian value of information, the Blackwell (1951) ordering, and notions of cognitive sophistication and bias.

Throughout this section, μ denotes an updating rule with Coarse Bayesian Representation $\langle \mathcal{P}, \mu^P \rangle$. Let \mathcal{A} denote the set of all nonempty, compact subsets of \mathbb{R}^Ω . 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. Action $x \in A$ yields payoff x_ω in state ω . For each $A \in \mathcal{A}$ and $s \in S$, let $c^s(A) := \operatorname{argmax}_{x \in A} x \cdot \mu^s$ denote the actions in A that maximize expected utility at beliefs μ^s .

Definition 5. Let $A \in \mathcal{A}$.

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

$$V^A(\sigma) := \max \sum_{\omega} \mu_{\omega}^e \sum_{s \in \sigma} s_{\omega} x_{\omega}^s \quad \text{subject to } x^s \in c^s(A). \quad (3)$$

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

$$\bar{V}^A(\sigma) := \max \sum_{\omega} \mu_{\omega}^e \sum_{s \in \sigma} s_{\omega} x_{\omega}^s \quad \text{subject to } x^s \in \operatorname{argmax}_{x \in A} x \cdot \frac{s \mu^e}{s \cdot \mu^e}. \quad (4)$$

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_{ω}^e to the average payoff in state ω given that signals are generated by σ . Consistent with the Cognizance axiom, the agent correctly forecasts his own signal-contingent beliefs and, hence,

signal-contingent choices. Equation (4) expresses a similar formula for a Bayesian agent: 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

$$c^{\hat{\mu}}(A) := \operatorname{argmax}_{x \in A} x \cdot \mu^P \text{ subject to } \hat{\mu} \in P \quad \text{and} \quad v^A(\hat{\mu}) := \max_{x \in c^{\hat{\mu}}(A)} x \cdot \hat{\mu}.$$

That is, $c^{\hat{\mu}}(A)$ consists of the actions in A that maximize expected utility for the Coarse Bayesian if the *Bayesian* posterior is $\hat{\mu}$. Similarly, $v^A(\hat{\mu})$ represents expected utility at A conditional on Bayesian posterior $\hat{\mu}$. 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

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

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

Proposition 5 states 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 posterior $\hat{\mu}$, so that the distribution of Bayesian posteriors can be used to calculate expected utility. This also facilitates comparisons between Bayesian and Coarse Bayesian payoffs (see Figure 5); clearly, $v^A(\hat{\mu}) \leq \bar{v}^A(\hat{\mu})$ for all $\hat{\mu}$ and, hence, $V^A(\sigma) \leq \bar{V}^A(\sigma)$ for all σ —the Bayesian always does better. Intuitively, Proposition 5 holds because a Coarse Bayesian updating rule is Homogeneous and, hence, a function of the Bayesian posterior;¹⁹ I omit the straightforward proof.

4.1 The Blackwell Ordering

This section 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). An experiment σ' is a **garbling** of σ if there is a matrix M with entries in $[0, 1]$ such that every row is a probability distribution and $\sigma' = \sigma M$.

¹⁹This is the fundamental assumption of de Clippel and Zhang (2022), who study persuasion with non-Bayesian agents. A similar result appears in Galperti (2019).

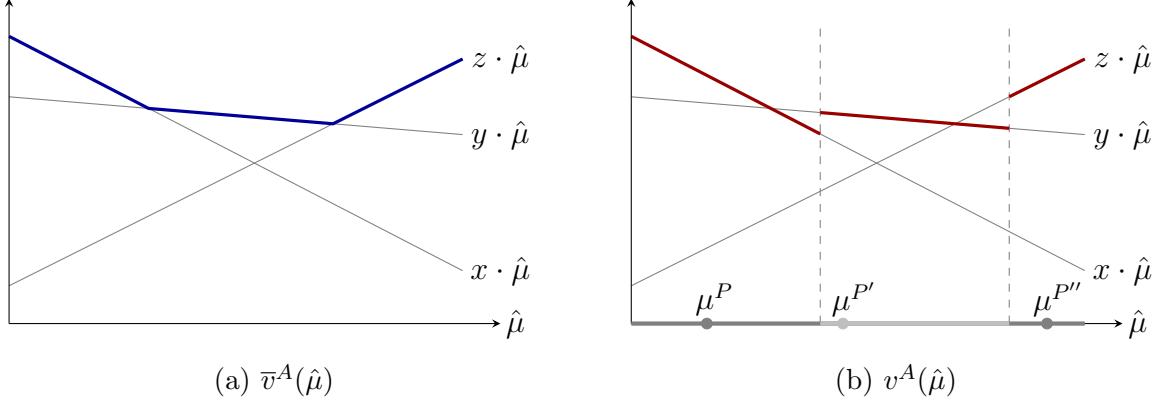


Figure 5: Bayesian vs. Coarse Bayesian value of information for $A = \{x, y, z\}$.

For the purposes of this paper, \sqsupseteq is defined by: $\sigma \sqsupseteq \sigma'$ 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's value of information satisfies the Blackwell ordering in all menus A —in fact, $\sigma \sqsupseteq \sigma'$ if and only if $\bar{V}^A(\sigma) \geq \bar{V}^A(\sigma')$ for all $A \in \mathcal{A}$. For Coarse Bayesians, this need not be 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 . Formally, $b^s(A) := \{x \in A : x \cdot \frac{s\mu^e}{s \cdot \mu^e} \geq y \cdot \frac{s\mu^e}{s \cdot \mu^e} \forall y \in X\}$. Let $c(A) = \bigcup_{s \in S} c^s(A)$ and $b(A) = \bigcup_{s \in S} b^s(A)$. That is, $c(A)$ is the set of actions in A that are chosen by the Coarse Bayesian—and $b(A)$ the set of actions chosen by the Bayesian—for at least one s . Observe that, by Confirmation, $c(A) \subseteq b(A)$.

Proposition 6. *Let $\langle \mathcal{P}, \mu^P \rangle$ be a regular Coarse Bayesian Representation and $A \in \mathcal{A}$. The following are equivalent:*

- (i) V^A satisfies the Blackwell ordering.
- (ii) v^A is convex.
- (iii) $c^s(A) \cap b^s(c(A)) \neq \emptyset$ for all s .

Proposition 6 characterizes, for regular Coarse Bayesians, the class of menus A such that V^A satisfies the Blackwell ordering.²⁰ The key property is (iii), asserting that Coarse Bayesian choices from A agree with Bayesian choices from the menu $c(A) \subseteq A$ (the submenu

²⁰The regularity requirement only serves to establish (i) \Rightarrow (iii). In particular, the implication (iii) \Rightarrow (i) holds for all Coarse Bayesian Representations, as does the equivalence of (i) and (ii). The implication (ii) \Rightarrow (i) is part of Blackwell's characterization, but the converse implication is not, and relies on the assumption that μ^e has full support (see Lemma 1 in the appendix).

of actions that are actually chosen at some signal realization). When (iii) is satisfied, Coarse Bayesian behavior at A coincides with Bayesian behavior at $c(A)$, making $v^A = \bar{v}^{c(A)}$ convex and $V^A = \bar{V}^{c(A)}$ satisfy the Blackwell ordering. Since (iii) is a rather strong requirement, Blackwell violations are a common occurrence.²¹

Example 2. Some non-Bayesians satisfy the Blackwell ordering in all menus. Suppose $N = 2$, so that Δ is represented by the interval $[0, 1]$ of values $\hat{\mu}_1$. First, consider $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ where \mathcal{P} contains two cells: $P = \{0\}$ and $P' = (0, 1]$. Assume $\mu^{P'} < 1$. Then, for every A , v^A is convex; this implies V^A satisfies the Blackwell ordering, even though choices generated by $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ violate condition (iii) of Proposition 6 in some menus. Next, let $\langle \mathcal{Q}, \mu^{\mathcal{Q}} \rangle$ consist of a cell $Q = [0, \mu^*]$ where $0 < \mu^* < 1$ and, for every $\hat{\mu} > \mu^*$, a singleton cell $\{\hat{\mu}\}$. Let $\mu^Q = \mu^*$. Choices generated by $\langle \mathcal{Q}, \mu^{\mathcal{Q}} \rangle$ satisfy condition (iii) of Proposition 6 for all A ; this implies the corresponding value of information function satisfies the Blackwell ordering in all menus, even though $\langle \mathcal{Q}, \mu^{\mathcal{Q}} \rangle$ violates the regularity assumption (see footnote 20).

Example 2 shows it is possible for non-Bayesian representations to generate functions V^A satisfying the Blackwell ordering for all A with or without condition (iii) of Proposition 6. Such representations are quite rare, however, in that small perturbations of the cells or representative points guarantee that V^A violates both the Blackwell ordering and condition (iii) for some A . Intuitively, violations of the Blackwell ordering arise through discontinuities in v^A because such discontinuities, except possibly on the boundary of Δ , make v^A non-convex. Most non-Bayesian representations have the property that any violation of (iii) introduces a non-convexity in v^A for some A because the gap between Bayesian and non-Bayesian choices creates points of discontinuity. For regular representations, violations of (iii) are both necessary and sufficient for the existence of such discontinuities.

While it is perhaps not too surprising that non-Bayesian updating can generate violations of the Blackwell ordering, it turns out that, for Coarse Bayesians, the connection to the Blackwell ordering runs much deeper:

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, such that $\mu^e = \dot{\mu}^e$. The following are equivalent:*

$$(i) \quad \langle \mathcal{P}, \mu^{\mathcal{P}} \rangle = \langle \mathcal{Q}, \dot{\mu}^{\mathcal{Q}} \rangle.$$

$$(ii) \quad \text{For all } \sigma \sqsupseteq \sigma' \text{ and } A \in \mathcal{A}, \quad V^A(\sigma) \geq V^A(\sigma') \Leftrightarrow \dot{V}^A(\sigma) \geq \dot{V}^A(\sigma').$$

²¹See Whitmeyer (2023) for a recent characterization of updating rules for which the associated value functions satisfy the Blackwell ordering in all menus.

Proposition 7 states that, for a regular Coarse Bayesian, the parameters $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ are pinned down by the agent's ranking of Blackwell-comparable experiments. Thus, by observing when the agent benefits (or expects to benefit) from a Blackwell improvement, one can uniquely identify the parameters of the representation. A key takeaway, then, is not just that Coarse Bayesians exhibit violations of the Blackwell ordering, but that they do so in a way that fully reveals their individual updating behavior.

4.2 Measures of Sophistication and Bias

In this section, I explore different notions of cognitive ability and how they relate to a Coarse Bayesian's value of information. In addition to providing basic comparative static results for the model, the findings are potentially relevant for endogenizing non-Bayesian updating rules and, hence, developing theories of where they "come from" (see the discussion at the end of the section). For any updating rule μ and signal $s \in S$, let

$$D_\mu(s) := \left\| \frac{s\mu^e}{\|s\mu^e\|} - \frac{\mu^s}{\|\mu^s\|} \right\|.$$

This is the Euclidean distance between μ^s and the Bayesian posterior $\frac{s\mu^e}{s \cdot \mu^e}$ after normalizing each vector to length 1. Thus, $D_\mu(s)$ provides a measure of how distorted the agent's beliefs are at signal s .

Definition 6. Suppose μ and $\dot{\mu}$ have 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:

- (i) $\dot{\mu}$ is **more sophisticated** than μ if every $P \in \mathcal{P}$ is a union of cells in \mathcal{Q} .
- (ii) $\dot{\mu}$ is **less biased** than μ if $D_{\dot{\mu}}(s) \leq D_\mu(s)$ for all $s \in S$.

Definition 6 provides two comparative notions of cognitive ability. Part (i) states that a Coarse Bayesian is more sophisticated if he employs a finer partition, while part (ii) states the agent is less biased if, for every signal, posterior beliefs are closer to the Bayesian posterior. Each ordering captures some aspect of what it means to be "more Bayesian," but the two concepts are quite different: higher sophistication entails higher responsiveness to information, while lower bias entails less skewness in the updating rule (see Figure 6).

The goal of this section is to characterize these orderings in terms of the welfare of the agent. A natural conjecture, for example, is that a more sophisticated agent always enjoys a higher expected utility than a less sophisticated agent, or benefits from more information whenever a less sophisticated agent does. As the next example shows, this conjecture is false.



Figure 6: An illustration of the bias ordering. The two updating rules employ the same pair feasible beliefs, but rule (b) is less biased than rule (a) because it exhibits smaller distortions away from Bayesian posteriors; this makes the cutoff between cells more “centered.”

Example 3. Consider a two-state setting, so that $\Delta = [0, 1]$. Let $\mathcal{P} = \{P, P'\}$ where $P = \{0\}$ and $P' = (0, 1]$ and $\mathcal{Q} = \{Q, Q', Q''\}$ where $Q = \{0\}$, $Q' = [\frac{3}{4}, 1]$, and $Q'' = (0, \frac{3}{4})$. Finally, let $\mu^P = \dot{\mu}^Q = 0$, $\mu^e = \mu^{P'} = \frac{4}{5} = \dot{\mu}^{Q'} = \dot{\mu}^e$, and $\dot{\mu}^{Q''} = \frac{1}{3}$. Clearly, $\dot{\mu}$ is more sophisticated than μ . Let $A = \{x, y\}$ where $x = (1, 0)$ and $y = (0, 1)$. Then

$$v^A(\hat{\mu}_1) = \begin{cases} 1 & \text{if } \hat{\mu}_1 = 0 \\ \hat{\mu}_1 & \text{otherwise} \end{cases} \quad \text{and} \quad \dot{v}^A(\hat{\mu}_1) = \begin{cases} 1 - \hat{\mu}_1 & \text{if } \hat{\mu}_1 < \frac{3}{4} \\ \hat{\mu}_1 & \text{otherwise} \end{cases},$$

so that $\dot{v}^A(\hat{\mu}_1) < v^A(\hat{\mu}_1)$ for $\frac{1}{2} < \hat{\mu}_1 < \frac{3}{4}$. Thus, $\dot{V}^A(\sigma) < V^A(\sigma)$ for some σ (for example, any σ such that $\tau^\sigma(\frac{2}{3}) = \frac{3}{5}$ and $\tau^\sigma(1) = \frac{2}{5}$). Moreover, v^A is convex but \dot{v}^A is not; thus, V^A satisfies the Blackwell ordering but \dot{V}^A does not.

In general, greater sophistication need not improve welfare because it does not rule out the possibility of wider gaps between Bayesian and Coarse Bayesian choices at some menu-signal pairs. Similarly, lower bias need not imply welfare improvements. At the end of this section, I return to this question and examine the conditions under which one Coarse Bayesian is better off than another in all decision problems (Proposition 10).

To characterize the sophistication ordering, an additional definition is required. Given $\langle \mathcal{P}, \mu^P \rangle$, a pair (A, σ) is μ^P -**decisive** if $c^s(A)$ is a singleton for all $s \in \sigma$; that is, if every posterior μ^P induced by σ yields a unique optimal action in A . For any $\sigma, \sigma' \in \mathcal{E}$, $V(\sigma) = V(\sigma')$ μ^P -**decisively** if $V^A(\sigma) = V^A(\sigma')$ for all A such that (A, σ) and (A, σ') are μ^P -decisive.

Proposition 8. Suppose $\langle \mathcal{P}, \mu^P \rangle$ and $\langle \mathcal{Q}, \dot{\mu}^Q \rangle$ are regular Coarse Bayesian Representations of μ and $\dot{\mu}$, respectively, and that $\mu^e = \dot{\mu}^e$. The following are equivalent:

- (i) $\dot{\mu}$ is more sophisticated than μ .
- (ii) If $\sigma, \sigma' \in \mathcal{E}$ and $\dot{V}(\sigma) = \dot{V}(\sigma')$ $\dot{\mu}^Q$ -decisively, then $V(\sigma) = V(\sigma')$ μ^P -decisively.

This result states that for regular Coarse Bayesians, greater sophistication means welfare is more responsive to information: as sophistication increases, fewer pairs σ, σ' yield identical

ex-ante expected utility for (almost) all menus A . The proof of Proposition 8 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. More-responsive welfare, of course, does not imply greater welfare.

The characterization of the bias ordering does not involve the responsiveness of welfare, but rather a comparison to that of a Bayesian. For each $s \in S$ and $A \in \mathcal{A}$, let $\bar{V}^A(s) := \bar{v}^A(B(\mu^e|s))$ and $V^A(s) := v^A(B(\mu^e|s))$ denote the Bayesian and Coarse Bayesian payoffs at menu A conditional on signal s . Let

$$L_\mu(s) := \sup_{A \in \mathcal{A}^*} \bar{V}^A(s) - V^A(s)$$

where \mathcal{A}^* denotes the set of menus A such that $\|x\| \leq 1$ for all $x \in A$. Intuitively, $L_\mu(s)$ is the maximum loss, relative to a Bayesian, that the Coarse Bayesian can incur under any decision problem A .²² Alternatively, $L_\mu(s)$ may be interpreted as the maximum rate at which a Bayesian agent can “money pump” the Coarse Bayesian agent under public information s . So, if actions x represent bets or gambles, and a Bayesian agent is free to specify a set $A \in \mathcal{A}^*$ after both agents have observed s , then $L_\mu(s)$ is the amount of money the Bayesian can extract from the Coarse Bayesian.²³

Proposition 9. *Suppose μ and $\dot{\mu}$ are Coarse Bayesian and $\mu^e = \dot{\mu}^e$. Then $L_{\dot{\mu}}(s) \leq L_\mu(s)$ if and only if $D_{\dot{\mu}}(s) \leq D_\mu(s)$. Thus, $\dot{\mu}$ is less biased than μ if and only if $L_{\dot{\mu}}(s) \leq L_\mu(s)$ for all $s \in S$.*

Proposition 9 establishes that $\dot{\mu}$ is less biased than μ if and only if $\dot{\mu}$ is less exploitable than μ : worst-case losses for $\dot{\mu}$, relative to a Bayesian, are smaller than those for μ .

As indicated above, neither greater sophistication nor lower bias guarantee higher payoffs in all decision problems. The next result establishes that, under mild regularity conditions, a particular refinement of these orderings is needed to improve payoffs in all decision problems.

Proposition 10. *Suppose $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ and $\langle \mathcal{Q}, \dot{\mu}^{\mathcal{Q}} \rangle$ are Coarse Bayesian Representations of μ and $\dot{\mu}$ such that $\mu^e = \dot{\mu}^e$ and non-singleton cells are regular. The following are equivalent:*

(i) $\dot{v}^A(\hat{\mu}) \geq v^A(\hat{\mu})$ for all $A \in \mathcal{A}$ and $\hat{\mu} \in \Delta$.

(ii) $\dot{\mu}$ is less biased, more sophisticated and, for every $\dot{\mu}^Q \in \dot{\mu}^{\mathcal{Q}} \setminus \mu^{\mathcal{P}}$, the cell Q is a singleton.

²²The restriction to normalized menus $A \in \mathcal{A}^*$ is needed because $V^{\lambda A} = \lambda V^A$ for all $\lambda > 0$.

²³Indeed, as shown in the appendix, one may restrict attention to menus of the form $A = \{0, x\}$ where, conditional on s , the Bayesian prefers the safe option 0 but the Coarse Bayesian strictly prefers x . On average, the Bayesian profits by $|x \cdot B(\mu^e|s)|$.

Proposition 10 states that payoffs increase at all menu-signal pairs if and only if the agent becomes more sophisticated and all “new” feasible posteriors $\dot{\mu}^Q$ represent singleton cells Q . This means the agent becomes perfectly Bayesian on a subset of Δ , blocking new or different distortions that yield lower payoffs in some menu-signal pair. It follows immediately that the agent is less biased and that $\dot{V}^A(\sigma) \geq V^A(\sigma)$ for all A and σ .

I conclude this section with a brief discussion of how my results might enable various approaches for selecting or endogenizing Coarse Bayesian updating rules. One approach is to solve for an optimal updating rule in a given *environment*—a menu and signaling structure—under some constraint (for example, a fixed number of cells or a cost per additional cell). Pioneered by Wilson (2014) and Brunnermeier and Parker (2005), versions of this approach can provide a theory of where the updating rule “comes from.” A drawback is that an updating rule adapted to one environment may be ill-suited for another. Only the robust ordering given by statement (ii) of Proposition 10 ensures weakly greater payoffs at all menu-signal pairs. So, rather than considering updating rules adapted to specific environments, one might instead endogenize them by selecting rules that are unimprovable (given costs or constraints) under the robust ordering. Alternatively, one might consider the weaker objective of minimizing worst-case losses (Proposition 9). These approaches are suitable if agents are unable to form probabilistic beliefs about their environment and, consequently, seek heuristics robust to such uncertainty. Naturally, different criteria yield different predictions about updating rules; minimization of worst-case losses, for example, leads to representations exhibiting less skewness. Analysis of endogenous updating rules is beyond the scope of this paper, but—as illustrated by the characterizations in this section—the framework of Coarse Bayesian updating provides a natural and tractable setting in which to carry it out.

5 Conclusion

In this paper, I have proposed a simple generalization of Bayes’ rule, *Coarse Bayesian updating*, that can account for a variety of biases and individual heterogeneity in updating behavior. Three axioms—*Homogeneity*, *Cognizance*, and *Confirmation*—fully characterize the model and have the property that strengthening any of them to an if-and-only-if form makes the agent fully Bayesian. Thus, Coarse Bayesian updating may be viewed as a “small” departure from Bayes’ rule, and there is a clear separation between the properties of Bayes’ rule that are satisfied by proper Coarse Bayesians and those that are necessarily violated.

An advantage of the framework is that it employs standard primitives that frequently appear in applications. The use of noisy signals over a state space, for example, allows one to import Coarse Bayesian updating into familiar settings in economics and game theory. I

illustrate this by embedding the model in a standard setting of decision under risk, leading to a close relationship with the Blackwell ordering and comparative notions of cognitive sophistication and bias. I leave further development of such applications to future work.

A Proofs

A.1 Proof of Theorem 2

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 $d : S \rightarrow S$ 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 satisfies properties (i)–(iii) of Theorem 2.

Conversely, suppose μ has a Signal Distortion Representation 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 . By (i) and (ii), 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 (iii), $\mu^P \in P$.

A.2 Proof of Proposition 1

It is straightforward to verify that Bayesian updating satisfies properties (i)–(iii). So, suppose μ has a Coarse Bayesian Representation $\langle \mathcal{P}, \mu^P \rangle$. We show that each of properties (i)–(iii) forces each cell of \mathcal{P} to be a singleton, making the agent Bayesian.

For (i), suppose $\mu^s = \mu^t$ implies $s \approx t$. Let $P \in \mathcal{P}$ and $\hat{\mu}, \hat{\mu}' \in P$. Choose signals s, t such that $B(\mu^e|s) = \hat{\mu}$ and $B(\mu^e|t) = \hat{\mu}'$. Then $\mu^s = \mu^t = \mu^P$, so that $s \approx t$ and, hence, $\hat{\mu} = B(\mu^e|s) = B(\mu^e|t) = \hat{\mu}'$. Thus, every cell $P \in \mathcal{P}$ is a singleton.

For (ii), suppose $\mu^{s+t} = \mu^s$ implies $\mu^s = \mu^t$. Suppose toward a contradiction that \mathcal{P} contains a non-singleton cell P . Since μ is non-constant, there exists $P' \in \mathcal{P}$ such that $\mu^P \neq \mu^{P'}$. Since μ^e has full support, there exist signals \hat{s}, \hat{t} such that $B(\mu^e|\hat{s}) = \mu^P$ and $B(\mu^e|\hat{t}) = \mu^{P'}$; thus, $\mu^{\alpha\hat{s}} = \mu^P$ and $\mu^{\beta\hat{t}} = \mu^{P'}$ for all $\alpha, \beta \in (0, 1)$. By equation (1) in the main text, it follows that if $\alpha\hat{s} + \beta\hat{t} \in S$, then $B(\mu^e|\alpha\hat{s} + \beta\hat{t}) = \frac{\alpha\hat{s}\cdot\mu^e}{(\alpha\hat{s} + \beta\hat{t})\cdot\mu^e}\mu^P + \frac{\beta\hat{t}\cdot\mu^e}{(\alpha\hat{s} + \beta\hat{t})\cdot\mu^e}\mu^{P'}$,

which converges to μ^P as $\beta \rightarrow 0$. By regularity of P , there is an ε -ball $B^\varepsilon \subseteq P$ around μ^P . Thus, for sufficiently small $\alpha, \beta \in (0, 1)$, we have $\mu^{\alpha\hat{s} + \beta\hat{t}} \in S$ and $B(\mu^e | \alpha\hat{s} + \beta\hat{t}) \in B^\varepsilon$; but then $\mu^{\alpha\hat{s} + \beta\hat{t}} = \mu^P = \mu^{\alpha\hat{s}}$ while $\mu^{\beta\hat{t}} = \mu^{P'} \neq \mu^P$, contradicting property (ii).

For (iii), suppose $\mu^t = \mu^s$ implies $t \approx \mu^s/\mu^e$. Consider the case $t = s$. Then $\mu^t = \mu^s$, so $s = t \approx \mu^s/\mu^e$. This implies $\mu^s \approx s\mu^e$, so that $\mu^s = B(\mu^e | s)$.

A.3 Proof of Proposition 2

Suppose μ is a Pooled Coarse Bayesian updating rule induced by a representation $\langle \mathcal{P}, \mu^P \rangle$ that is stable at ω . Let $\{s^n\}_{n=1}^\infty$ denote a sequence of signal realizations from σ such that $B(\mu^e | s^1 \dots s^n) \rightarrow \delta_\omega$, and $\{B^n\}_{n=1}^\infty$ the associated sequence of Pooled Coarse Bayesian beliefs; formally, $B^n := \mu^P$ such that $B(\mu^e | s^1 \dots s^n) \in P \in \mathcal{P}$. By stability at ω , there is an $\varepsilon > 0$ and a cell $P^* \in \mathcal{P}$ such that the ε -ball in Δ around δ_ω is a subset of P^* . Thus, for all n sufficiently large, $B(\mu^e | s^1 \dots s^n) \in P^*$ and, hence, $B^n = \mu^{P^*}$. So, if $B(\mu^e | s^1 \dots s^n) \rightarrow \delta_\omega$ almost surely, then $\mu^{(s^1, \dots, s^n)} = B^n \rightarrow \mu^{P^*}$ almost surely.

A.4 Proof of Proposition 3

For (i), let μ be a Sequential Signal Distortion rule. 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 ordering. However, μ need not be invariant to signal ordering. For example, consider a model with two states and distortion function

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

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 $\mu^{(s,t)} \neq \mu^{st}$.

For (ii), consider a Sequential Coarse Bayesian updating rule satisfying all requirements in the second part of the statement. Since μ^e has full support, there is a signal r such that $B(\mu^e | r) = \mu^P$. Similarly, there is a signal t such that $B(\mu^P | t) = \mu^{P'}$ because μ^P has full support. It follows that $\mu^{(r,t,s^*)} = \mu^P \neq \mu^{P'} = \mu^{(r,s^*,t)}$.

A.5 Proof of Proposition 4

Given a finite sequence $s^1, \dots, s^n \in \sigma = [t^1, \dots, t^J]$ and $1 \leq j \leq J$, let n_j denote the number of signals s^i such that $s^i = t^j$. Then $\mu^{(s^1, \dots, s^n)} = B(\mu^e | r^n)$ where $r^n := d(s^1)d(s^2) \dots d(s^n) =$

$d(t^1)^{n_1}d(t^2)^{n_2}\dots d(t^J)^{n_J}$. Observe that, in state ω , $\frac{n_j}{n} \rightarrow t_\omega^j$ almost surely. Thus,

$$(r^n)^{1/n} := d(t^1)^{n_1/n}d(t^2)^{n_2/n}\dots d(t^J)^{n_J/n} \rightarrow d(t^1)^{t_\omega^1}d(t^2)^{t_\omega^2}\dots d(t^J)^{t_\omega^J} := t^*$$

almost surely. Consider the likelihood ratio $\ell_{\omega',\omega''}^n := \frac{r_{\omega'}^n}{r_{\omega''}^n}$. If $\frac{t_{\omega'}^*}{t_{\omega''}^*} < 1$, then $\ell_{\omega',\omega''}^n \rightarrow 0$ almost surely because $\ell_{\omega',\omega''}^n = \left(\frac{(r_{\omega'}^n)^{1/n}}{(r_{\omega''}^n)^{1/n}}\right)^n$ and $\frac{(r_{\omega'}^n)^{1/n}}{(r_{\omega''}^n)^{1/n}} \rightarrow \frac{t_{\omega'}^*}{t_{\omega''}^*} \in [0, 1)$ almost surely. So, take any $\omega^* \in E^*$. Then, as $n \rightarrow \infty$, we have

$$B(\mu^e|r^n) = B\left(\mu^e \middle| \frac{1}{t_{\omega^*}^*}r^n\right) = \frac{\mu^e \frac{r^n}{t_{\omega^*}^*}}{\mu^e \cdot \frac{r^n}{t_{\omega^*}^*}} \xrightarrow{\text{a.s.}} \frac{\mu^e 1_{[\omega' \in E^*]}}{\mu^e \cdot 1_{[\omega' \in E^*]}} = B(\mu^e|t_{E^*}).$$

A.6 Proof of Proposition 6

Lemma 1. *Let $\varphi : \Delta \rightarrow \mathbb{R}$ and $\Phi : \mathcal{E} \rightarrow \mathbb{R}$ such that $\Phi(\sigma) = \sum_{\hat{\mu}} \varphi(\hat{\mu})\tau^\sigma(\hat{\mu})$. Suppose Φ satisfies the Blackwell ordering: $\sigma \sqsupseteq \sigma'$ implies $\Phi(\sigma) \geq \Phi(\sigma')$. Then φ is convex.*

Proof. Let $\hat{\mu}, \hat{\mu}' \in \Delta$, $\alpha \in (0, 1)$, and $\hat{\mu}^\alpha := \alpha\hat{\mu} + (1 - \alpha)\hat{\mu}'$. Since μ^e has full support, there exists $\hat{\mu}^* \in \Delta$ and $\lambda \in (0, 1]$ such that $\lambda\hat{\mu}^* + (1 - \lambda)\hat{\mu}^\alpha = \mu^e$. Let $\sigma = [s^*, s, s']$ and $\sigma' = [s^*, s + s']$ where $s^* = \lambda\frac{\hat{\mu}^*}{\mu^e}$, $s = (1 - \lambda)\alpha\frac{\hat{\mu}^*}{\mu^e}$, and $s' = (1 - \lambda)(1 - \alpha)\frac{\hat{\mu}'}{\mu^e}$. Clearly, $\sigma \sqsupseteq \sigma'$, so that $\Phi(\sigma) \geq \Phi(\sigma')$. Moreover, $\mu^e \cdot s^* = \lambda$, $\mu^e \cdot s = (1 - \lambda)\alpha$, $\mu^e \cdot s' = (1 - \lambda)(1 - \alpha)$, and $\mu^e \cdot (s + s') = 1 - \lambda$, while $B(\mu^e|s^*) = \hat{\mu}^*$, $B(\mu^e|s) = \hat{\mu}$, $B(\mu^e|s') = \hat{\mu}'$, and $B(\mu^e|s + s') = \hat{\mu}^\alpha$. Thus, $\Phi(\sigma) = \varphi(\hat{\mu}^*)\lambda + \varphi(\hat{\mu})(1 - \lambda)\alpha + \varphi(\hat{\mu}')(1 - \lambda)(1 - \alpha)$ and $\Phi(\sigma') = \varphi(\hat{\mu}^*)\lambda + \varphi(\hat{\mu}^\alpha)(1 - \lambda)$, so that $\Phi(\sigma) \geq \Phi(\sigma')$ yields $\alpha\varphi(\hat{\mu}) + (1 - \alpha)\varphi(\hat{\mu}') \geq \varphi(\hat{\mu}^\alpha)$, as desired. \square

To prove Proposition 6, let $A \in \mathcal{A}$ and observe that (i) \Rightarrow (ii) by Lemma 1 (taking $\varphi = v^A$). The converse implication, (ii) \Rightarrow (i), follows from Blackwell's theorem. To see that (iii) \Rightarrow (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 $V^A(\sigma) = V^{A'}(\sigma) = \bar{V}^{A'}(\sigma)$ for all σ . That is, V^A coincides with the Bayesian value of information in some menu, and therefore satisfies the Blackwell ordering.

Finally, we prove that (i) \Rightarrow (iii). Suppose (iii) is violated; that is, there exists $s \in S$ such that $c^s(A) \cap b^s(c(A)) = \emptyset$. Let $\hat{\mu} = B(\mu^e|s)$. Then there exists $x \in c(A)$ such that $v^A(\hat{\mu}) = x \cdot \hat{\mu} < y \cdot \hat{\mu}$ for all $y \in b^s(c(A))$. Choose any $y \in b^s(c(A))$ and $P \in \mathcal{P}$ such that $y \in c^P(A)$. Let $t \in S$ such that $B(\mu^e|t) = \mu^P$. By regularity, P has full dimension in Δ and μ^P belongs to the interior of P ; therefore, we may assume $B(\mu^e|s + t) \in P$ (if necessary,

scale s and t down by some $\lambda > 0$ sufficiently small). Observe that

$$B(\mu^e|s+t) = \frac{s \cdot \mu^e}{(s+t) \cdot \mu^e} \hat{\mu} + \frac{t \cdot \mu^e}{(s+t) \cdot \mu^e} \mu^P := \hat{\mu}',$$

and that there exists $y' \in c^{\mu^P}(A)$ such that

$$v^A(\hat{\mu}') = y' \cdot \hat{\mu}' = \frac{s \cdot \mu^e}{(s+t) \cdot \mu^e} y' \cdot \hat{\mu} + \frac{t \cdot \mu^e}{(s+t) \cdot \mu^e} y' \cdot \mu^P.$$

In particular, y' maximizes the above expression, so we have $y' \cdot \hat{\mu} \geq y \cdot \hat{\mu}$ and $y' \cdot \mu^P = y \cdot \mu^P$ because $y \in c^{\mu^P}(A)$. Now let $\sigma = [s, t, e - s - t]$ and $\sigma' = [s + t, e - s - t]$. Clearly, $\sigma \sqsupseteq \sigma'$. Let $V^A(e - s - t) := v^A(B(\mu^e|e - s - t))[(e - s - t) \cdot \mu^e]$. Then

$$\begin{aligned} V^A(\sigma') &= v^A(\hat{\mu}')[s + t) \cdot \mu^e] + V^A(e - s - t) \\ &= (y' \cdot \hat{\mu})(s \cdot \mu^e) + (y' \cdot \mu^P)(t \cdot \mu^e) + V^A(e - s - t) \\ &\geq (y \cdot \hat{\mu})(s \cdot \mu^e) + (y \cdot \mu^P)(t \cdot \mu^e) + V^A(e - s - t) \\ &> (x \cdot \hat{\mu})(s \cdot \mu^e) + (y \cdot \mu^P)(t \cdot \mu^e) + V^A(e - s - t) \\ &= V^A(\sigma). \end{aligned}$$

A.7 Proof of Proposition 7

The implication (i) \Rightarrow (ii) is clear; the converse follows immediately from the next lemma.

Lemma 2. *Suppose $\langle \mathcal{P}, \mu^P \rangle$ and $\langle \mathcal{Q}, \dot{\mu}^Q \rangle$ are regular representations of μ and $\dot{\mu}$, respectively, such that $\mu^e = \dot{\mu}^e$. Furthermore, suppose that for all $\sigma \sqsupseteq \sigma'$ and $A \in \mathcal{A}$, $\dot{V}^A(\sigma) \geq \dot{V}^A(\sigma') \Rightarrow V^A(\sigma) \geq V^A(\sigma')$. Then \mathcal{Q} is finer than \mathcal{P} and $\mu^P \subseteq \dot{\mu}^Q$.*

Proof of Lemma 2. The proof is divided into three steps.

Step 1: for every $Q \in \mathcal{Q}$, there is a unique $P \in \mathcal{P}$ such that $\text{int}(Q) \subseteq \text{int}(P)$.

First, observe that for every $Q \in \mathcal{Q}$ there is at least one $P \in \mathcal{P}$ such that $\text{int}(Q) \cap \text{int}(P) \neq \emptyset$; this holds because at least one P intersects the (nonempty, by regularity) set $\text{int}(Q)$, which implies $\text{int}(Q) \cap \text{int}(P) \neq \emptyset$ by regularity of Q and P .

So, suppose toward a contradiction that there exist $Q \in \mathcal{Q}$ and distinct $P, P' \in \mathcal{P}$ such that $\text{int}(Q) \cap \text{int}(P) \neq \emptyset$ and $\text{int}(Q) \cap \text{int}(P') \neq \emptyset$. Then there exist $\hat{\mu}, \hat{\mu}' \in \text{int}(Q)$ such that $\hat{\mu} \in \text{int}(P)$ and $\hat{\mu}' \in \text{int}(P')$. Note that $\hat{\mu} \neq \hat{\mu}'$ since $P \cap P' = \emptyset$. Moreover, we may assume $\mu^P \notin \text{co}\{\hat{\mu}, \hat{\mu}'\}$ since, by regularity, we can replace $\hat{\mu}$ with a point in the interior of $\text{co}\{\mu^P, \hat{\mu}'\} \cap P$ if $\mu^P \in \text{co}\{\hat{\mu}, \hat{\mu}'\}$. Similarly, we may assume $\mu^{P'} \notin \text{co}\{\hat{\mu}, \hat{\mu}'\}$.

Next, we argue that it is without loss to assume that either $\mu^P \notin \text{co}\{\hat{\mu}, \hat{\mu}', \mu^{P'}\}$ or $\mu^{P'} \notin \text{co}\{\hat{\mu}, \hat{\mu}', \mu^P\}$. First, consider the case $N = 2$ (2 states). Since $\mu^P \notin \text{co}\{\hat{\mu}, \hat{\mu}'\}$ and $\mu^{P'} \notin \text{co}\{\hat{\mu}, \hat{\mu}'\}$, it follows immediately that $\mu^P \notin \text{co}\{\hat{\mu}, \hat{\mu}', \mu^{P'}\}$ because otherwise $\mu^P \in \text{co}\{\mu^{P'}, \hat{\mu}'\} \subseteq P'$. Similarly, $\mu^{P'} \notin \text{co}\{\hat{\mu}, \hat{\mu}', \mu^P\}$. Now consider the case $N \geq 3$. By regularity, we may assume that the points $\hat{\mu}$, $\hat{\mu}'$, μ^P , and $\mu^{P'}$ are distinct and not collinear (regularity allows us to perturb the points if necessary). It follows immediately that $\mu^P \notin \text{co}\{\hat{\mu}, \hat{\mu}', \mu^{P'}\}$ or $\mu^{P'} \notin \text{co}\{\hat{\mu}, \hat{\mu}', \mu^P\}$.

Suppose $\mu^P \notin \text{co}\{\hat{\mu}, \hat{\mu}', \mu^{P'}\}$ (the argument for the other case is similar). Then we may strictly separate μ^P and $\text{co}\{\hat{\mu}, \hat{\mu}', \mu^{P'}\}$; in particular, there exists x such that $x \cdot \mu^P < 0$ and $x \cdot \tilde{\mu} > 0$ for $\tilde{\mu} \in \{\hat{\mu}, \hat{\mu}', \mu^{P'}\}$. If necessary, perturb x so that $x \cdot \mu^Q \neq 0$. Let $A = \{x, 0\}$ and $s, t \in S$ such that $B(\mu^e|s) = \hat{\mu}$ and $B(\mu^e|t) = \hat{\mu}'$. For sufficiently small $\alpha, \beta > 0$, we have $\alpha s + \beta t \in S$; moreover, by equation (1) in the main text, $B(\mu^e|\alpha s + \beta t) \rightarrow \hat{\mu}'$ as $\alpha \rightarrow 0$. Thus, we assume without loss of generality (replacing s and t with appropriate αs and βt) that $B(\mu^e|s + t) \in \text{int}(P')$. It follows that $c^s(A) = c^{\mu^P}(A) = 0$ while $c^t(A) = c^{s+t}(A) = c^{\mu^{P'}}(A) = x$. Finally, let $\sigma = [s, t, e - s - t]$ and $\sigma' = [s + t, e - s - t]$. Clearly, $\sigma \sqsupseteq \sigma'$ and $\dot{V}^A(\sigma) = \dot{V}^A(\sigma')$ since $\hat{\mu}$, $\hat{\mu}'$, and $B(\mu^e|s + t)$ belong to the same cell $Q \in \mathcal{Q}$. However, $V^A(\sigma') > V^A(\sigma)$ because $V^A(s + t) > V^A(s) + V^A(t)$, where $V^A(\tilde{s}) := v^A(B(\mu^e|\tilde{s}))(\tilde{s} \cdot \mu^e)$. This contradicts the second assumption of the lemma.

We have shown that for every $Q \in \mathcal{Q}$, there is a unique $P \in \mathcal{P}$ such that $\text{int}(Q) \cap \text{int}(P) \neq \emptyset$. Since \mathcal{P} partitions Δ and cells are regular, it follows that, in fact, $\text{int}(Q) \subseteq \text{int}(P)$.

Step 2: $\mu^P \subseteq \dot{\mu}^Q$.

Suppose toward a contradiction that there is a cell $P \in \mathcal{P}$ such that $\mu^P \neq \dot{\mu}^Q$ for all $Q \in \mathcal{Q}$. Let Q denote the (unique) cell in \mathcal{Q} such that $\mu^P \in Q$. By regularity, there is a neighborhood of μ^P contained in $\text{int}(P)$; since $\mu^P \in Q$, such a neighborhood intersects $\text{int}(Q)$. Thus, by Step 1, $\text{int}(Q) \subseteq \text{int}(P)$. Moreover, since $\mu^e = \dot{\mu}^e$, we have $P \neq P^e$ and $Q \neq Q^e$, where $\mu^e \in P^e \in \mathcal{P}$, $\dot{\mu}^e \in Q^e \in \mathcal{Q}$, and $\text{int}(Q^e) \subseteq \text{int}(P^e)$. There are two cases: either $\mu^P \notin \text{co}\{\dot{\mu}^Q, \mu^e\}$ or $\mu^P \in \text{co}\{\dot{\mu}^Q, \mu^e\}$.

If $\mu^P \notin \text{co}\{\dot{\mu}^Q, \mu^e\}$, there exists x such that $x \cdot \mu^P < 0$ and $x \cdot \tilde{\mu} > 0$ for $\tilde{\mu} \in \text{co}\{\dot{\mu}^Q, \mu^e\}$. Let $A = \{x, 0\}$ and $s, t \in S$ such that $B(\mu^e|s) = \dot{\mu}^Q$ and $B(\mu^e|t) = \mu^e$. As in Step 1, we may choose s and t so that $s + t \in S$ and $B(\mu^e|s + t) \in \text{int}(Q^e) \subseteq \text{int}(P^e)$. Thus, $c^s(A) = c^{\mu^P}(A) = 0$ and $c^t(A) = c^{s+t}(A) = c^{\mu^e}(A) = x$. Letting $\sigma = [s, t, e - s - t]$ and $\sigma' = [s + t, e - s - t]$, it follows that $\sigma \sqsupseteq \sigma'$, $\dot{V}^A(\sigma) = \dot{V}^A(\sigma')$, and $V^A(\sigma') > V^A(\sigma)$, contradicting the second assumption of the lemma.

If instead $\mu^P \in \text{co}\{\dot{\mu}^Q, \mu^e\}$, we may strictly separate μ^e from $\text{co}\{\dot{\mu}^Q, \mu^P\}$: there exists x such that $x \cdot \mu^e < 0$ and $x \cdot \tilde{\mu} > 0$ for $\tilde{\mu} \in \text{co}\{\dot{\mu}^Q, \mu^P\}$. Moreover, we may choose x so that

the line $x \cdot \hat{\mu}' = 0$ passes through $\text{int}(P)$ and, therefore, so that there exists $\hat{\mu} \in P$ so that $x \cdot \hat{\mu} < 0$. Let $s, t \in S$ so that $s+t \in S$, $B(\mu^e|s) = \hat{\mu}$, $B(\mu^e|t)$, and $B(\mu^e|s+t) \in \text{int}(Q^e)$. Let $A = \{x, 0\}$. Then $c^s(A) = c^{\mu^P}(A) = x$ and $c^t(A) = c^{s+t}(A) = 0$. Letting $\sigma = [s, t, e-s-t]$ and $\sigma' = [s+t, e-s-t]$, it follows that $\sigma \sqsupseteq \sigma'$, $\dot{V}^A(\sigma) = \dot{V}^A(\sigma')$, and $V^A(\sigma') > V^A(\sigma)$, contradicting the second assumption of the lemma.

Step 3: for every $Q \in \mathcal{Q}$, there exists $P \in \mathcal{P}$ such that $Q \subseteq P$.

Let $Q \in \mathcal{Q}$. By Step 1, there is a unique $P \in \mathcal{P}$ such that $\text{int}(Q) \subseteq \text{int}(P)$. Suppose toward a contradiction that there exists $\hat{\mu} \in Q$ such that $\hat{\mu} \notin P$; such a $\hat{\mu}$ must be on the boundary of Q , so $\hat{\mu} \neq \dot{\mu}^Q$ by regularity. Since $\text{int}(Q) \subseteq \text{int}(P)$, we also have that $\hat{\mu}$ is on the boundary of P (otherwise there is a neighborhood of $\hat{\mu}$ contained in the complement of P ; but every such neighborhood intersects $\text{int}(Q)$, contradicting $\text{int}(Q) \subseteq \text{int}(P)$).

Since \mathcal{P} partitions Δ and $\hat{\mu} \notin P$, there is a cell $P' \in \mathcal{P}$ ($P' \neq P$) such that $\hat{\mu} \in P'$. By regularity, $\mu^{P'} \in \text{int}(P')$. Moreover, since $\hat{\mu}$ is on the boundary of P , $\hat{\mu}$ is also on the boundary of P' . Thus, we may strictly separate $\mu^{P'}$ from the closure of P ; in particular, there exists x such that $x \cdot \mu^{P'} < 0$ and $x \cdot \tilde{\mu} > 0$ for $\tilde{\mu} \in \text{co}\{\hat{\mu}, \dot{\mu}^Q\}$. Choose $s, t \in S$ so that $B(\mu^e|s) = \hat{\mu}$, $B(\mu^e|t) = \dot{\mu}^Q$, and $B(\mu^e|s+t) \in \text{int}(Q)$. Letting $\sigma = [s, t, e-s-t]$ and $\sigma' = [s+t, e-s-t]$, it follows that $\sigma \sqsupseteq \sigma'$, $\dot{V}^A(\sigma) = \dot{V}^A(\sigma')$, and $V^A(\sigma') > V^A(\sigma)$, contradicting the second assumption of the lemma. \square

A.8 Proof of Proposition 8

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

Lemma 3. *Suppose $\langle \mathcal{P}, \mu^P \rangle$ 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 A such that (A, σ) and (A, σ') are μ^P -decisive.*

Proof. Suppose σ and σ' are \mathcal{P} -equivalent. Observe that for every μ^P -decisive pair $(A, \hat{\sigma})$, $V^A(\hat{\sigma}) = \sum_{P \in \mathcal{P}} (\mu^e s^{P,\hat{\sigma}}) \cdot c^{\mu^P}(A)$ because decisiveness implies $c^{\mu^P}(A)$ is a singleton for all $P \in \mathcal{P}$ where $s^{P,\hat{\sigma}} \neq 0$. 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 menu A such that (A, σ) and (A, σ') are μ^P -decisive but $V^A(\sigma) \neq V^A(\sigma')$. 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 $\mu^\delta \setminus \{\mu^{P^*}\}$; that is, there exists x such that $x \cdot \mu^{P^*} > 0 > x \cdot \mu^{P'}$ for all $\mu^{P'} \in \mu^\delta \setminus \{\mu^{P^*}\}$. By regularity, we may assume that x is such that the menu $A = \{x, 0\}$ makes (A, σ) and

(A, σ') μ^P -decisive (if necessary, perturb x so that $c^s(A)$ is a singleton for all $s \in \sigma \cup \sigma'$). 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 $\mu^P \in \mu^\delta \setminus \{\mu^{P^*}\}$. Thus, $V^A(\sigma) \neq V^A(\sigma')$ provided $(\mu^e \delta^{P^*}) \cdot x \neq 0$. Since the separation is strict and $\langle \mathcal{P}, \mu^P \rangle$ is regular, we may perturb x if necessary to ensure $(\mu^e \delta^{P^*}) \cdot x \neq 0$. \square

Proof that (i) implies (ii). Let $\sigma, \sigma' \in \mathcal{E}$ and suppose $\dot{V}^A(\sigma) = \dot{V}^A(\sigma')$ for all A such that (A, σ) and (A, σ') are $\dot{\mu}^Q$ -decisive. By Lemma 3, σ and σ' are Q -equivalent. Since Q is finer than \mathcal{P} , it follows that σ and σ' are \mathcal{P} -equivalent. Thus, by Lemma 3, $V^A(\sigma) = V^A(\sigma')$ for all $\dot{\mu}^Q$ -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 Q -equivalent. By Lemma 3 and the hypothesis of (ii), this implies σ and σ' are μ^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, Q is finer than \mathcal{P} .

A.9 Proof of Proposition 9

Fix $s \in S$ and let $\mu^* = B(\mu^e|s)$ and $\mu^P = \mu^s$ where $\mu^s \in P \in \mathcal{P}$. If $A \in \mathcal{A}^*$, then there exist $x^*, y^* \in A$ such that $\bar{V}^A(s) = x^* \cdot \mu^*$ and $V^A(s) = y^* \cdot \mu^P$. In particular, $x^* \cdot \mu^* \geq x \cdot \mu^*$ and $y^* \cdot \mu^P \geq y \cdot \mu^P$ for all $x, y \in A$. Let $A^* = \{x^* - x^*, y^* - x^*\} = \{0, y^* - x^*\}$. Then $\bar{V}^A(s) - V^A(s) = \bar{V}^{A^*}(s) - V^{A^*}(s)$. Hence, to compute $L_\mu(s)$, it is without loss of generality to consider menus of the form $\{0, y\}$ where $\|y\| \leq 1$. We therefore rewrite the $L_\mu(s)$ as

$$\begin{aligned} L_\mu(s) &= \sup_{\|y\| \leq 1} 0 \cdot \mu^* - y \cdot \mu^* \text{ subject to: } 0 \cdot \mu^* \geq y \cdot \mu^* \text{ and } y \cdot \mu^P > 0 \cdot \mu^P \\ &= \inf_{\|y\| \leq 1} y \cdot \mu^* \text{ subject to: } 0 \geq y \cdot \mu^* \text{ and } y \cdot \mu^P > 0. \end{aligned}$$

The first constraint ensures the Bayesian prefers action 0 over y at signal s while the second ensures the Coarse Bayesian prefers y over 0 at s . Hence, we seek the infimum of $y \cdot \mu^*$ over all y on the unit (hyper)sphere bounded by the planes $y \cdot \mu^* \leq 0$ and $y \cdot \mu^P > 0$. Clearly, the infimum is characterized by a point y^* on the plane $y \cdot \mu^P = 0$. Thus, we seek a point on the disc $\{y : y \cdot \mu^P = 0 \text{ and } \|y\| \leq 1\}$ tangent to a plane $y \cdot \mu^* = c$ with normal μ^* . There are two such points; one maximizes $y \cdot \mu^*$, the other minimizes it.

Restricting attention to the case $\mu^* \neq \mu^P$, the first constraint does not bind. Thus, the Lagrangian is

$$\mathcal{L} = -y \cdot \mu^* + \lambda_1(y \cdot \mu^P) + \lambda_2(y \cdot y - 1).$$

Setting $\frac{\partial \mathcal{L}}{\partial y_\omega} = 0$ gives $2\lambda_2 y = \mu^* - \lambda_1 \mu^P$. Then $y \cdot \mu^P = 0$ implies $0 = \mu^* \cdot \mu^P - \lambda_1 \|\mu^P\|^2$ and $y \cdot y = 1$ implies $2\lambda_2 = \mu^* \cdot y - \lambda_1 \mu^P \cdot y = \mu^* \cdot y$. Thus, $\lambda_1 = \frac{\mu^* \cdot \mu^P}{\|\mu^P\|^2}$, so that

$$2\lambda_2 y = \mu^* - \left(\frac{\mu^* \cdot \mu^P}{\|\mu^P\|^2} \right) \mu^P.$$

Since $2\lambda_2 = \mu^* \cdot y$, this implies $(\mu^* \cdot y)y = \mu^* - \left(\frac{\mu^* \cdot \mu^P}{\|\mu^P\|^2} \right) \mu^P$. Thus, any solution y satisfies

$$(\mu^* \cdot y)^2 = \|\mu^*\|^2 - \frac{(\mu^* \cdot \mu^P)^2}{\|\mu^P\|^2} = \|\mu^*\|^2 - \frac{\|\mu^*\|^2 \|\mu^P\|^2 \cos^2 \theta}{\|\mu^P\|^2} = \|\mu^*\|^2 \sin^2 \theta$$

where $\theta \in (0, \frac{\pi}{2}]$ is the angle (in radians) between μ^* and μ^P . Thus, $L_\mu(s) = |y \cdot \mu^*| = \|\mu^*\| \sin \theta$, which is increasing in θ . Observe that $D_\mu(s) = \left\| \frac{\mu^*}{\|\mu^*\|} - \frac{\mu^P}{\|\mu^P\|} \right\|$ is the length of the chord connecting the points $\frac{\mu^*}{\|\mu^*\|}$ and $\frac{\mu^P}{\|\mu^P\|}$ on the unit circle. The length of a chord with central angle θ is $2 \sin(\frac{\theta}{2})$, which is strictly increasing on $[0, \frac{\pi}{2}]$. Thus, $D_\mu(s) = 2 \sin(\frac{\theta}{2})$ increases if and only if θ increases, so that $D_\mu(s)$ increases if and only if $L_\mu(s)$ increases.

A.10 Proof of Proposition 10

To see that (ii) implies (i), observe that $\dot{v}^A(\hat{\mu}) \neq v^A(\hat{\mu})$ only if $\hat{\mu}$ belongs to a cell Q such that $\dot{\mu}^Q \notin \mu^P$. Every such Q is a singleton because $\dot{\mu}$ is less biased than μ , which implies $\mu^P \subseteq \dot{\mu}^Q$ and, hence, that Q is a “new” cell. Thus, $\dot{v}^A(\hat{\mu}) = \bar{v}^A(\hat{\mu}) \geq v^A(\hat{\mu})$.

To prove that (i) implies (ii), first apply Proposition 9 to get that $\dot{\mu}$ is less biased than μ . Therefore, $\mu^P \subseteq \dot{\mu}^Q$. We need to show that Q is finer than \mathcal{P} and that every cell Q such that $\dot{\mu}^Q \in \dot{\mu}^Q \setminus \mu^P$ is a singleton.

First, we verify that Q is finer than \mathcal{P} . Suppose toward a contradiction that there is a cell $Q \in \mathcal{Q}$ that intersects two or more distinct cells of \mathcal{P} . There is a unique $P \in \mathcal{P}$ such that $\dot{\mu}^Q \in P$. Let $P' \neq P$ be another cell of \mathcal{P} such that $Q \cap P' \neq \emptyset$. Clearly, $\dot{\mu}^Q \notin P'$. Let $\partial P'$ denote the boundary of P' . There are two cases.

Case 1: $\dot{\mu}^Q \notin \partial P'$. Then, since P' is convex, there exists $x \in \mathbb{R}^N$ that strictly separates $\dot{\mu}^Q$ and P' : $x \cdot \dot{\mu}^Q > 0 > x \cdot \hat{\mu}$ for all $\hat{\mu} \in P'$. Let $A = \{x, 0\}$. Then $v^A(\hat{\mu}') = 0$ for all $\hat{\mu}' \in Q \cap P'$ because $0 > x \cdot \mu^P$. However, $\dot{v}^A(\hat{\mu}') = x \cdot \hat{\mu}'$ for all $\hat{\mu}' \in Q \cap P'$ because $x \cdot \dot{\mu}^Q > 0$. Since $0 > x \cdot \hat{\mu}'$ for all $\hat{\mu}' \in Q \cap P'$, it follows that $\dot{v}^A(\hat{\mu}') < v^A(\hat{\mu}')$ for such $\hat{\mu}'$, a contradiction.

Case 2: $\dot{\mu}^Q \in \partial P'$. Then P' is not a singleton (otherwise $\mu^{P'} = \dot{\mu}^Q \notin P'$), forcing P' to be regular. Moreover, Q is regular because it intersects the (disjoint) sets P and P' . Thus, there are disjoint open neighborhoods $N_Q \subseteq Q$ and $N_{P'} \subseteq P'$ of $\dot{\mu}^Q$ and $\mu^{P'}$. Since N_Q and $N_{P'}$ are convex, there exists $x \in \mathbb{R}^N$ that strictly separates them: $x \cdot \hat{\mu} > 0 > x \cdot \hat{\mu}'$ for all $\hat{\mu} \in N_Q$ and $\hat{\mu}' \in N_{P'}$. Moreover, $\dot{\mu}^Q \in Q \cap \partial P'$ implies $N_Q \cap P' \neq \emptyset$, where $\partial P'$; by regularity, $N_Q \cap P'$ is a full-dimensional subset of $Q \cap P'$. Perturb x so that the plane $x \cdot \hat{\mu} = 0$ passes through the interior of $N_Q \cap P'$ (but not the point $\dot{\mu}^Q$); this can be done by shifting the plane toward the point $\dot{\mu}^Q$. Then x no longer separates N_Q and $N_{P'}$, but the set $C := \{\hat{\mu} \in N_Q \cap P' : 0 > x \cdot \hat{\mu}\}$ is nonempty, and we still have $x \cdot \dot{\mu}^Q > 0$ and $0 > x \cdot \hat{\mu}'$ for all $\hat{\mu}' \in N_{P'}$. Letting $A = \{x, 0\}$, it follows that $v^A(\hat{\mu}) = 0 > x \cdot \hat{\mu} = \dot{v}^A(\hat{\mu})$ for all $\hat{\mu} \in C$, a contradiction.

Next, we verify that every cell Q such that $\dot{\mu}^Q \in \dot{\mu}^Q \setminus \mu^P$ is a singleton. Suppose toward a contradiction that there exists $\dot{\mu}^Q \in \dot{\mu}^Q \setminus \mu^P$ such that Q is not a singleton. Since $\dot{\mu}$ is more sophisticated than μ , there is a unique $P \in \mathcal{P}$ such that $Q \subseteq P$. Note that $\mu^P = \dot{\mu}^P \in \dot{\mu}^P$. Since μ^Q belongs to the relative interior of Q , there exists $\mu^* \in Q$ such that $\dot{\mu}^Q \notin \{\alpha\mu^* + (1 - \alpha)\mu^P : \alpha \in [0, 1]\} := L$. The set L is closed and convex, and therefore can be strictly separated from $\dot{\mu}^Q$: there exists $x \in \mathbb{R}^N$ such that $x \cdot \dot{\mu}^Q > 0 > x \cdot \hat{\mu}$ for all $\hat{\mu} \in L$. In particular, both $x \cdot \mu^* < 0$ and $x \cdot \mu^P < 0$. Let $A = \{0, x\}$. Then, at (Bayesian) posterior $\mu^* \in Q \subseteq P$, the $\langle \mathcal{P}, \mu^P \rangle$ representation selects 0 from A : $v^A(\mu^*) = 0$. Under representation $\langle \mathcal{Q}, \mu^Q \rangle$, however, x is selected from A at posterior μ^* because $\mu^* \in Q$ and $x \cdot \dot{\mu}^Q > 0$. Thus, $\dot{v}^A(\mu^*) = x \cdot \mu^* < 0$, so that $\dot{v}^A(\mu^*) < v^A(\mu^*)$.

References

Barberis, N. and A. Shleifer (2003). Style investing. *Journal of financial Economics* 68(2), 161–199.

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.

Bordalo, P., K. Coffman, N. Gennaioli, and A. Shleifer (2016). Stereotypes. *The Quarterly Journal of Economics* 131(4), 1753–1794.

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

Bursztyn, L. and D. Y. Yang (2022). Misperceptions about others. *Annual Review of Economics* 14(1), 425–452.

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

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

De Bondt, W. F. and R. Thaler (1985). Does the stock market overreact? *The Journal of finance* 40(3), 793–805.

de Clippel, G. and X. Zhang (2022). Non-bayesian persuasion. *Journal of Political Economy (forthcoming)*.

Dominiak, A., M. Kovach, and G. Tserenjigmid (2023). Inertial updating. *arXiv preprint arXiv:2303.06336*.

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. M. (2021). An axiomatic model of persuasion. *Econometrica* 89(5), 2081–2116.

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

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. (1954). *The foundations of statistics*. New York: Wiley.

Schwartzstein, J. and A. Sunderam (2021). Using models to persuade. *American Economic Review* 111(1), 276–323.

Sharot, T. and N. Garrett (2016). Forming beliefs: Why valence matters. *Trends in cognitive sciences* 20(1), 25–33.

Sharpe, W. F. et al. (1992). Asset allocation: Management style and performance measurement. *Journal of portfolio Management* 18(2), 7–19.

Teeple, K. (2022). Support, resistance, and technical trading. *Working Paper*.

Thaler, M. (2021). Overinference from weak signals, underinference from strong signals, and the psychophysics of interpreting information. *arXiv preprint arXiv:2109.09871*.

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

Tversky, A. and D. Kahneman (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological review* 90(4), 293.

Whitmeyer, M. (2023). Bayes= blackwell, almost. *arXiv preprint arXiv:2302.13956*.

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

Zhao, C. (2022). Pseudo-bayesian updating. *Theoretical Economics* 17(1), 253–289.