

The Empirical Content of Bayesianism*

Pooya Molavi[†]

January 18, 2024

Abstract

This paper characterizes the conditions under which the observed beliefs of a group of agents are consistent with Bayesian updating. Beliefs are consistent with Bayesianism if they arise from the application of Bayes' rule given some subjective distribution for the state and the signals agents observe between periods. The paper's main finding is that beliefs are consistent with Bayesianism if and only if the mean of the distribution of posteriors is uniformly absolutely continuous with respect to the prior. Furthermore, the paper shows that the existing results on the empirical content of Bayesianism rely on additional restrictions on permissible subjective distributions, such as the requirement that agents have correct beliefs about the distribution of signals.

*This paper is based on the second chapter of my PhD dissertation at MIT. An earlier version of this paper was previously circulated under the title "Tests of Bayesian Rationality." I am grateful to Daron Acemoglu, Marios Angeletos, Sandeep Baliga, Drew Fudenberg, Arda Gitmez, Parag Pathak, Larry Samuelson, Alvaro Sandroni, Eran Shmaya, Alireza Tahbaz-Salehi, Juuso Toikka, Olivier Wang, Iván Werning, Alex Wolitzky, Muhammet Yildiz, and seminar participants at MIT, the Becker Friedman Institute, and Aalto University for their valuable comments.

[†]Northwestern University, pmolavi@kellogg.northwestern.edu.

1 Introduction

Following the treatise of [Savage \(1972\)](#), the Bayesian theory of probability has become the dominant paradigm in the modeling of decision-making under uncertainty. This paradigm's dominance in economics is not unwarranted. It allows one to assign probabilities to unique or rare events. It has an elegant foundation in the study of rational choice under uncertainty. And it is appealing from a normative perspective—as [Epstein and Le Breton \(1993\)](#) proclaim, “dynamically consistent beliefs must be Bayesian.” What is less clear is whether Bayesianism is a good positive model of individual behavior. Settling this question requires characterizing the testable predictions of Bayesian rationality.

This paper characterizes the empirical content of Bayesianism. It considers an analyst who observes how the beliefs of a group of agents evolve. The analyst can perfectly observe agents' beliefs about an arbitrary state of the world but not how agents plan to update their beliefs based on their signals. The analyst aims to determine if the observed belief sequence is consistent with Bayesian updating given a joint subjective distribution for the state and signal. The paper's main result is that the analyst can rationalize his observation as consistent with Bayesianism if and only if the mean of the distribution of posteriors is uniformly absolutely continuous with respect to the prior.

This result suggests that Bayesianism has limited testable implications. In particular, when the state space is finite, a belief sequence is consistent with Bayesian updating if and only if the posterior mean is supported on a subset of the prior's support. If the prior has full support on a finite state space, then Bayesianism imposes *no* restrictions on the distribution of posteriors. These results question the feasibility of testing Bayesianism in observational data.

The paper's characterization result is obtained under assumptions that make it *easier* for the analyst to reject agents' Bayesianism. The analyst is assumed to directly observe agents' beliefs, instead of having to infer them from their actions. The analyst is free to elicit what agents believe about an arbitrary state belonging to an arbitrary state space, and he can observe those beliefs without any observation noise. I further assume that there is a large number of ex ante identical agents, and agents observe i.i.d. signals. All these assumptions make it easier to disprove agents' Bayesianism; yet, any observation meeting the uniform absolute continuity condition can be rationalized. However, the assumptions make it possible to show that this condition is both necessary and sufficient for consistency with Bayesianism.

This finding may appear at odds with the existing results in the literature. [Aumann and](#)

Maschler (1995) and Kamenica and Gentzkow (2009) argue that a belief sequence is Bayes' plausible if and only if the posterior mean equals the prior. Shmaya and Yariv (2016) argue that any belief sequence in which the prior is in the relative interior of the convex hull of posteriors is consistent with agents' use of Bayes' rule. Both of these conditions are more restrictive than the uniform absolute continuity condition derived in this paper. However, those results are obtained under additional restrictions on what constitutes a reasonable subjective belief for agents.

I prove two additional theorems to clarify the relationship between this paper's characterization and those in the literature. The theorems adapt the existing results to the more general setting of the current paper, thus making them directly comparable to the paper's result. They demonstrate that the earlier results characterize the empirical content of Bayesianism only under additional assumptions on agents' subjective beliefs. Aumann and Maschler (1995) and Kamenica and Gentzkow (2009) do so by requiring agents to have correct beliefs about the distribution of signals, whereas Shmaya and Yariv (2016) require the subjective belief to have the same support as the true distribution.

Beyond the literature discussed above, this paper also contributes to the literature on deviations from rational expectations. This literature can be roughly divided into two strands. The first strand, such as Esponda and Pouzo (2016, 2021), maintains the assumption of Bayesian updating but admits the possibility that agents hold misspecified priors.¹ The second strand studies the implications of non-Bayesian updating rules such as representativeness and availability heuristics (Tversky and Kahneman, 1974), confirmation bias (Rabin and Schrag, 1999), and diagnostic expectations (Bordalo, Gennaioli, and Shleifer, 2018).² This paper's main result clarifies the relationship between these two strands of the literature by showing that almost any non-Bayesian updating rule is observationally equivalent to Bayesian updating given a misspecified prior about the distribution of signals. Bohren and Hauser (2023) also study the question of when non-Bayesian updating rules can be represented as misspecification. They extend Shmaya and Yariv (2016)'s analysis by making agents' forecasts of their future beliefs observable and deriving necessary and sufficient conditions for an updating rule and a forecast to have a misspecified-model representation. In contrast, this paper's focus is characterizing the empirical content of Bayesianism absent information on agents' beliefs about how they will update their beliefs.

¹See also Bohren (2016), Fudenberg, Romanyuk, and Strack (2017), Frick, Iijima, and Ishii (2020), Fudenberg, Lanzani, and Strack (2021), and the references therein.

²See Epstein, Noor, and Sandroni (2010), Molavi, Tahbaz-Salehi, and Jadbabaie (2018), and Cripps (2019) for other examples of non-Bayesian updating rules.

2 Setup

This section introduces the environment and defines what it means for belief sequences to be consistent with Bayesianism.

2.1 The Environment

I consider an analyst who examines whether a group of agents updates their beliefs about a fixed state of the world using Bayes' rule. The state is denoted by x and belongs to a separable metric space X .

There is a large number of agents indexed by $i \in I$. In each of two periods $t = 0, 1$, the analyst elicits what each agent believes about the state of the world. Agent i 's time- t belief about the value of x is a probability distribution, denoted by $\mu_{it} \in \Delta(X)$. I assume that the analyst can perfectly observe μ_{it} for all i and $t = 0, 1$.

Agents' beliefs might evolve between the periods due to new information. I let s_i denote the signal observed by agent i between the two periods and assume without loss of generality that agents' signals belong to the set $S \equiv \Delta(X)$. Each agent i uses a measurable mapping $\varphi_i : s_i \mapsto \mu_{i1}$ to form her posterior based on her realized signal s_i .^{3 4}

I make several assumptions that all help the analyst conclude that agents must *not* be Bayesian—these assumptions make the negative result of the paper even more striking. First, agents are ex ante identical. In particular, $\mu_{i0} = \mu_0^*$ for some $\mu_0^* \in \Delta(X)$ and all $i \in I$. Second, agents' signals are independent and identically distributed, with $\mathbb{P} \in \Delta(S)$ denoting the true distribution of signals given the fixed state of the world. Third, agents all use the same mapping $\varphi = \varphi_i$ to form their beliefs as a function of their signals. Fourth, the number of agents is large enough that the empirical distribution of observed posteriors $\{\mu_{i1}\}_{i \in I}$ provides an arbitrarily good approximation to the corresponding population distribution. Specifically, I assume that the analyst can perfectly observe the population distribution of agents' posterior beliefs, denoted by $F_1^* \in \Delta(\Delta(X))$. Finally, the analyst is assumed to know everything described so far. While these assumptions lead to a tight characterization result, they are not required for the paper's finding that Bayesianism only imposes a weak restriction on belief sequences. Section 4.1 elaborates on this point.

The analyst's question is whether he can interpret the pair (μ_0^*, F_1^*) as being consistent with Bayesian updating given some subjective belief held by agents. The paper's main result

³The term "posterior" here refers to agents' beliefs at time one, regardless of whether those beliefs are derived from agents' prior via Bayes' rule.

⁴The assumption that an agent's posterior is a *deterministic* function of her signal is without loss of generality. Any random updating rule is equivalent to a deterministic updating rule with a true distribution of signals \mathbb{P} that does the randomization for agents.

establishes that any pair (μ_0^*, F_1^*) that satisfies an absolute continuity condition can be rationalized as consistent with Bayes' rule. This result is obtained despite the aforementioned assumptions being biased towards rejecting Bayesianism.

2.2 Bayes Plausibility

Before presenting the main result, it is necessary to define what it means for an observed pair (μ_0^*, F_1^*) to be consistent with Bayesianism. Agents are Bayesian if they (i) possess a well-defined subjective belief over the set of state-signal pairs; (ii) assign positive probabilities to signals that occur with positive probabilities; and (iii) update their beliefs following each signal that occurs with a positive probability using Bayes' rule.

The first two requirements can be formalized in the following way: First, agents must have a subjective distribution $\mathbb{Q} \in \Delta(X \times S)$ over the set of states and signals.⁵ Second, their subjective distribution must assign a positive probability to any signal that is realized with a positive probability. This is to ensure that agents can use Bayes' rule following every contingency. This requirement can be expressed using the following notion:

Definition 1. If P and Q are probability distributions over the same measurable space, P is *uniformly absolutely continuous* with respect to Q if there exists a positive constant c such that $P(E) \leq cQ(E)$ for any measurable set E .⁶

I require the subjective distribution \mathbb{Q} to be such that the true distribution of signals \mathbb{P} is uniformly absolutely continuous with respect to the S -marginal \mathbb{Q}_S . When \mathbb{P} has a finite support, this condition reduces to the requirement that \mathbb{Q} does not rule out any signal that is realized with a positive probability. More generally, the uniform absolute continuity requirement ensures that agents' beliefs do not disproportionately discount the likelihood of certain signals. Although this condition is not necessary for the paper's main finding, having a more demanding notion of Bayesianism strengthens the result by highlighting the fact that the conclusion does not rely on the inapplicability of Bayes' rule after zero-probability events. It also allows me to turn the statement of the main result into an "if and only if" statement.

The third criterion for Bayesianism is agents' use of Bayes' rule to update their beliefs. This criterion is formally expressed through the concept of regular conditional probability. Given the measurable space $(X \times S, \mathcal{X} \times \mathcal{S})$ and probability distribution $\mathbb{Q} \in \Delta(X \times S)$, a *regular conditional probability* of \mathbb{Q} given \mathcal{S} is a mapping $\nu : S \times X \rightarrow [0, 1]$ such that (i) $\nu(s, \cdot)$ is a probability distribution on X for every $s \in S$, (ii) the mapping $s \mapsto \nu(s, D)$ is measurable for

⁵Note that the subjective distribution \mathbb{Q} does not have an i subscript since I have assumed that agents are ex ante identical.

⁶See Lemma 1 of the appendix for an equivalent way of defining uniform absolute continuity.

all $D \in \mathcal{X}$, and (iii) the kernel ν satisfies

$$\mathbb{Q}(D \times E) = \int_E \nu(s, D) \mathbb{Q}_S(ds) \quad (1)$$

for all $D \in \mathcal{X}$ and $E \in \mathcal{S}$, where \mathbb{Q}_S is the S -marginal of \mathbb{Q} . The regular conditional probability ν defines a mapping $\varphi : s \mapsto \nu(s, \cdot)$ from agents' signals to their posteriors. Agents are Bayesian given subjective distribution \mathbb{Q} if they use this mapping to update their beliefs.

The regular conditional probability ν determines agents' posterior beliefs as a function of their subjective belief and the realized signal. However, it does not specify the distribution of those posterior beliefs. In particular, for any event $D \in \mathcal{X}$, $\nu(s, D)$ is a random variable whose distribution depends on the distribution of the signal s . To determine the probability with which each posterior is realized, one needs to use the true distribution of signals. Given a regular conditional probability ν and the true distribution of signals \mathbb{P} , agents' posterior about the state x is distributed according to the probability distribution $F_\nu \in \Delta(\Delta(X))$, defined as

$$F_\nu(\{\mu_1 \in \Delta(X) : \mu_1 \in E\}) \equiv \mathbb{P}(\{s \in S : \nu(s, \cdot) \in E\}) \quad (2)$$

for all $E \in \mathcal{S}$. This is the observed distribution of posteriors if agents are Bayesian with subjective distribution \mathbb{Q} (and the corresponding ν) and the true distribution of signals is given by \mathbb{P} .

I can now define what it means for agents' observed belief sequence to be consistent with Bayesianism.

Definition 2. Given the state space X , signal space $S = \Delta(X)$, and true distribution of signals $\mathbb{P} \in \Delta(S)$, a pair of observations (μ_0^*, F_1^*) , consisting of agents' prior and the distribution of their posteriors about the state, is *consistent with Bayesianism* if there exists a subjective distribution $\mathbb{Q} \in \Delta(X \times S)$ for agents that satisfies the following conditions:

- (a) $\mathbb{Q}_X = \mu_0^*$,
- (b) \mathbb{P} is uniformly absolutely continuous with respect to \mathbb{Q}_S ,
- (c) \mathbb{Q} has a regular conditional probability ν such that $F_\nu = F_1^*$,

where \mathbb{Q}_X and \mathbb{Q}_S are the X - and S -marginals of the subjective distribution \mathbb{Q} , respectively, and F_ν is the distribution of posteriors defined in (2).

This definition formalizes the intuitive notion of Bayesianism laid out at the beginning of this subsection. The analyst's task is to find a subjective distribution \mathbb{Q} that explains the observed changes in agents' beliefs. This conjectured \mathbb{Q} is a joint distribution for the state

and signal that must satisfy three conditions: Condition (a) of Definition 2 simply requires the conjectured distribution to be consistent with the observed prior. Condition (b) is the requirement that agents assign non-vanishing probabilities to signals that are realized with positive probabilities. Condition (c) requires that the observed distribution of posteriors matches the distribution obtained when agents start with the conjectured distribution \mathbb{Q} , observe signals as per \mathbb{P} , and update their beliefs using Bayes' rule.

3 The Empirical Content of Bayesianism

3.1 The Main Result

The paper's main result provides a necessary and sufficient condition for the observed pair (μ_0^*, F_1^*) to be consistent with Bayesianism:

Theorem 1. *A pair (μ_0^*, F_1^*) is consistent with Bayesianism if and only if the mean of the distribution of posteriors $\bar{\mu}_1^* \equiv \int \mu F_1^*(d\mu)$ is uniformly absolutely continuous with respect to the prior μ_0^* .*

Proof of the "if" direction. The proof of this direction is constructive. Given the measurable space (X, \mathcal{X}) and the true signal distribution \mathbb{P} , I construct the subjective distribution \mathbb{Q} that rationalizes an observed pair (μ_0^*, F_1^*) that satisfies

$$\bar{\mu}_1^*(D) \leq c\mu_0^*(D) \tag{3}$$

for all $D \in \mathcal{X}$.⁷ Since $\bar{\mu}_1^*$ and μ_0^* are both probability measures on (X, \mathcal{X}) , the constant c in equation (3) must be at least weakly larger than one. If $c < 1$, then $\bar{\mu}_1^*(X) < c\mu_0^*(X) = c < 1$, a contradiction. If $c = 1$, then $\bar{\mu}_1^* = \mu_0^*$. This is because $\bar{\mu}_1^*(D) \leq \mu_0^*(D)$ implies $\bar{\mu}_1^*(D^c) \geq \mu_0^*(D^c)$, where $D^c \in \mathcal{X}$ denotes the complement of D . But $\bar{\mu}_1^*(D^c) \leq \mu_0^*(D^c)$ by equation (3). Therefore, $\bar{\mu}_1^*(D) = \mu_0^*(D)$. Since D is an arbitrary measurable set, $\bar{\mu}_1^* = \mu_0^*$. In the remainder of the proof, I construct the subjective distribution \mathbb{Q} that rationalizes (μ_0^*, F_1^*) , separately for the $c > 1$ and $c = 1$ cases.

I first prove the result for the $c > 1$ case. I start by constructing the regular conditional probability $\nu : S \times \mathcal{X} \rightarrow [0, 1]$ that represents agents' posterior about the state $x \in X$ conditional on the signal s . Let $\ominus \in S$ denote a signal such that $\mathbb{P}(\ominus) = 0$. Such a signal always exists since $S = \Delta(X)$ is uncountable, but there are at most countably many signals $s \in S$ such that $\mathbb{P}(s) > 0$. For any $s \in S$ such that $\varphi(s) \in \text{supp } F_1^*$ and $s \neq \ominus$, set $\nu(s, D) = \varphi(s)(D)$ for all $D \in \mathcal{X}$.

⁷For any metric space X , the sigma-algebra on X is assumed to be the Borel sigma-algebra, denoted by \mathcal{X} , and $\Delta(X)$ denotes the set of probability distributions on (X, \mathcal{X}) , endowed with the topology of weak convergence and the corresponding Borel sigma-algebra.

Set

$$v(\ominus, D) = \frac{c}{c-1} \mu_0^*(D) - \frac{1}{c-1} \bar{\mu}_1^*(D)$$

for all $D \in \mathcal{X}$. Finally, set $v(s, D) = \mu_0^*(D)$ for any $s \in S$ such that $\varphi(s) \notin \text{supp } F_1^* \cup \{\ominus\}$ and all $D \in \mathcal{X}$, indicating that agents' posterior equals their prior conditional on any signal realized with zero probability. Note that, by construction, the mapping $s \mapsto v(s, D)$ is measurable for any $D \in \mathcal{X}$. Therefore, to show that v is a kernel, it is sufficient to show that $v(s, \cdot)$ is a probability distribution on (X, \mathcal{X}) for all $s \in S$. This holds by construction for all $s \neq \ominus$. In order to verify that $v(\ominus, \cdot)$ is a probability measure, first note that,

$$v(\ominus, X) = \frac{c}{c-1} \mu_0^*(X) - \frac{1}{c-1} \bar{\mu}_1^*(X) = 1,$$

and

$$v(\ominus, \emptyset) = \frac{c}{c-1} \mu_0^*(\emptyset) - \frac{1}{c-1} \bar{\mu}_1^*(\emptyset) = 0,$$

where I am using the facts that $\mu_0^*(X) = \bar{\mu}_1^*(X) = 1$ and $\mu_0^*(\emptyset) = \bar{\mu}_1^*(\emptyset) = 0$ since μ_0^* and $\bar{\mu}_1^*$ are both probability measures on X . Next note that, for any set $D \in \mathcal{X}$,

$$v(\ominus, D) = \frac{c}{c-1} \mu_0^*(D) - \frac{1}{c-1} \bar{\mu}_1^*(D) \geq \frac{c}{c-1} \mu_0^*(D) - \frac{c}{c-1} \mu_0^*(D) = 0,$$

where the inequality follows equation (3). Finally, $v(\ominus, \cdot)$ is countably additive since both μ_0^* and $\bar{\mu}_1^*$ are probability measures and thus are countably additive. Therefore, $v(\ominus, \cdot)$ is a probability measure on (X, \mathcal{X}) .

I can now define the subjective distribution \mathbb{Q} , starting with its S -marginal distribution \mathbb{Q}_S .

Let

$$\mathbb{Q}_S(E) \equiv \frac{1}{c} \mathbb{P}(E) + \frac{c-1}{c} \mathbb{1}\{\ominus \in E\},$$

for all $E \in \mathcal{S}$, and let

$$\mathbb{Q}(D \times E) \equiv \int_E v(s, D) \mathbb{Q}_S(ds) \tag{4}$$

for all $D \in \mathcal{X}$ and $E \in \mathcal{S}$. Note that since the sigma-algebra $(\mathcal{X} \times \mathcal{S})$ over $(X \times S)$ is generated by sets of the form $D \times E$ with $D \in \mathcal{X}$ and $E \in \mathcal{S}$, the above expression fully specifies the probability distribution \mathbb{Q} . Moreover, comparing equations (1) and (4) shows that v is indeed a regular conditional probability of \mathbb{Q} given \mathcal{S} . Lastly, since $\mathbb{Q}_S(E) = \frac{1}{c} \mathbb{P}(E) + \frac{c-1}{c} \mathbb{1}\{\ominus \in E\} \geq \frac{1}{c} \mathbb{P}(E)$ for all $E \in \mathcal{S}$, the true distribution \mathbb{P} is uniformly absolutely continuous with respect to \mathbb{Q}_S .

It remains to show that $\mathbb{Q}_X = \mu_0^*$ and that the distribution of posteriors F_v , defined in equation (2), coincides with the observed posterior distribution F_1^* . Note that, by definition, $F_1^* = \mathbb{P} \circ \varphi^{-1}$. Let

$$\hat{S} \equiv \text{supp } \mathbb{P} = \{s \in S : \varphi(s) \in \text{supp } F_1^*\}.$$

For any $D \in \mathcal{X}$,

$$\begin{aligned}
\mathbb{Q}_X(D) &= \int_S \nu(s, D) \mathbb{Q}_S(ds) \\
&= \frac{1}{c} \int_{\hat{S}} \varphi(s)(D) \mathbb{P}(ds) + \frac{c-1}{c} \nu(\Theta, D) \\
&= \frac{1}{c} \int_{\varphi(\hat{S})} \mu(D) \mathbb{P} \circ \varphi^{-1}(d\mu) + \mu_0^*(D) - \frac{1}{c} \bar{\mu}_1^*(D) \\
&= \frac{1}{c} \int_{\text{supp } F_1^*} \mu(D) F_1^*(d\mu) + \mu_0^*(D) - \frac{1}{c} \bar{\mu}_1^*(D) \\
&= \frac{1}{c} \bar{\mu}_1^*(D) + \mu_0^*(D) - \frac{1}{c} \bar{\mu}_1^*(D) = \mu_0^*(D).
\end{aligned}$$

I next show that $F_\nu = F_1^*$. Note that

$$\mathbb{P}(\{s \in S : \varphi(s) \notin \text{supp } F_1^*\}) = \mathbb{P}(\{s \in S : s \notin \text{supp } \mathbb{P}\}) = 0.$$

On the other hand, by construction, $\mathbb{P}(\Theta) = 0$. Therefore, for any $E \in \mathcal{S}$,

$$\begin{aligned}
F_\nu(E) &= \mathbb{P}(\{s \in S : \nu(s, \cdot) \in E\}) \\
&= \mathbb{P}(\{s \in S : \nu(s, \cdot) \in E, \varphi(s) \in \text{supp } F_1^*, s \neq \Theta\}) \\
&= \mathbb{P}(\{s \in S : \varphi(s) \in E\}) = F_1^*(E).
\end{aligned}$$

This completes the proof for the $c > 1$ case.

I next consider the case where $c = 1$. I construct the regular conditional probability $\nu : S \times \mathcal{X} \rightarrow [0, 1]$ by setting $\nu(s, D) = \varphi(s)(D)$ for all $s \in S$ such that $\varphi(s) \in \text{supp } F_1^*$ and all $D \in \mathcal{X}$ and setting $\nu(s, D) = \mu_0^*(D)$ for all $s \in S$ such that $\varphi(s) \notin \text{supp } F_1^*$ and all $D \in \mathcal{X}$. By construction, $\nu(s, \cdot)$ is a probability distribution on (X, \mathcal{X}) , and the mapping $s \mapsto \nu(s, D)$ is measurable for all $D \in \mathcal{X}$. I set the S -marginal \mathbb{Q}_S of the subjective distribution equal to the true distribution \mathbb{P} of signals and define \mathbb{Q} as in (4). By construction, ν is a regular conditional probability of \mathbb{Q} given \mathcal{S} , and \mathbb{P} is uniformly absolutely continuous with respect to \mathbb{Q}_S . Next, note that, for any $D \in \mathcal{X}$,

$$\mathbb{Q}_X(D) = \int_S \nu(s, D) \mathbb{P}(ds) = \int_{\text{supp } F_1^*} \mu(D) F_1^*(d\mu) = \bar{\mu}_1^*(D) = \mu_0^*(D),$$

where the last equality follows the fact that $\bar{\mu}_1 = \mu_0^*$ when $c = 1$, established in the first paragraph of the proof. Moreover, by an argument similar to the $c > 1$ case,

$$F_\nu(E) = \mathbb{P}(\{s \in S : \nu(s, \cdot) \in E\}) = \mathbb{P}(\{s \in S : \varphi(s) \in E\}) = F_1^*(E)$$

for all $E \in \mathcal{S}$. This shows that the subjective distribution \mathbb{Q} constructed above rationalizes the observed pair (μ_0^*, F_1^*) . \square

Proof of the “only if” direction. Let \mathbb{Q} denote agents’ subjective distribution on $(X \times S)$ and ν denote the regular conditional probability of \mathbb{Q} given \mathcal{S} . The existence of ν follows the assumption that \mathbb{Q} satisfies condition (c) of Definition 2. Since \mathbb{Q} satisfies condition (a) and ν is a regular conditional probability of \mathbb{Q} given \mathcal{S} ,

$$\mu_0^*(D) = \mathbb{Q}_X(D) = \int_S \nu(s, D) \mathbb{Q}_S(ds)$$

for all $D \in \mathcal{X}$. On the other hand, for any $D \in \mathcal{X}$,

$$\bar{\mu}_1^*(D) = \int_S \mu(D) F_1^*(d\mu) = \int_S \mu(D) F_\nu(d\mu) = \int_S \nu(s, D) \mathbb{P}(ds),$$

where the second equality is due to the assumption that \mathbb{Q} satisfies condition (c) of Definition 2, and the third equality follows the definition of F_ν . Finally, \mathbb{P} is uniformly absolutely continuous with respect to \mathbb{Q}_S by condition (b) of Definition 2. Therefore, by Lemma 1 of the appendix, there exists a Radon–Nikodym derivative $f \equiv \frac{d\mathbb{P}}{d\mathbb{Q}_S}$ that is bounded up to sets of \mathbb{Q}_S -measure zero. Therefore,

$$\int_S \nu(s, D) \mathbb{P}(ds) = \int_S \nu(s, D) f(s) \mathbb{Q}_S(ds) \leq c \int_S \nu(s, D) \mathbb{Q}_S(ds),$$

where c denotes a positive constant such that $f \leq c$ with \mathbb{Q}_S -probability one. Combining the last three displays implies that $\bar{\mu}_1^*(D) \leq c \mu_0^*(D)$. The fact that c can be chosen independently of the set $D \in \mathcal{X}$ establishes that $\bar{\mu}_1^*$ is uniformly absolutely continuous with respect to $\mu_0^*(D)$ and completes the proof. \square

The uniform absolute continuity of the posterior mean with respect to the prior encompasses the entire empirical content of Bayesianism. Absent additional a priori restrictions on what constitutes a reasonable subjective distribution, any belief sequence that satisfies this condition is consistent with Bayesian updating. It is easy to see that the absolute continuity condition is necessary for Bayesianism: If the prior of a Bayesian agent assigns zero probability to an event, her posterior must also assign zero probability to the event—regardless of agents’ subjective belief and the true distribution of signals. What is more surprising is that absolute continuity is also sufficient for consistency with Bayesianism. The proof establishes this result by constructing a subjective distribution \mathbb{Q} starting from a pair (μ_0^*, F_1^*) , which satisfies the absolute continuity condition, and showing that the conjectured \mathbb{Q} satisfies the properties set out in Definition 2.

3.2 Special Cases

Building on Theorem 1, I next discuss two of its corollaries. The first corollary addresses scenarios where the state space X is finite, a common situation in applications. The corollary

follows from the theorem given the following observation: For two distributions P and Q over a finite set, P is uniformly absolutely continuous with respect to Q if the support of P is a subset of the support of Q .

Corollary 1. *Suppose the state space X is finite. A pair (μ_0^*, F_1^*) is consistent with Bayesianism if and only if*

$$\text{supp } \bar{\mu}_1^* \subseteq \text{supp } \mu_0^*,$$

where $\bar{\mu}_1^* \equiv \int \mu F_1^*(d\mu)$ denotes the mean of the distribution of posteriors.

This result summarizes the empirical content of Bayesianism in discrete settings. The mean posterior belief cannot assign positive probabilities to states that have zero probability according to the prior. The necessity of this property for beliefs to be consistent with Bayesianism is apparent given Bayes' rule. The corollary goes a step further by establishing that this property is also sufficient for belief sequences to be consistent with Bayesianism. A further corollary of Corollary 1 is that *any* posterior distribution is consistent with Bayesianism when the prior has full support over a finite state space.

Corollary 2. *Suppose the state space X is finite and μ_0^* has full support over X . Then the pair (μ_0^*, F_1^*) is consistent with Bayesianism for any distribution F_1^* of posteriors.*

The following example illustrates an application of the theorem and its corollaries:

Example 1. The state takes values in the set $X = \{H, L\}$. Agents' observed prior is the uniform distribution over X . The observed distribution of posteriors F_1^* is as follows: For a quarter of agents, the belief that the state is H goes up to 0.8. The remaining agents become certain that the state is H . Should the observation of this belief sequence lead the analyst to conclude that agents are not Bayesian? The answer is no; the observed pair (μ_0^*, F_1^*) is indeed consistent with Bayesianism. This conclusion follows Corollary 2 of Theorem 1 by noting that μ_0^* has full support over X .

The observation of F_1^* imposes some restrictions on the true distribution of signals \mathbb{P} and the mapping φ used by agents to update their beliefs. Since agents' posteriors take on two values, there are at least two signals that are realized with positive probabilities. The observation of one set of signals moves agents' posterior belief that the state is H to 0.8. Since a quarter of agents end up with the posterior $\mu_1(H) = 0.8$, the signals that lead to this posterior must have probability $\mathbb{P}(\{s : \varphi(s) = (\mu_1(H) = 0.8)\}) = 0.25$. Likewise, there is a set of signals that has true probability $\mathbb{P}(\{s : \varphi(s) = (\mu_1(H) = 1)\}) = 0.75$ and makes agents certain that the state is H . With slight abuse of notation, in the remainder of the example, I refer to the

$\{s : \varphi(s) = (\mu_1(H) = 0.8)\}$ and $\{s : \varphi(s) = (\mu_1(H) = 1)\}$ events simply as the $s = 0.8$ and $s = 1$ signals, respectively.⁸

I illustrate how (μ_0^*, F_1^*) can be rationalized by finding a subjective distribution \mathbb{Q} such that the belief sequence of a Bayesian agent with subjective distribution \mathbb{Q} matches the observed prior and distribution of posteriors. The distribution \mathbb{Q} needs to satisfy three requirements for it to rationalize the observed prior μ_0^* and posterior distribution F_1^* . First, \mathbb{Q} must be consistent with the observed prior μ_0^* ; i.e., $\mathbb{Q}(H) = \mu_0^*(H) = 0.5$. Second, agents must assign positive probabilities to the $s = 0.8$ and $s = 1$ signals for Bayes' rule to be applicable after the observation of those signals. Third, agents' posterior conditional on the observation of the $s = 0.8$ and $s = 1$ signals must be consistent with their observed posteriors; i.e., $\mathbb{Q}(H|s = 0.8) = 0.8$ and $\mathbb{Q}(H|s = 1.0) = 1.0$.

One needs to also specify what agents believe about the probability of observing signals other than 0.8 and 1.0. I start by assuming that agents believe the signal can only take values $s = 0.8$ and $s = 1.0$. This assumption constrains \mathbb{Q} to satisfy $\mathbb{Q}(\{(x, s) : s \in \{0.8, 1.0\}\}) = 1$. This constraint, together with the requirements previously discussed and the requirement that $\mathbb{Q}(x, s) \geq 0$ for any (x, s) , yields a mixed system of equalities and inequalities for the four unknown probabilities $\mathbb{Q}(H, 0.8)$, $\mathbb{Q}(L, 0.8)$, $\mathbb{Q}(H, 1.0)$, and $\mathbb{Q}(L, 1.0)$:

$$\mathbb{Q}(H, 0.8) + \mathbb{Q}(L, 0.8) > 0, \quad (5)$$

$$\mathbb{Q}(H, 1.0) + \mathbb{Q}(L, 1.0) > 0, \quad (6)$$

$$\frac{\mathbb{Q}(H, 0.8)}{\mathbb{Q}(H, 0.8) + \mathbb{Q}(L, 0.8)} = 0.8, \quad (7)$$

$$\frac{\mathbb{Q}(H, 1.0)}{\mathbb{Q}(H, 1.0) + \mathbb{Q}(L, 1.0)} = 1.0, \quad (8)$$

$$\mathbb{Q}(H, 0.8), \mathbb{Q}(L, 0.8), \mathbb{Q}(H, 1.0), \mathbb{Q}(L, 1.0) \geq 0, \quad (9)$$

$$\mathbb{Q}(H, 0.8) + \mathbb{Q}(H, 1.0) = 0.5, \quad (10)$$

$$\mathbb{Q}(L, 0.8) + \mathbb{Q}(L, 1.0) = 0.5. \quad (11)$$

It is easy to verify that this system does not have a solution.

Thus, for the observed belief sequence to arise from Bayesian updating by agents, they must believe the possibility that the signal takes values outside the set $\{0.8, 1.0\}$. With slight abuse of notation, I let $s = \ominus$ denote the event that the signal takes a value outside the set

⁸This would not be an abuse of notation under the assumption that φ is the identity mapping. Note that the assumption that φ is the identity mapping is innocuous in this example since observing F_1^* only identifies $\mathbb{P} \circ \varphi^{-1} = F_1^*$ —but not \mathbb{P} or φ . Nonetheless, the construction in the example can be easily modified to allow for the possibility that F_1^* , \mathbb{P} , and φ are separately identified by the analyst. I forgo this extension here since it would lead to additional notational complexity without offering any new insights. See the proof of Theorem 1 for the general construction.

$\{0.8, 1.0\}$. Constraints (10) and (11) now have to be modified as follows:

$$\mathbb{Q}(H, 0.8) + \mathbb{Q}(H, 1.0) + \mathbb{Q}(H, \ominus) = 0.5, \quad (12)$$

$$\mathbb{Q}(L, 0.8) + \mathbb{Q}(L, 1.0) + \mathbb{Q}(L, \ominus) = 0.5. \quad (13)$$

The remaining requirements, expressed in equations (5)–(9), remain intact. However, \mathbb{Q} must now additionally satisfy the two non-negativity requirements:

$$\mathbb{Q}(H, \ominus), \mathbb{Q}(L, \ominus) \geq 0. \quad (14)$$

Equations (5)–(9) and (12)–(14) constitute a mixed system of equalities and inequalities for the six unknown probabilities $\mathbb{Q}(H, 0.8)$, $\mathbb{Q}(L, 0.8)$, $\mathbb{Q}(H, 1.0)$, $\mathbb{Q}(L, 1.0)$, $\mathbb{Q}(H, \ominus)$, and $\mathbb{Q}(L, \ominus)$. The fact that the mean of the posterior distribution has the same support as the prior is sufficient to ensure that this system has a solution. One such solution—and the one corresponding to the proof of Theorem 1—is as follows:

	0.8	1.0	\ominus
H	0.25	0.25	0
L	0.0625	0	0.4375

Note that the observed pair (μ_0^*, F_1^*) can be rationalized only if agents are allowed to have a misspecified belief about the distribution of signals. If agents were to hold a correctly specified belief, the subjective distribution \mathbb{Q} would have to agree with the true distribution \mathbb{P} on the probabilities of different signals. But then the systems of equalities and inequalities that determine \mathbb{Q} would have no solution.

4 Discussion

This section discusses the assumptions that underpin the paper’s main result and how relaxing them changes the conclusion.

4.1 Assumptions

The paper makes a number of explicit and implicit assumptions in order to obtain the characterization in Theorem 1. These assumptions can be divided into three categories:

1. Homogeneity assumptions: Agents are ex ante identical, use the same function to map their signals to posteriors, and receive i.i.d. signals.

2. Knowledge assumptions: The analyst directly and perfectly observes agents' beliefs—instead of having to identify them from choice data. He knows the true distribution of signals and the mapping agents use to update their beliefs. He observes the population distribution of posteriors.
3. No restrictions on the subjective belief: The analyst puts no restrictions on what constitutes a reasonable subjective belief. Moreover, the analyst does not observe agents' beliefs concerning the likelihood of various signals.

These sets of assumptions serve different purposes. The homogeneity assumptions help with the identification of agents' subjective beliefs. The analyst observes each agent only after the realization of a single signal. Without the homogeneity assumptions, how an agent behaves after a signal would not be informative of how other agents would have behaved if they observed that same signal.

The knowledge assumptions limit what the analyst can freely choose in order to rationalize his observations. When the analyst knows an object such as the true signal distribution, he has no flexibility in choosing that object to rationalize agents' behavior. This limits the set of observations that can be rationalized by the analyst and makes it easier for him to reject his observations as consistent with Bayesianism.

The first two sets of assumptions strengthen the “if” direction of Theorem 1. Under these assumptions, the analyst cannot reject agents' Bayesianism as long as his observations satisfy the uniform absolute continuity condition. This remains true despite the constraints these assumptions place on the analyst's ability to rationalize observations. If these assumptions were to be relaxed, certain observations—even those that do not adhere to the uniform absolute continuity condition—might still be explainable as consistent with Bayesianism. Hence, the uniform absolute continuity condition may be seen as an upper bound on the empirical content of Bayesianism. The main role of the first two sets of assumptions is to make this bound tight by turning the theorem into an “if and only if” result.

The final set of assumptions enables the analyst to freely choose the subjective belief \mathbb{Q} to rationalize his observations. The analyst is assumed to perfectly observe what agents believe about the state but not how they intend to update those beliefs. Therefore, he is free to form any conjecture about how agents are planning to do so. This degree of freedom makes it more challenging for the analyst to disprove the Bayesianism of agents. Unlike the first two sets of assumptions, which narrow the range of observations consistent with Bayesianism, the third set introduces more flexibility, thus expanding this range.

The assumption that any well-defined subjective distribution is permissible is made in keeping with [Savage \(1972\)](#)'s idea of purely subjective beliefs. I assume that any subjective belief \mathbb{Q} that is consistent with agents' elicited beliefs is a valid subjective distribution. Agents' rationality is judged not based on their subjective beliefs but based on how they update those beliefs. In the next subsection, I discuss two alternatives to this assumption proposed in the literature and how they change the conclusion of [Theorem 1](#).

4.2 Alternative Notions of Bayesianism

The first alternative I consider assumes that agents have a correctly specified belief about the distribution of signals. This assumption leads to the martingale property of Bayesian beliefs: The expectation of the posterior must equal the prior. [Aumann and Maschler \(1995\)](#) and [Kamenica and Gentzkow \(2009\)](#) show that this is indeed the only restriction imposed on beliefs by the requirement that agents are Bayesian. The following theorem generalizes this result to general metric spaces. More importantly, however, it highlights the fact that the martingale property characterizes the empirical content of Bayesianism *only* when agents are required to have correct beliefs about the distribution of signals.

Theorem 2. *The pair (μ_0^*, F_1^*) is consistent with Bayesianism given a subjective distribution \mathbb{Q} with the S -marginal satisfying $\mathbb{Q}_S = \mathbb{P}$ if and only if $\mu_0 = \bar{\mu}_1^* \equiv \int \mu F_1^*(d\mu)$.*

Proof. Proof of the “if” direction is identical to the $c = 1$ case in the proof of [Theorem 1](#)'s “if” direction. The proof of the “only if” direction follows similar steps as in the proof of [Theorem 1](#)'s “only if” direction. In particular, by an identical argument, $\mu_0^*(D) = \int_S \nu(s, D) \mathbb{Q}_S(ds)$ and $\bar{\mu}_1^*(D) = \int_S \nu(s, D) \mathbb{P}(ds)$. The assumption that $\mathbb{Q}_S = \mathbb{P}$ completes the proof. \square

A more permissive notion of Bayesianism is proposed by [Shmaya and Yariv \(2016\)](#). They allow Bayesian agents to have incorrect beliefs about the distribution of signals— as long as the supports of those beliefs coincide with the support of the true distribution. The following theorem generalizes [Shmaya and Yariv \(2016\)](#)'s Lemma 1 to general metric state spaces and arbitrary true signal distributions. It reduces to their result when both X and $\text{supp } \mathbb{P}$ are finite sets. However, its main significance is to clarify that [Shmaya and Yariv \(2016\)](#)'s conclusion relies on an a priori restriction on what constitutes a reasonable subjective distribution.

Theorem 3. *The following statements are equivalent:*

- A. *The pair (μ_0^*, F_1^*) is consistent with Bayesianism given a subjective distribution \mathbb{Q} with a S -marginal \mathbb{Q}_S that is uniformly absolutely continuous with respect to \mathbb{P} .*

B. There exists a probability measure $\lambda \in \Delta(S)$ such that λ and F_1^* are mutually uniformly absolutely continuous and $\mu_0^* = \int \mu \lambda(d\mu)$.

Proof. First, suppose there exists a probability measure $\lambda \in \Delta(S)$ such that λ and F_1^* are mutually uniformly absolutely continuous and $\mu_0^* = \int \mu \lambda(d\mu)$. By Lemma 1 of the appendix, there exist Radon–Nikodym derivatives $f \equiv \frac{d\lambda}{dF_1^*}$ and $\frac{1}{f} \equiv \frac{dF_1^*}{d\lambda}$ such that $c \leq f \leq C$ for some positive constants c, C (up to sets of F_1^* - and λ -measure zero). Set $\nu(s, D) = \varphi(s)(D)$ for all $s \in S$ and $D \in \mathcal{X}$, and set $\mathbb{Q}_S(ds) = f(\varphi(s))\mathbb{P}(ds)$. I need to show that \mathbb{Q}_S , as defined above, is indeed a probability distribution on (S, \mathcal{S}) . By construction, $\mathbb{Q}_S(E) \geq 0$ for all $E \in \mathcal{S}$, and $\mathbb{Q}_S(\emptyset) = 0$. Next, note that

$$\int_S \mathbb{Q}_S(ds) = \int_S f(\varphi(s))\mathbb{P}(ds) = \int_S f(\mu)\mathbb{P} \circ \varphi^{-1}(d\mu) = \int_S f(u)F_1^*(d\mu) = \int_S \lambda(d\mu) = 1,$$

where the first equality is by definition, the second one uses the change-of-variables formula for pushforward measures, the third equality is due to the fact that $F_1^* = \mathbb{P} \circ \varphi^{-1}$, the fourth one uses the definition of f , and the last equality is because λ is a probability measure on S . Finally, \mathbb{Q}_S is countably additive since \mathbb{P} is countably additive. Therefore, \mathbb{Q}_S is a well-defined probability distribution. I finish the construction by defining \mathbb{Q} as in equation (1). Note that, by construction, ν is a conditional probability of \mathbb{Q} given \mathcal{S} . Furthermore, by an argument similar to the one in the above display,

$$\mathbb{Q}_X = \int_S \nu(s, \cdot)\mathbb{Q}_S(ds) = \int_S \varphi(s)f(\varphi(s))\mathbb{P}(ds) = \int_S \mu f(\mu)F_1^*(d\mu) = \int_S \mu \lambda(d\mu) = \mu_0^*,$$

where the last equality is by assumption. Therefore, condition (a) of Definition 2 is satisfied. Furthermore, since $\mathbb{Q}_S(ds) = f(\varphi(s))\mathbb{P}(ds)$ and $c \leq f \leq C$ almost surely, \mathbb{Q}_S and \mathbb{P} are mutually uniformly absolutely continuous. That is, condition (b) of Definition 2 is satisfied, and \mathbb{Q}_S is uniformly absolutely continuous with respect to \mathbb{P} . On the other hand,

$$F_\nu(E) = \mathbb{P}(\{s \in S : \nu(s, \cdot) \in E\}) = \mathbb{P}(\{s \in S : \varphi(s) \in E\}) = \mathbb{P} \circ \varphi^{-1}(E) = F_1^*(E)$$

for all $E \in \mathcal{S}$, implying that condition (c) is also satisfied.

Next, suppose (μ_0^*, F_1^*) is consistent with Bayesianism given a subjective distribution \mathbb{Q} with a S -marginal \mathbb{Q}_S that is uniformly absolutely continuous with respect to \mathbb{P} , and let ν denote the regular conditional probability of \mathbb{Q} given \mathcal{S} . The existence of ν follows the assumption that \mathbb{Q} satisfies condition (b) of Definition 2. I define $\lambda \in \Delta(S)$ as follows:

$$\lambda(E) \equiv \mathbb{Q}_S(\{s \in S : \nu(s, \cdot) \in E\})$$

for all $E \in \mathcal{S}$. I next show that λ and F_1^* are mutually uniformly absolutely continuous and $\mu_0^* = \int \mu \lambda(d\mu)$. Since \mathbb{Q} satisfies condition (a) and ν is a regular conditional probability of \mathbb{Q}

given \mathcal{S} ,

$$\mu_0^* = \mathbb{Q}_X = \int_{\mathcal{S}} \nu(s, \cdot) \mathbb{Q}_S(ds) = \int_{\mathcal{S}} \mu \lambda(d\mu),$$

where the last equality is by definition. On the other hand, for all $E \in \mathcal{S}$,

$$\lambda(E) = \mathbb{Q}_S(\{s \in \mathcal{S} : \nu(s, \cdot) \in E\}) \geq \frac{1}{c} \mathbb{P}(\{s \in \mathcal{S} : \nu(s, \cdot) \in E\}) = \frac{1}{c} F_\nu(E) = \frac{1}{c} F_1^*(E)$$

for some positive constant c , where the first two equalities are by definition, the inequality is by condition (b) of Definition 2, and the third equality is by condition (c) of Definition 2. Likewise, since \mathbb{Q}_S is uniformly absolutely continuous with respect \mathbb{P} , for all $E \in \mathcal{S}$,

$$\lambda(E) = \mathbb{Q}_S(\{s \in \mathcal{S} : \nu(s, \cdot) \in E\}) \leq C \mathbb{P}(\{s \in \mathcal{S} : \nu(s, \cdot) \in E\}) = C F_\nu(E) = C F_1^*(E)$$

for some positive constant C . The fact that c and C can be chosen independently of $E \in \mathcal{S}$ in the above two displays establishes that λ and F_1^* are mutually uniformly absolutely continuous. \square

Technical Appendix

Lemma 1. *Let P and Q be probability distributions over the same measurable space. P is uniformly absolutely continuous with respect to Q if and only if there exists a Radon–Nikodym derivative $f \equiv \frac{dP}{dQ}$ and a positive constant c such that $f \leq c$ up to sets of Q -measure zero.*

Proof. First, suppose there exists a Radon–Nikodym derivative $f \equiv \frac{dP}{dQ}$ and a positive constant c such that $f \leq c$ up to sets of Q -measure zero. For any measurable set E ,

$$P(E) = \int_E dP = \int_E f dQ \leq c \int_E dQ = cQ(E).$$

That is, P is uniformly absolutely continuous with respect to Q .

Next, suppose P is uniformly absolutely continuous with respect to Q . Then, by definition, $P(E) \leq cQ(E)$ for some c and any measurable set E . In particular, $P(E) = 0$ for any E for which $Q(E) = 0$. Therefore, P is absolutely continuous with respect to Q , and so, by the Radon–Nikodym theorem, there exists a derivative $f \equiv \frac{dP}{dQ}$. I finish the proof by showing that f is bounded Q -almost surely. Toward a contradiction, suppose that for any positive constant C there exists a measurable set E with $Q(E) > 0$ such that $f > C$ on E . Then,

$$P(E) = \int_E dP = \int_E f dQ > C \int_E dQ = CQ(E).$$

Since C is arbitrary, there exists no constant c such that $P(E) \leq cQ(E)$ for all E , a contradiction to the assumption that P is uniformly absolutely continuous with respect to Q . \square

References

- Aumann, Robert J. and Michael Maschler (1995), *Repeated Games with Incomplete Information*. MIT press.
- Bohren, J. Aislinn (2016), “Informational Herding with Model Misspecification.” *Journal of Economic Theory*, 163, 222–247.
- Bohren, J. Aislinn and Daniel N. Hauser (2023), “Behavioral Foundations of Model Misspecification.” Working Paper.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018), “Diagnostic Expectations and Credit Cycles.” *The Journal of Finance*, 73, 199–227.
- Cripps, Martin W. (2019), “Divisible Updating.” Working Paper.
- Epstein, Larry G. and Michel Le Breton (1993), “Dynamically Consistent Beliefs Must Be Bayesian.” *Journal of Economic Theory*, 61, 1–22.
- Epstein, Larry G., Jawwad Noor, and Alvaro Sandroni (2010), “Non-Bayesian Learning.” *The BE Journal of Theoretical Economics*, 10, 1–20.
- Esponda, Ignacio and Demian Pouzo (2016), “Berk–Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models.” *Econometrica*, 84, 1093–1130.
- Esponda, Ignacio and Demian Pouzo (2021), “Equilibrium in Misspecified Markov Decision Processes.” *Theoretical Economics*, 16, 717–757.
- Frick, Mira, Ryota Iijima, and Yuhta Ishii (2020), “Misinterpreting Others and the Fragility of Social Learning.” *Econometrica*, 88, 2281–2328.
- Fudenberg, Drew, Giacomo Lanzani, and Philipp Strack (2021), “Limit Points of Endogenous Misspecified Learning.” *Econometrica*, 89, 1065–1098.
- Fudenberg, Drew, Gleb Romanyuk, and Philipp Strack (2017), “Active Learning with a Misspecified Prior.” *Theoretical Economics*, 12, 1155–1189.
- Kamenica, Emir and Matthew Gentzkow (2009), “Bayesian Persuasion.” *American Economic Review*, 101, 2590–2615.
- Molavi, Pooya, Alireza Tahbaz-Salehi, and Ali Jadbabaie (2018), “A Theory of Non-Bayesian Social Learning.” *Econometrica*, 86, 445–490.
- Rabin, Matthew and Joel L. Schrag (1999), “First Impressions Matter: A Model of Confirmatory Bias.” *The Quarterly Journal of Economics*, 114, 37–82.
- Savage, Leonard J. (1972), *The Foundations of Statistics*. Courier Corporation.
- Shmaya, Eran and Leeat Yariv (2016), “Experiments on Decisions Under Uncertainty: A Theoretical Framework.” *American Economic Review*, 106, 1775–1801.
- Tversky, Amos and Daniel Kahneman (1974), “Judgment Under Uncertainty: Heuristics and Biases.” *Science*, 185, 1124–1131.