

LEARNING UNDER UNAWARENESS

Simon Grant

Idione Meneghel

Rabee Tourky

November 3, 2021

Abstract

We propose a model of learning when experimentation is possible, but unawareness and ambiguity matter. In this model, complete lack of information regarding the underlying data generating process is expressed as a (maximal) family of priors. These priors yield posterior inferences that become more precise as more information becomes available. As information accumulates, however, the individual's level of awareness as encoded in the state space may expand. Such newly learned states are initially seen as ambiguous, but as evidence accumulates there is a gradual reduction of ambiguity.

1 Introduction

State spaces and probabilities are ubiquitous in economic models. They provide an unrivaled analytical tool to study situations in which uncertainty matters. Yet, it is not always clear how state spaces and beliefs emerged. If such knowledge is founded upon experience, how so?

Research School of Economics, Australian National University, Canberra ACT 2601, Australia. Emails: simon.grant@anu.edu.au, idione.meneghel@anu.edu.au, rabee.tourky@anu.edu.au.

* We thank George Mailath, Larry Epstein, Larry Samuelson, and Marie-Louise Vierø for invaluable comments. We also thank participants of SAET 2015, Zif Robust Finance & Beyond Conference 2016, FUR 2016, D-TEA 2017, and Virginia Tech Advances in Decision Theory Workshop 2018.

Standard models of learning take for granted (an exogenously given) state space, impose an objective prior, and update it using Bayes' rule in face of new information. Practical problems, however, rarely fit this mould. For all intents and purposes, states of the world are essentially *abstract* representations of resolutions of uncertainty. Moreover, whereas Bayesian updating is an effective tool to sort out the probability of the number of heads in a finite sequence of tosses of a fair coin, in many practical situations, it may not be possible to specify a prior probability. The standard Bayesian machinery is ill-equipped to provide guidance when the new information is surprising—and falls on a zero probability event—or when the new information contradicts past experience—and cannot be categorized in any of the previously considered events. Moreover, it leaves no room for confidence in the probability assessments made.

This paper considers learning when there is incomplete information about the structure of the state space, but further information can be obtained through experimentation. We take for granted that the final objective of such analysis, after observations have been made, is the choice among possible courses of action whose consequences depend on the state of the world. For the present purposes, we are only concerned with the assessment phase of the analysis: the process of building the state space on the go and updating beliefs in a way that satisfies some basic principles of coherence and consistency.

For the problem to be well-defined, we aim to answer two questions. First, we need to sort out what constitutes a state of the world. If it is possible to learn new sources of randomness, then the model should allow for extensions of the state space to accommodate them. In particular, the model of the state space should permit two operations:

1. *Creation*: when new states are added without changing the structure of old events;
and
2. *Refinement*: when old states turn into events that can be partitioned into more richly described states.

Second, we need to work out how to form and update a prior belief on this space that

truly reflects complete ignorance. To answer the first question, we invoke the approach of Karni and Vierø [2013, 2017]; to answer the second, we propose an imprecise version of the Dirichlet process prior, defined by Ferguson [1973, 1974].

More specifically, the theory presented here is adapted to the following kind of problem. Suppose that an expert is challenged to express her opinions about possible plans of action whose consequences depend on the unknown state of the world. Before making her assessment, she has the opportunity to carry out a sequential experiment to learn about possible states of the world and their plausibility. Experiments are described by an underlying stochastic process, that embodies the physical law governing the machinery of the experiment. Ideally, the realizations of these sequential trials provide a full description of the outcome, thus resolving all uncertainty regarding the state of the world. However, the expert's level of awareness restricts her perception of the realized state of the world, she can only partially observe these realizations. Thus the experimenter's awareness level determines which events she can conceive and delimits the boundaries of her conceivable state space. As new evidence becomes available, the experimenter may discover new states. Consequently, her awareness level increases and her conceivable state space expands.

Beliefs are represented by an evolving set of predictive distributions over future, conceivable states of the world; inferences about events can be summarized by upper and lower probabilities. As new states are discovered, probability mass may be shifted from old, non-null events to the events just created. Moreover, newly learned events are initially perceived as ambiguous. As evidence accumulates, however, the experimenter becomes more familiar with these events. The ambiguity associated with those events gradually disappears and the assessment made by the individual converges to their true posterior probability.

The resulting theory captures several desirable, intuitive features:

1. *Rich class of stochastic environments*: the model is able to capture a wide range of data generating processes.

2. *Internal consistency*: the individual's beliefs are revised in a coherent way.
3. *External consistency*: the evolution of beliefs reflects data (frequentist validation).
4. *Possibility of surprises*: the individual is cognizant of the possibility of consequences and actions of which she is currently unaware, but which may be revealed over time; as a result, learning never ends.
5. *Event-specific ambiguity*: ambiguity is related to lack of familiarity; newly learned events are seen as more ambiguous than old events.
6. *Large-sample confidence*: as evidence accumulates, the ambiguity perceived by the individual is gradually reduced and her confidence in her assessment increases.

Our approach relates to four different literatures. First, we provide a model of learning that, in line with Marinacci [2002] and Epstein and Schneider [2007], accommodates ambiguous beliefs. Secondly, we explicitly model the process of inductive reasoning implied by the dynamics of growing awareness described in Dominiak and Tserenjigmid [2017] that utilizes the framework developed in Karni and Vierø [2013, 2017]. Because ambiguity emerges endogenously, when information is surprising, we provide a dynamic foundation for the unanimity rule preference representation axiomatized in Bewley [2002] and Gilboa, Macheroni, Marinacci, and Schmeidler [2010] as well as the partial “comparative likelihood relation” in Nehring [2009]. Finally, we provide an axiomatization of the model's probability kernel, which can be compared to the result in Billot, Gilboa, Samet, and Schmeidler [2005].

There are two other papers that make the connection between levels of awareness and perceptions of ambiguity. Halpern, Rong, and Saxena [2010] introduces unawareness in the context of a Markov decision problem and provides a characterization as to when the individual can ‘learn’ to play nearly optimally. Kochov [2016], on the other hand, proposes a ‘revealed preference’ test to distinguish between those contingencies the individual is unaware of versus those she foresees but whose likelihood she perceives as ambiguous.

The paper is organized as follows. Section 2 describes the underlying data generating process and how actions and consequences are discovered. Section 3 explains how the individual's conception of the world evolves in view of this acquired information. Section 4 investigates the properties of the model. Section 5 explores the implications of the results to the problem of consensus formation of beliefs. Finally, Section 6 discusses the Bayesian interpretation of the model. All the proofs are collected in the Appendix.

2 The discovery process

First, we introduce the elements of the data generating process and the nature of the observations made by the individual. Let $T = \{0, 1, \dots, t, \dots\}$ denote time. There exists a countable set \mathbb{A} of actions, which corresponds to the set of alternatives that are or may become known to the individual. There also exists a set \mathbb{C} of consequences, which we take to be a separable, completely metrizable space.

2.1 Underlying stochastic process

We are given a sequence

$$X_1, X_2, \dots, X_t, \dots$$

of random variables defined on the ambient probability space $(\Omega, \mathcal{F}, \mu)$ and taking values in the common measurable space $(\mathbb{C}^{\mathbb{A}}, \mathcal{B})$, where \mathcal{B} denotes the Borel σ -algebra on the metrizable product $\mathbb{C}^{\mathbb{A}}$. If $x_t \in \mathbb{C}^{\mathbb{A}}$ is a realization of X_t , then action $a \in \mathbb{A}$ is associated with the consequence $(x_t)_a \in \mathbb{C}$, which we also denote by $a(x_t)$. We may also think of the action a as the sequence of random variables $(a(X_1), a(X_2), \dots, a(X_t), \dots)$ with range \mathbb{C} . An element of $\mathbb{C}^{\mathbb{A}}$ specifies the unique consequence associated with every possible action in \mathbb{A} . A realization x_t of X_t , being an element of $\mathbb{C}^{\mathbb{A}}$, thus resolves all uncertainty.

Throughout the paper, we assume the following stochastic dependence condition on $(X_t)_{t \geq 1}$, originally studied by Kallenberg [1988].

Definition 2.1. The sequence of random variables $(X_t)_{t \geq 1}$ is said to be *conditionally exchangeable* if for every $n \geq 1$

$$(X_1, \dots, X_n, X_{n+1}) \sim (X_1, \dots, X_n, X_{n+2}),$$

where \sim means “distributed as.”

We view conditional exchangeability as a structural judgement, that is, it is an assessment regarding the structural properties of the underlying stochastic process, which are the result of the design of the experiment. It is related to the more familiar notion of exchangeability. Recall that a sequence of random variables $(X_t)_{t \geq 1}$ is said to be *exchangeable* whenever every permutation of every finite subsequence has the same distribution, i.e., for every $n \geq 1$

$$(X_{k_1}, \dots, X_{k_n}) \sim (X_{k_{\sigma(1)}}, \dots, X_{k_{\sigma(n)}}),$$

where $\sigma: \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ is a permutation on $\{1, \dots, n\}$.

It is immediate that exchangeable sequences are conditionally exchangeable, but the converse is true only under stationarity. Recall that a sequence of random variables $(X_t)_{t \geq 1}$ is said to be *stationary* whenever the distribution of finite subsequences is invariant over time, i.e., for every $k \geq 1$

$$(X_1, \dots, X_n) \sim (X_{k+1}, \dots, X_{k+n}).$$

By a result due to Kallenberg [1988, Proposition 2.1], if a stationary sequence of random variables is conditionally exchangeable, then it is exchangeable. Kallenberg [1988] provides an example of a non-stationary, but conditionally exchangeable sequence that is not exchangeable.

Furthermore, conditional exchangeability can be understood in terms of the martingale aspect of the process $(X_t)_{t \geq 1}$. In particular, Kallenberg [1988, Proposition 2.2] shows that conditional exchangeability is equivalent to a property later coined as conditionally identically distributed.

Let $\mathcal{G} = (\mathcal{G}_t)_{t \geq 0}$ denote the nested sequence of sub- σ -algebras \mathcal{G}_t of \mathcal{F} such that $\mathcal{G}_0 = \{\emptyset, \Omega\}$ and $\mathcal{G}_t = \sigma(X_1, \dots, X_t)$ for every $t \geq 1$. That is, \mathcal{G} is the filtration induced or generated by the sequence of random variables $(X_t)_{t \geq 1}$. It is the smallest filtration with the property that X_t is \mathcal{G}_t -measurable for every $t \geq 1$.

Definition 2.2. We say that $(X_t)_{t \geq 1}$ is *conditionally identically distributed with respect to* \mathcal{G} (or \mathcal{G} -c.i.d.) if the conditional expectation satisfies

$$\mathbb{E}_\mu[g(X_k) \mid \mathcal{G}_t] = \mathbb{E}_\mu[g(X_{t+1}) \mid \mathcal{G}_t] \quad \mu\text{-a.s.}$$

for every $k > t \geq 0$ and every bounded measurable $g: \mathbb{C}^{\mathbb{A}} \rightarrow \mathbb{R}$.

Essentially, this property means that, at every time t , future realizations $(X_k)_{k > t}$ are identically distributed conditional on past realizations $(X_k)_{k \leq t}$. We refer to Berti, Pratelli, and Rigo [2004] for examples of non-exchangeable, c.i.d. stochastic processes.

2.2 Discovery of actions and consequences

The individual in period 0 is aware of a nonempty, finite set $A_0 \subseteq \mathbb{A}$ of actions, and a nonempty, finite set $C_0 \subseteq \mathbb{C}$ of consequences. We now describe how other actions and consequences are discovered.

Discovery of actions. We suppose that actions in \mathbb{A} are discovered over time, via a given nested sequence

$$A_0 \subseteq A_1 \subseteq \dots \subseteq A_t \subseteq \dots \subseteq \mathbb{A},$$

where each A_t , $t \geq 0$, is a finite, nonempty set of actions. We think of A_0 as the individual's prior knowledge of the available actions. At time $t \geq 1$, the individual discovers or becomes aware of the set $A_t \setminus A_{t-1}$ of additional actions. We do not model how the sequence $(A_t)_{t \geq 0}$ is realized. We leave open the possibility that the sequence is a sample path of some independent stochastic process or the result of some independent learning process. As an example, the actions in $(A_t)_{t \geq 0}$ could represent the sequence of treatment

options for a particular illness as new alternatives are discovered and made available to the specialist. Further, we do not assume that all actions in \mathbb{A} are eventually revealed to the individual. That is, it could be the case that everything is intrinsically knowable, but our model also allows for “unknown unknowns” that remain unknowable indefinitely.

Stochastic discovery of consequences. Throughout the rest of the section, we fix a sample path

$$x_1, x_2, \dots, x_t, \dots$$

of realizations of $(X_t)_{t \geq 1}$.

The realization x_t sets the perfect level of description of the true state of the world. What the individual observes, however, is restricted by her level of unawareness. If $x \in \mathbb{C}^{\mathbb{A}}$, then we let $x|_{A_t} = (x_a)_{a \in A_t}$. In period t , the individual’s information consists of the finite history of all observations up to period t for the given sample path, which can be written as:

$$\widehat{x}_t = (x_1|_{A_1}, x_2|_{A_2}, \dots, x_t|_{A_t}).^1$$

In particular, the individual discovers consequences as she (partially) observes the realizations of the collection of random variables $a(X_t)$, induced by the actions known to her.² Thus the set of consequences known to the individual is the set of all observed consequences given the finite history \widehat{x}_t . Formally, the set of consequences the individual considers possible at the end of period t is given by

$$\widehat{C}_t := C_0 \cup \left\{ \bigcup_{k=1}^t \{a(x_k) : a \in A_k\} \right\} \cup \{\theta_t\},$$

with $\widehat{C}_0 := C_0 \cup \{\theta_0\}$. We interpret θ_t , for $t \geq 0$, as the intangible consequence that represents the possibility of observing an outcome that is not in the set $C_0 \cup \left\{ \bigcup_{k=1}^t \{a(x_k) : a \in A_k\} \right\}$

¹We mark variables that depend on the finite history of observations with a hat ($\widehat{\cdot}$).

²An alternative way to model the discovery of consequences is to assume that the individual chooses an action each period and observes only the consequence associated to that chosen action. Our results can be adapted to that case, under the additional assumption of independence across actions.

of known consequences. That is, we also allow for the individual to conceive of the possibility that some actions may yield, in future periods, consequences which are as yet unknown to her in period t .³ Of course, we assume that $\mathbb{C} \cap \{\theta_0, \theta_1, \dots\} = \emptyset$.

3 The inference problem

In this section, we explain how the individual's perception of the state space evolves over time. Recall that, given the sample path

$$x_1, x_2, \dots, x_t, \dots,$$

the information available to the individual at period t is the finite history

$$\widehat{x}_t = (x_1|_{A_1}, x_2|_{A_2}, \dots, x_t|_{A_t})$$

of all observations up to period t . Given this finite history of observations, the individual's conceivable state space expands to convey her increasing level of awareness.

3.1 Evolution of the set of conceivable states

Following the approach of Karni and Schmeidler [1991], and Karni and Vierø [2013, 2017], conceivable states represent the possible resolution of uncertainty restricted by the awareness level of the individual. That is, having observed the finite history \widehat{x}_t and being aware of the consequences in \widehat{C}_t and actions in A_t , the individual can conceive of states that are elements of the finite set

$$\widehat{S}_t := \widehat{C}_t^{A_t},$$

which we refer to as the set of *conceivable states (at the end of period t)*, for the given history \widehat{x}_t .

³Our analysis does not require the *a priori* introduction of the intangible consequences represented by θ_t . However, this more general set-up brings the model in line with Karni and Vierø [2017].

In particular, for every period t , the observation $x_t|_{A_t}$ is understood as a conceivable state in \widehat{S}_t , that is, $x_t|_{A_t} = (a(x_t))_{a \in A_t} \in \widehat{S}_t$. Moreover, for every $t \geq 1$, the set of conceivable states \widehat{S}_{t-1} induces a partition of \widehat{S}_t , whereby each element in \widehat{S}_{t-1} is mapped to an event in \widehat{S}_t .

The following example illustrates how the conceivable state space evolves over time and how conceivable states in previous periods are mapped to events in later periods. Suppose that the individual is aware of one action a and one consequence c_1 in period $t-1$. The set \widehat{S}_{t-1} comprises two states s_1 and s_2 :

$$\begin{array}{cc} s_1 & s_2 \\ \hline a & c_1 \quad \theta_{t-1} \end{array}$$

In the next period, we consider three alternative scenarios. If $A_t = \{a\}$ and $a(x_t) = c_2$, then a new, previously unknown consequence is revealed in period t . Hence, at the end of period t , the individual conceives of three states in \widehat{S}_t :

$$\begin{array}{ccc} \underbrace{s_1} & & \underbrace{s_2} \\ & s_1^i & s_2^i \quad s_2^{ii} \\ \hline a & c_1 & c_2 \quad \theta_t \end{array}$$

In this scenario, the state s_1 in \widehat{S}_{t-1} becomes s_1^i , and the state s_2 in \widehat{S}_{t-1} is “split” into s_2^i and s_2^{ii} . The partition of \widehat{S}_t induced by \widehat{S}_{t-1} is thus $\{\{s_1^i\}, \{s_2^i, s_2^{ii}\}\}$.

If alternatively $A_t = \{a, b\}$, $a(x_t) = b(x_t) = c_1$, then a new action is discovered but no new consequence, and s_1 and s_2 correspond to events in \widehat{S}_t :

$$\begin{array}{cccc} & \underbrace{s_1} & & \underbrace{s_2} \\ & s_1^i & s_1^{ii} & s_2^i \quad s_2^{ii} \\ \hline a & c_1 & c_1 & \theta_t \quad \theta_t \\ b & c_1 & \theta_t & c_1 \quad \theta_t \end{array}$$

In this alternative scenario, s_1 splits into s_1^i , and s_1^{ii} and s_2 splits into s_2^i and s_2^{ii} . The partition of \widehat{S}_t induced by \widehat{S}_{t-1} is thus $\{\{s_1^i, s_1^{ii}\}, \{s_2^i, s_2^{ii}\}\}$.

Finally, if $A_t = \{a, b\}$, $a(x_t) = c_1$ and $b(x_t) = c_2$, then the new set of conceivable states \widehat{S}_t is:

	s_1			s_2					
	s_1^i	s_1^{ii}	s_1^{iii}	s_2^i	s_2^{ii}	s_2^{iii}	s_2^{iv}	s_2^v	s_2^{vi}
<i>a</i>	c ₁	c ₁	c ₁	c ₂	c ₂	c ₂	θ_t	θ_t	θ_t
<i>b</i>	c ₁	c ₂	θ_t	c ₁	c ₂	θ_t	c ₁	c ₂	θ_t

In this case, the state s_1 is split into s_1^i , s_1^{ii} and s_1^{iii} , and the state s_2 is split into the remaining six states. Therefore, the partition of \widehat{S}_t induced by the state space \widehat{S}_{t-1} conceivable in the previous period is $\{\{s_1^i, s_1^{ii}, s_1^{iii}\}, \{s_2^i, s_2^{ii}, s_2^{iii}, s_2^{iv}, s_2^v, s_2^{vi}\}\}$.


In a nutshell, there are two circumstances that surprise the individual and prompts her to expand the set of states she believes to be possible:

- i. the discovery of new consequences: formally, that corresponds to $\widehat{C}_{t-1} \setminus \{\theta_{t-1}\}$ being a proper subset of $\widehat{C}_t \setminus \{\theta_t\}$; and
- ii. the discovery of new feasible actions: that corresponds to A_{t-1} being a proper subset of A_t .

The advantage of working with this canonical state space and its product structure is that it distinguishes between the information deduced by logical inference—which states can be conceived—and the information explicitly conveyed by the data—the distributions over consequences induced by the actions. In particular, nothing suggests to the individual that the drawings of consequences induced by different actions is independent. However, as will be shown in the next section, the individual's (imprecise) prior reflects her total ignorance regarding possible correlations between different actions. It is thus left to the evidence to reveal the true correlation structure among the states.⁴

We now formalize the embedding illustrated by the previous examples for the general structure of the states. For any pair of periods t and t' , with $t < t'$, we can express the associated embeddings and projections as follows.

⁴The first formal treatment of imprecise probability is from Boole [1854], who suggested it as a way to reconcile the theory of logic, which can express complete ignorance, and the theory of probability.



For every period t , let $\kappa_t: \mathbb{C} \rightarrow \widehat{\mathbb{C}}_t$ denote the mapping that takes consequences in \mathbb{C} that are unknown in period t to θ_t in $\widehat{\mathbb{C}}_t$, and keeps the other, known consequences unchanged. That is,

$$\kappa_t(c) = \begin{cases} c & \text{if } c \in \widehat{\mathbb{C}}_t \setminus \{\theta_t\}, \\ \theta_t & \text{otherwise.} \end{cases}$$

The mapping κ_t induces a natural embedding $\psi_t: \mathbb{C}^{A_t} \rightarrow \widehat{\mathbb{C}}_t^{A_t}$, whereby

$$\psi_t(x) = (\kappa_t(x_a))_{a \in A_t}.$$

Finally, for $t < t'$, we define $\varphi_{t':t}: \widehat{\mathbb{S}}_{t'} \rightarrow \widehat{\mathbb{S}}_t$ as the unique mapping such that

$$\psi_t \circ \text{proj}_{A_t} = \varphi_{t':t} \circ \psi_{t'}.$$

Notice that, because $\widehat{\mathbb{C}}_t \setminus \{\theta_t\} \subseteq \mathbb{C}$, the mapping ψ_t is surjective. Moreover, because $A_t \subseteq A_{t'}$, and $\widehat{\mathbb{C}}_t \setminus \{\theta_t\} \subseteq \widehat{\mathbb{C}}_{t'} \setminus \{\theta_{t'}\}$, it follows that $\varphi_{t':t}$ is also surjective, and hence it has a set-valued inverse. The partition of $\widehat{\mathbb{S}}_{t'}$ induced by $\widehat{\mathbb{S}}_t$ is the one given by the collection of sets $\{\varphi_{t':t}^{-1}(s) : s \in \widehat{\mathbb{S}}_t\}$. In particular, for each state $s \in \widehat{\mathbb{S}}_t$, the event $\varphi_{t':t}^{-1}(s) \subseteq \widehat{\mathbb{S}}_{t'}$ is the collection of conceivable states in $\widehat{\mathbb{S}}_{t'}$ into which the state s has been “split.”

3.2 Evolution of beliefs

We now propose a prior, based on the Dirichlet process, that represents (almost) complete ignorance with respect to the distribution μ .⁵

⁵Benavoli, Mangili, Ruggeri, and Zaffalon [2015] propose a prior similar to the one we describe. They, however, apply it to Bayesian hypothesis testing, obtaining a method with nice asymptotic properties and, at the same time, that is more robust compared to the usual tests.

The Dirichlet process. The simplest and one of the most commonly used nonparametric statistical model is the Dirichlet process, defined by Ferguson [1973, 1974].⁶ It was introduced as a prior over probability distributions and, due to the tractability of the resulting posterior inferences, it is widely employed for Bayesian nonparametric inference.

We use the following notation. If S is a separable, completely metrizable space, then let $\Delta(S)$ denote the set of Borel probability measures on S . For every s in S , let $\delta_s \in \Delta(S)$ denote the (Dirac or degenerate) probability measure that assigns probability one to s obtaining.

To describe the Dirichlet process, let $G(\alpha, \beta)$, with $\alpha > 0$ and $\beta > 0$, denote the Gamma distribution on \mathbb{R}_+ , with Lebesgue density

$$\frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x},$$

where $\Gamma(\alpha)$ is the complete gamma function. For $G(\alpha, \beta)$, α is called the shape parameter and β is the scale parameter. The n -dimensional Dirichlet distribution with parameter $(\alpha_1, \dots, \alpha_n)$, with $\alpha_k \geq 0$ and $\sum_k \alpha_k > 0$, is the distribution of the random vector $\left(\frac{Z_1}{\sum_k Z_k}, \dots, \frac{Z_n}{\sum_k Z_k}\right)$ taking values in the unit simplex Δ^{n-1} , where

$$Z_k \stackrel{\text{i.i.d.}}{\sim} G(\alpha_k, 1).$$

Definition 3.1. Let π be a finite non-null (probability) measure on $(\mathbb{X}, \mathcal{B}_{\mathbb{X}})$, where $\mathcal{B}_{\mathbb{X}}$ is the Borel σ -algebra of subsets of \mathbb{X} , and $\alpha > 0$. Then a stochastic process Θ is a *Dirichlet process with base measure π and concentration parameter α* , denoted by $\Theta \sim DP(\alpha, \pi)$, if for every finite measurable partition $\{B_1, \dots, B_n\}$ of \mathbb{X} , the random vector $(\Theta(B_1), \dots, \Theta(B_n))$ has a Dirichlet distribution with parameter $(\alpha\pi(B_1), \dots, \alpha\pi(B_n))$.

⁶We refer to Ferguson [1973, 1974] for the first description of the Dirichlet process and its use in Bayesian nonparametric inference. More recently, this literature expanded rapidly; we refer to Walker, Damien, Laud, and Smith [1999] and Escobar and West [1995] for more modern treatments, as well as generalizations of the Dirichlet process.

Under the Dirichlet process, data is assumed to be generated according to the law

$$\begin{aligned}\Theta &\sim DP(\alpha, \pi) \\ X_1, X_2, \dots &| \Theta \stackrel{\text{i.i.d.}}{\sim} \Theta.\end{aligned}$$

That is, the parameter that explains the data is the random measure Θ itself. The Dirichlet Process is thus a probability distribution on $\Delta(\mathbb{X})$, the space of probability measures over $(\mathbb{X}, \mathcal{B}_{\mathbb{X}})$. We record the following properties of the Dirichlet process.

Theorem 3.2 (Ferguson [1973, Theorem 1], Ghosh and Ramamoorthi [2003, Chapter 3]).

Let (X_1, \dots, X_t) be i.i.d. samples from Θ and suppose Θ has a $DP(\alpha, \pi)$ distribution. Then:

1. The conditional distribution of Θ given the finite sample (X_1, \dots, X_t) is a Dirichlet process with parameters $(\alpha + t, \frac{\alpha}{\alpha+t}\pi + \frac{1}{\alpha+t} \sum_k \delta_{X_k})$:

$$\Theta | (X_1, \dots, X_t) \sim DP\left(\alpha + t, \frac{\alpha}{\alpha+t}\pi + \frac{1}{\alpha+t} \sum_k \delta_{X_k}\right).$$

2. The predictive distribution of the next observation X_{t+1} , given the finite sample (X_1, \dots, X_t) , is

$$\mathbb{P}[X_{t+1} | X_1, \dots, X_t] = \frac{\alpha}{\alpha+t}\pi + \frac{1}{\alpha+t} \sum_{k=1}^t \delta_{X_k}.$$

We can re-write the predictive distribution of the next observation as

$$\mathbb{P}[X_{t+1} | X_1, \dots, X_t] = \frac{\alpha}{\alpha+t}\pi + \frac{t}{\alpha+t} \frac{\sum_{k=1}^t \delta_{X_k}}{t}.$$

We interpret this conditional distribution as follows. The term $\frac{\sum_k \delta_{X_k}}{t}$ is the empirical distribution and represents the contribution of experience. The term π represents the initial guess. Thus the conditional distribution of the next observation is a weighed average of the initial guess and the empirical distribution. The relative weights $\frac{\alpha}{\alpha+t}$ and $\frac{t}{\alpha+t}$ balance out our confidence on the prior beliefs versus the data.

The imprecise Dirichlet process. We extend the Dirichlet process to allow for ambiguity, arising from the lack of information about the law of the underlying process, as well as from the incompleteness of the observed data. Specifically, we construct a class of Dirichlet processes by allowing the base measure π to vary in the set of probability measures.

From the perspective of the individual, in every period $t \geq 0$, prior uncertainty regarding the probability law of the underlying stochastic process is expressed by the following class of Dirichlet priors:

$$\widehat{\Lambda}_t = \{DP(\alpha, \pi) : \pi \in \Delta(\widehat{S}_t)\},$$

for some fixed $\alpha > 0$. That is, $\widehat{\Lambda}_t$ is the class of Dirichlet processes with common parameter α that allows the base measure π to be any element of the set of probability measures $\Delta(\widehat{S}_t)$. For a fixed concentration parameter $\alpha > 0$, the set $\widehat{\Lambda}_t$ combines all Dirichlet priors that the individual can conceive in a given period t , thus representing her unconditional belief assessment over the conceivable state space \widehat{S}_t under the veil of ignorance, before making any observations. We note that the base measure π in the precise Dirichlet process $DP(\alpha, \pi)$ is interpreted as the initial guess. Therefore, the set $\widehat{\Lambda}_t$ reflects complete ignorance about the law of the underlying process, since the set of initial guesses is maximal: every conceivable (nontrivial) event $E \subseteq \widehat{S}_t$ has upper probability of 1 and lower probability of 0 in that set.

In subsequent periods $t \geq 1$, the individual updates these unconditional assessments $\widehat{\Lambda}_t$ in light of the history of observations \widehat{x}_t . However, when reassessing her beliefs, the individual's level of awareness increases over time. With the wisdom of hindsight, she can conceive of many sample paths that are consistent with her observations \widehat{x}_t . She then updates her belief assessment conservatively, allowing for all those possible sample paths.

More precisely, in period t , if the individual's conceivable state space is \widehat{S}_t , then she can conceive of all histories of length t , that is,

$$h = (h_1, \dots, h_t) \in \times_{k=1}^t \widehat{S}_t,$$

with $h_k \in \widehat{S}_t$ for every $k = 1, \dots, t$. Given observations \widehat{x}_t , we let $\widehat{H}_t \subseteq \times_{k=1}^t \widehat{S}_t$ denote

the set of histories that are *conceivable* in period t and *consistent* with the finite history of observations \widehat{x}_t , that is,

$$\widehat{H}_t = \left\{ (h_1, \dots, h_t) \in \times_{k=1}^t \widehat{S}_t : h_k|_{A_k} = (\widehat{x}_t)_k \text{ for } k = 1, \dots, t \right\}.$$

Notice that, since the state space may expand over time, usually there will be more than one history consistent with \widehat{x}_t . In particular, for every $t \geq 2$, if at least two consequences are known, $|\widehat{C}_t| \geq 2$, and the individual becomes aware of at least one new action in period t , so that $A_{t-1} \subset A_t$, then there will be at least two different histories in \widehat{H}_t . Moreover, regardless of the cardinality of the set \widehat{H}_t , the last observation $(\widehat{x}_t)_t$ is $x_t|_{A_t}$, which is an element of \widehat{S}_t . Thus $h_t = x_t|_{A_t}$ for every $h = (h_1, \dots, h_t)$ in \widehat{H}_t .

The following theorem describes the individual's posterior inferences in the presence of the ambiguity generated by growing awareness.

Theorem 3.3. *Let (X_1, \dots, X_t) be i.i.d. samples from Θ and suppose Θ has one of the distributions $DP(\alpha, \pi)$ in $\widehat{\Lambda}_t$. Let also \widehat{H}_t denote the set of histories consistent with the observations \widehat{x}_t generated by Θ . Then:*

1. *The set of conditional distributions of Θ , given the histories \widehat{H}_t , is the set of Dirichlet processes:*

$$\widehat{\Lambda}_t | \widehat{H}_t = \left\{ DP \left(\alpha + t, \frac{\alpha}{\alpha + t} \pi + \frac{1}{\alpha + t} \sum_k \delta_{h_k} \right) : \pi \in \Delta(\widehat{S}_t) \text{ and } h \in \widehat{H}_t \right\}.$$

2. *The set of predictive distributions, given the set of histories \widehat{H}_t , is*

$$\mathbb{P}[X_{t+1}|_{A_t} | \widehat{H}_t] = \left\{ \frac{\alpha}{\alpha + t} \pi + \frac{1}{\alpha + t} \sum_{k=1}^t \delta_{h_k} : \pi \in \Delta(\widehat{S}_t) \text{ and } h \in \widehat{H}_t \right\}.$$

Remark 3.4. Henceforth, we write $\mathbb{P} | \widehat{H}_t$ for the set of predictive distributions $\mathbb{P}[X_{t+1}|_{A_t} | \widehat{H}_t]$ above.

Having observed \widehat{x}_t , the individual can conceive of many finite histories $h \in \widehat{H}_t$. Theorem 3.3 says that she updates every prior in $\widehat{\Lambda}_t$ given every sample history consistent with her observations.

Notice that the influence of the probability measures in $\widehat{\Lambda}_t$, the individual's unconditional assessment, on her conditional assessment $\widehat{\Lambda}_t \mid \widehat{H}_t$ is determined by the hyperparameter α and the length of the history of observations she has seen, that is, t . We shall interpret α as the learning parameter. Since the weight given to the empirical distributions is $t/(\alpha+t)$, higher values of α are associated with a lower degree of confidence of the individual on the accumulated data.⁷

We also note that the degree of imprecision of $\widehat{\Lambda}_t \mid \widehat{H}_t$ comes from two sources. The first, not surprisingly, relates to the lack of prior information to guide the choice of the base measure π of the (unconditional) Dirichlet prior $DP(\alpha, \pi)$. The second arises from the lack of familiarity with the newly discovered events, since the discovery of new actions in a subsequent period t increases the number of elements in the set of conceivable histories \widehat{H}_t of length t that are consistent with the observed data \widehat{x}_t .

The set of predictive distributions $\mathbb{P} \mid \widehat{H}_t$ represents the conditional assessment the individual has about the next observation. For every event $E \subseteq \widehat{S}_t$, this conditional assessment can be summarized by the upper probability

$$\bar{\rho}_t(E) = \max\{\rho(E) : \rho \in \mathbb{P} \mid \widehat{H}_t\}$$

and the lower probability

$$\underline{\rho}_t(E) = \min\{\rho(E) : \rho \in \mathbb{P} \mid \widehat{H}_t\}.$$

Notice that if, for some $k \leq t$, $h_k \in E$ for every $h \in \widehat{H}_t$, then $\underline{\rho}_t(E) > 0$. That is, upon making the observation that event E has occurred at least once unambiguously, for every conceivable history consistent with the observations, the individual's revised beliefs regarding E are bounded away from zero.

⁷We refer to Walley [1996] and Raiffa and Schlaifer [1961] for discussions about possible choices of α .

Remark 3.5. Given Theorem 3.3, it readily follows that for every pair of periods $t > t'$:

$$\mathbb{P} \mid \widehat{H}_t = \left\{ \left(\frac{\alpha + t'}{\alpha + t} \right) \pi + \left(\frac{t - t'}{\alpha + t} \right) \sum_{k=t'+1}^t \frac{\delta_{h_k}}{t - t'} : h \in \widehat{H}_t, \pi \in \Delta(\widehat{S}_t) \text{ for which} \right. \\ \left. \text{there exists } \nu \in \mathbb{P} \mid \widehat{H}_{t'} \text{ such that } \pi \left(\varphi_{t:t'}^{-1}(s) \right) = \nu(s) \text{ for every } s \in \widehat{S}_{t'} \right\}.$$

So, in particular, by taking $t' = t-1$ we obtain the following equivalent recursive definition for the individual's predictive assessment in period t :

$$\mathbb{P} \mid \widehat{H}_t = \left\{ \left(\frac{\alpha + t - 1}{\alpha + t} \right) \pi + \frac{\delta_{x_t | A_t}}{\alpha + t} : \pi \in \Delta(\widehat{S}_t) \text{ for which} \right. \\ \left. \text{there exists } \nu \in \mathbb{P} \mid \widehat{H}_{t-1} \text{ such that } \pi \left(\varphi_{t:t-1}^{-1}(s) \right) = \nu(s) \text{ for every } s \in \widehat{S}_{t-1} \right\}.$$

4 Properties of the model

In this section, we study the properties of the model. We first provide a characterization of the updating rule that generates the predictive distribution in Theorem 3.3–2. We then study the large sample properties of the individual's belief assessment.

4.1 Internal consistency

We begin by noting that in any period $t \geq 0$, if the decision maker's conceivable state space is S , then she can conceive of all future sample paths of any finite length n , that is,

$$f = (f_1, \dots, f_n) \in \times_{k=1}^n S,$$

with $f_k \in S$ for every $k = 1, \dots, n$. Let F_S^n denote the set of all such conceivable future sample paths of length n and set $F_S = \cup_{n \geq 1} F_S^n$. In addition, let \widehat{S}_∞ denote the set of conceivable state spaces the individual will become aware of as a result of the observations she will make. That is, $\widehat{S}_\infty = \{\widehat{S}_t : t \geq 0\}$.

Adopting the nomenclature of Billot et al. [2005], we define a *probabilistic belief in period t* to be a mapping $\rho_t^S : \Delta(S) \times F_S \rightarrow \Delta(S)$. Given a belief $\pi \in \Delta(S)$, we interpret the

probabilistic belief $\rho_t^S(\pi; f)$ as telling us how the individual anticipates that such a belief π in period t would be revised as a result of observing the conceivable future sample path $f \in F_S$.

We consider the following axioms on the family of probabilistic beliefs $\{\rho_t^S : t \geq 0, S \in \widehat{S}_\infty\}$.

Axiom 1 (RESPONSIVENESS). For every belief $\pi \in \Delta(S)$ and every future sample path $f = (f_1, \dots, f_n) \in F_S$:

$$\{f_1, \dots, f_n\} \subseteq \text{supp } \rho_t^S(\pi; f) \subseteq \text{supp } \pi \cup \{f_1, \dots, f_n\}.$$

Axiom 2 (SYMMETRIC TREATMENT [OF PREVIOUSLY UNCONSIDERED STATES]). For every pair of state spaces S and S' and every pair of beliefs $\pi \in \Delta(S)$, $\pi' \in \Delta(S')$:

$$\rho_t^S(\pi; s)(s) = \rho_t^{S'}(\pi'; s')(s'),$$

whenever $s \notin \text{supp } \pi$ and $s' \notin \text{supp } \pi'$.

Axiom 3 (INVARIANCE). For every belief $\pi \in \Delta(S)$ and every future sample path $f = (f_1, \dots, f_n) \in F_S$:

$$\rho_t^S(\pi; f) = \rho_t^S(\pi; \sigma f),$$

where $\sigma f = (f_{\sigma(1)}, \dots, f_{\sigma(n)})$ and $\sigma: \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ is a permutation of $\{1, \dots, n\}$.

Axiom 4 (LINEARITY IN BELIEFS). For every pair of beliefs $\pi, \pi' \in \Delta(S)$ and every future sample path $f \in F_S$:

$$\rho_t^S(\beta\pi + (1 - \beta)\pi'; f) = \beta\rho_t^S(\pi; f) + (1 - \beta)\rho_t^S(\pi'; f),$$

for every $\beta \in [0, 1]$.

Axiom 5 (INTERTEMPORAL COHERENCE). For every pair of periods $t < t'$, every belief $\pi \in \Delta(S)$ and every future sample path $f = (f_1, \dots, f_n) \in F_S^n$ with $n > t' - t$:

$$\rho_t^S(\pi; f) = \rho_{t'}^S \left(\rho_t^S(\pi; f_1, \dots, f_{t'-t}); f_{t'-t+1}, \dots, f_n \right).$$

As the name suggests, Axiom 1 (responsiveness), requires that ρ_t^S assigns strictly positive weight to every state that appears in the future sample path f . Furthermore, a state s cannot be in the support of $\rho_t^S(\pi; f)$ if it is neither in the support of the belief π nor in the future sample path f .

Axiom 2 (symmetric treatment), may be interpreted as requiring, should the decision-maker observe next period a previously unconsidered state, that the (strictly positive) weight that she anticipates being shifted to this previously unconsidered state depends only on the period t in which this probabilistic belief is formed.

Axiom 3 (invariance) is an exchangeability property: future information is interpreted in the same way, regardless of the order in which it arrives. It reflects the conditional exchangeability property of the underlying data generating process.

To explain the normative appeal of Axiom 4 (linearity in beliefs), recall that any initial belief π may be expressed as the probability weighted sum

$$\pi = \sum_{s \in S} \pi(s) \delta_s .$$

That is, π is a weighted sum of the Dirac probability measures associated with each element in its support. Noting that $\rho_t^S(\delta_s; f)$ is the probabilistic belief the individual anticipates she would have starting from the degenerate belief concentrated on the state s after observing the sample path f , we see from the repeated application of Axiom 4 that

$$\rho_t^S(\pi; f) = \sum_s \pi(s) \rho_t^S(\delta_s; f) .$$

That is, the probabilistic belief $\rho_t^S(\pi; f)$ can be expressed as a mixture of the beliefs in the set $\{\rho_t^S(\delta_s; f) : s \in S\}$ with the weights corresponding to the weight she assigns in her belief π to each particular state s .

Our final axiom, Axiom 5 (intertemporal coherence), is a consistency property. It ensures that the family of probabilistic beliefs exhibits an appropriate “law of iterative conditioning.”

Together, these five axioms characterize the anticipated revision of beliefs being undertaken in accordance with a Dirichlet process.

Theorem 4.1. *Suppose that $|\widehat{S}_t| \geq 3$ for some $t \geq 0$. Then the following are equivalent.*

1. *The family of probabilistic beliefs $\{\rho_t^S : t \geq 0, S \in \widehat{S}_\infty\}$, satisfies Axioms 1–5 (responsiveness, symmetric treatment, invariance, linearity in beliefs and intertemporal coherence).*
2. *In every period $t \geq 0$ and for each $S \in \widehat{S}_\infty$ the probabilistic belief ρ_t^S , takes the following form: for each belief $\pi \in \Delta(S)$, each future sample path $f = (f_1, \dots, f_n) \in F_S$*

$$\rho_t^S(\pi; f) = \frac{\alpha + t}{\alpha + t + n} \pi + \frac{n}{\alpha + t + n} \sum_{k=1}^n \frac{1}{n} \delta_{f_k},$$

where $\alpha \geq 0$.

For any family of probabilistic beliefs satisfying the five axioms above, Theorem 4.1 implies that the individual's conditional assessment in period t as specified in part 2 of Theorem 3.3 may be reexpressed in terms of the probabilistic belief $\rho_0^{\widehat{S}_t}$ as follows.

$$\mathbb{P} \mid \widehat{H}_t = \left\{ \rho_0^{\widehat{S}_t}(\pi; h) : h \in \widehat{H}_t, \pi \in \Delta(\widehat{S}_t) \right\}.$$

That is, we can view $\mathbb{P} \mid \widehat{H}_t$ as the set of probability beliefs the individual in period 0 would anticipate her assessment in period t would be, if she had been aware in period 0 of what her conceivable state space in period t would be, as well as what were all the conceivable histories in \widehat{H}_t that would be consistent with her observations up to and including period t .

Alternatively, employing the second of the equivalent expressions for $\mathbb{P} \mid \widehat{H}_t$ from Remark 3.5, we obtain the recursive formulation given by:

$$\mathbb{P} \mid \widehat{H}_t = \left\{ \rho_{t-1}^{\widehat{S}_t}(\pi; x_t|_{A_t}) : \pi \in \Delta(\widehat{S}_t) \text{ for which there exists } \nu \in \mathbb{P} \mid \widehat{H}_{t-1} \right. \\ \left. \text{such that } \pi \left(\varphi_{t:t-1}^{-1}(s) \right) = \nu(s) \text{ for every } s \in \widehat{S}_{t-1} \right\}.$$

We conclude this section by comparing our probability assignment rule with the similarity-based probability assignment rule introduced and axiomatized by Billot et al. [2005]. Formally, the intersection of our probability assignment rule and theirs is the limiting case in

which our learning parameter α is equal to zero and their similarity weighting function is constant. That is, the case in which all the elements of any “database” (the analog of our conceivable histories) in their setting are (considered) equally similar. As there is no increasing awareness in their model, their probability assignment is always precise.

4.2 External consistency

We now investigate whether, as the sample size increases, the decision maker’s assessment regarding known events converges in some meaningful sense to the true posterior probability of the event, according to the data generating process.

The ambient sample space $(\Omega, \mathcal{F}, \mu)$ describes every possible outcome of all sources of randomness. However, from the perspective of the individual, the particular choice of this space is immaterial. The individual only cares about the sequence of induced distributions defined by

$$\mu_t(B) = \mu(X_t \in B)$$

for every B in the σ -algebra \mathcal{B} of Borel subsets of $\mathbb{C}^{\mathbb{A}}$.

Furthermore, the individual’s perception of these distributions is restricted by her current level of awareness, which has two implications. First, the individual can conceive of the idea of the unknown, that is, she is cognizant of her unawareness. She can thus conceive of random variables that may take the intangible value θ_t in period t . Let $\tilde{\mathbb{C}} = \mathbb{C} \cup \{\theta_0, \theta_1, \dots\}$ denote the set of extended consequences and consider the set $\tilde{\mathbb{C}}^{\mathbb{A}}$ of functions from \mathbb{A} into $\tilde{\mathbb{C}}$. Endow $\tilde{\mathbb{C}}$ with the σ -algebra \mathcal{C} generated by the union of the Borel σ -algebra on \mathbb{C} and the discrete σ -algebra on $\{\theta_0, \theta_1, \dots\}$. Let $\tilde{\mathcal{B}} = \sigma(B^a : a \in \mathbb{A})$ denote the σ -algebra of subsets of $\tilde{\mathbb{C}}^{\mathbb{A}}$ generated by the cylinder sets

$$B^a = \text{proj}_a^{-1}(B)$$

for every set B in \mathcal{C} and every action $a \in \mathbb{A}$. We shall abuse notation and identify the random variable X_t , which takes values in $\mathbb{C}^{\mathbb{A}}$, with the coextension of X_t that takes values in the larger range space $\tilde{\mathbb{C}}^{\mathbb{A}}$.

Second, in every period, the individual can only observe the realized consequences associated with the actions in the set known to her. Because the individual can only learn the distributions of consequences induced by the actions already discovered by her, learning is not uniform. To deal with this issue, take as given period t' . The set $A_{t'}$ consists of all actions known to the individual in the given period t' . Consider the set $\widetilde{\mathcal{C}}^{A_{t'}}$ of functions from $A_{t'}$ into $\widetilde{\mathcal{C}}$ and endow it with $\mathcal{B}_{t'} = \sigma(B^a : a \in A_{t'})$, that is, the σ -algebra of subsets of $\widetilde{\mathcal{C}}^{A_{t'}}$ generated by the cylinder sets

$$B^a = \text{proj}_a^{-1}(B)$$

for every set B in \mathcal{C} and every $a \in A_{t'}$. If $g: \widetilde{\mathcal{C}}^{A_{t'}} \rightarrow \mathbb{R}$ is a bounded, $\mathcal{B}_{t'}$ -measurable function, then the individual's (conditional) assessment of g in period $t \geq t'$ is given by the collection of conditional expectations

$$\mathbb{E}_\mu \left[\sum g(\psi_t(X_{t+1}|A_t)|_{A_{t'}}) \rho[\psi_t(X_{t+1}|A_t)] \mid \mathcal{G}_t \right]$$

for every $\rho \in \mathbb{P} \mid \widehat{H}_t$. There is learning if these assessments become closer to the true expectation $\mathbb{E}_\mu[g(X_{t+1}|A_{t'}) \mid \mathcal{G}_t]$ as the sample size increases.⁸

Our main result, Theorem 4.2, formalizes these ideas. It conveys the sense in which the individual learns in the long run. Essentially, it says that the limiting distribution induced over consequences by the actions in $A_{t'}$ and the individual's assessment of this distribution become indistinguishable from the point of view of integrating bounded, measurable functions.

If $(\Omega, \mathcal{F}, \mu)$ is a measure space and $(X_t)_{t \geq 1}$ is a sequence of random variables with values in an arbitrary measurable space, then we can define the following notion of countably generated sets of converging functions.

Consider a family \mathbb{G} of real-valued, measurable functions such that for every $g \in \mathbb{G}$

⁸We note that the function g could be any bounded, measurable function. In particular, if g is an indicator function on an event, then the individual learns the probability of that event. If g is a payoff function over the consequences of a particular action, then the individual learns the expected payoff of that action.

there exists a random variable V_g such that

$$\mathbb{E}_\mu[g(X_{t+1}) \mid \mathcal{G}_t] \rightarrow V_g \quad \mu\text{-a.s.}$$

We say that \mathbb{G} is *countably generated* whenever it contains a countable subset \mathbb{G}^0 such that

$$\mathbb{E}_\mu[g(X_{t+1}) \mid \mathcal{G}_t] \rightarrow V_g \quad \mu\text{-a.s.}$$

uniformly for every $g \in \mathbb{G}$ if and only if the convergence happens uniformly for every $g \in \mathbb{G}^0$.

Theorem 4.2. *Fix some $t' \geq 1$. If $(X_t)_{t \geq 1}$ is conditionally exchangeable and \mathbb{G} is a countably generated family of bounded, $\mathcal{B}_{t'}$ -measurable functions $g: \tilde{\mathcal{C}}^{A_{t'}} \rightarrow \mathbb{R}$, then for every $g \in \mathbb{G}$ there exists a random variable V_g such that $\mathbb{E}_\mu[g(X_{t+1}|_{A_{t'}}) \mid \mathcal{G}_t] \rightarrow V_g$ μ -a.s. and*

1. $\mathbb{E}_\mu \left[\max_{\rho \in \mathbb{P}|\hat{H}_t} \sum g(\psi_t(X_{t+1}|_{A_t})|_{A_{t'}}) \rho[\psi_t(X_{t+1}|_{A_t})] \mid \mathcal{G}_t \right] \rightarrow V_g \quad \mu\text{-a.s. and}$
2. $\mathbb{E}_\mu \left[\min_{\rho \in \mathbb{P}|\hat{H}_t} \sum g(\psi_t(X_{t+1}|_{A_t})|_{A_{t'}}) \rho[\psi_t(X_{t+1}|_{A_t})] \mid \mathcal{G}_t \right] \rightarrow V_g \quad \mu\text{-a.s.}$

A useful feature of this model is that it is possible to measure the degree of imprecision or ambiguity of the individual's beliefs for every event she is aware of. We say that the *degree of ambiguity* associated with an event $E \subseteq \hat{S}_t$ in period t is the difference between the upper and the lower probabilities:

$$\Delta_t(E) = \bar{\rho}_t(E) - \underline{\rho}_t(E).$$

The following corollary is an easy consequence of Theorem 4.2. It establishes that ambiguity is associated with lack of familiarity or information regarding conceivable events. There is resolution of ambiguity over time, but it is only partial, in the sense that the individual sees the events she just became aware of as more ambiguous.

Corollary 4.3. *Fix some $t' \geq 1$ and let $E \subseteq \hat{S}_{t'}$. Then for $t > t'$,*

$$\Delta_t(\varphi_{t':t}(E)) \rightarrow 0 \quad \mu\text{-a.s.}$$

5 Merging of opinions

This section considers the problem of a group of agents who need to make a collective decision involving the (uncertain) value of possible plans of action. Initially, their beliefs may not be consistent, but they will have some opportunity to jointly gather evidence regarding the stochastic outcome of each plan of action available before they make a final decision.

The group consensus formation problem has been studied extensively. Among Bayesians, there is a lot of controversy regarding which is the best procedure for aggregating beliefs—many desirable properties lead to dictatorial aggregation rules (see Genest and Zidek [1986] for a survey of classical aggregation methods and impossibility theorems as well as Mongin [1995]). Shafer [1986] and Walley [1991, Chapter 4] propose belief aggregation methods, based on theories of imprecise probability, that seem to perform better and avoid dictatorial rules [see also Crès, Gilboa, and Vieille, 2011]. Going beyond the realm of belief aggregation and into the territory of multi-Bayesian decision problems does not help. These impossibility results carry through if one incorporates preferences into the problem—as illustrated by Hylland and Zeckhauser [1979]. For surveys, we refer to Weerahandi and Zidek [1981, 1983], and Zidek [1988]; Gilboa, Samet, and Schmeidler [2004] is a rare positive result in this area.

We focus on the aggregation rule determined by unanimity voting, whereby an action is acceptable if and only if it is acceptable for every individual in the group. We show that the beliefs of the members of the group merge with increasing information, as a corollary of Theorem 4.2.

Consider a group of n individuals who must reach a consensus regarding the value of each act or plan of action in a common, finite set of possible actions \mathbb{A} . Suppose that each individual i has a prior assessment of the expected value of each action in \mathbb{A} . For the purposes of this section, the choice of the set \mathbb{A} is arbitrary. For example, it could happen that only actions “observable” by all individuals are considered. In this case, \mathbb{A} is the intersection of the individual sets of actions known by each expert, that is, $\mathbb{A} = \bigcap_{i=1}^n A^i$.

Alternatively, it could also happen that the individuals are set out to find a cooperative solution to a complex problem, about which each expert has only a partial understanding. In this case, they could have engaged in a pre-play communication and agreed on the set that represents the combined knowledge of all in the group. Then the appropriate description of the problem is to take \mathbb{A} to be the union of all individual sets of feasible acts, thus $\mathbb{A} = \cup_{i=1}^n A^i$. We also assume that there is a common set \mathbb{C} of consequences and that the group agrees with respect to the return associated with each of these consequences. That is, there exists a bounded and measurable function $v: \mathbb{C} \rightarrow \mathbb{R}$ that expresses the common value assigned by all individuals to each consequence in \mathbb{C} . Let $C_0 \subseteq \mathbb{C}$ denote the finite subset of consequences known to the experts in period 0.

We assume that each individual in the group is Bayesian and makes inferences from the observed data by updating a Dirichlet process prior. In period 0, expert i 's assessment is represented by a Dirichlet prior on $\mathbb{C}^{\mathbb{A}}$ with base measure π_0^i , with support on the set $C_0^{\mathbb{A}}$, and common concentration parameter $\alpha > 0$.⁹ The interpretation is that π_0^i represents the initial guess of expert i and α represents the experts' unanimous confidence in their initial estimation.

The evolution of the experts assessment goes as follows. In period 0, expert i 's belief is represented by the probability measure $\pi_0^i \in \Delta(\mathbb{C}^{\mathbb{A}})$. For every period $t \geq 1$, after the group publicly observes the sample $X_1 = x_1, \dots, X_t = x_t$, expert i updates her assessment by computing the predictive probability under the Dirichlet posterior, that is,

$$\mathbb{P}_t^i[X_{t+1} \in \cdot \mid X_1, \dots, X_t] = \frac{\alpha}{\alpha + t} \pi_0^i[\cdot] + \frac{1}{\alpha + t} \sum_{k=1}^t \delta_{X_k}(\cdot).$$

The group's belief is then represented by the set of predictive probabilities

$$\mathbb{P}_t^G \mid X_1, \dots, X_t = \left\{ \mathbb{P}_t^i[X_{t+1} \in \cdot \mid X_1, \dots, X_t] : i = 1, \dots, n \right\}.$$

Notice that the group's beliefs $\mathbb{P}_t^G \mid X_1, \dots, X_t$ is a subset of the set of predictive distributions $\mathbb{P} \mid \widehat{H}_t$ described in Theorem 3.3, when we let $A_t = \mathbb{A}$ for every period t .

⁹The assumption that the experts share a common concentration parameter α is made for simplicity and it could be relaxed.

Furthermore, each individual i evaluates action a by the expectation

$$\mathbb{E}_t^i[a \mid X_1, \dots, X_t] = \sum_{x \in \text{supp } \mathbb{P}_t^i} \mathbb{P}_t^i[x \mid X_1, \dots, X_t] v \circ a(x).$$

Therefore, the consensus of the group at t is represented by the lower and upper expectations. That is, the aggregated opinion of the group is given by the pair of functionals $\underline{E}_t: \mathbb{C}^A \rightarrow \mathbb{R}$ and $\bar{E}_t: \mathbb{C}^A \rightarrow \mathbb{R}$, defined by

$$\underline{E}_t(a) = \min \{ \mathbb{E}_t^i[a \mid X_1, \dots, X_t] : i = 1, \dots, n \},$$

and

$$\bar{E}_t(a) = \max \{ \mathbb{E}_t^i[a \mid X_1, \dots, X_t] : i = 1, \dots, n \}.$$

The following result follows from Theorem 4.2. It establishes that the group's lower and upper expectations merge as the experts revise their opinions in view of the information gathered over time.

Corollary 5.1. *If $(X_t)_{t \geq 1}$ is conditionally exchangeable, then for every action a*

$$|\bar{E}_t(a) - \underline{E}_t(a)| \rightarrow 0 \quad \mu\text{-a.s.}$$

Corollary 5.1 provides, to some extent, circumstances under which the common prior assumption may be justified. It implies that the posterior beliefs of different agents merge, after observing a sufficiently long history of past, public signals. The common prior could be understood as the limiting posterior of agents who have been observing the realization of a common signal for a sufficiently long time before the economic interaction being studied begins.

It is worth noting that the argument proposed here avoids, at least partially, the shortcomings of alternative arguments in support of the common prior assumption.¹⁰ From a Bayesian statistics point of view, even in the single-person decision problem, agents need to assign positive probability to the true parameter. In the multi-person problem, these

¹⁰For a survey of such arguments, we refer to Morris [1995].

results require that individuals agree on which events should be assigned positive probability (mutual absolute continuity condition).¹¹ These assumptions seem barely weaker than the common prior assumption itself. By applying the consistency results of our learning model, however, it is possible to show convergence of posterior beliefs even when the individuals' prior beliefs are mutually singular.

The common prior assumption has also been justified from a frequentist point of view, by the argument that past experiences wash away differences in beliefs, because limiting relative frequencies can be commonly learned. Corollary 5.1 has the same flavor. It should be noted, however, that contrary to most of the results available, our convergence results do not require stationarity of the underlying stochastic process.¹²

6 Bayesian interpretation

So far, the discussion focused on the classical Bayesian view that there is an objective model of probability which explains the data. According to this view, there is a true but unknown parameter that is to be estimated from the data. The prior assessments in $\widehat{\Lambda}_t$ thus represent initial guesses about this parameter. Consistency matters to classical Bayesians because they would like the posterior to converge in a meaningful way to the true objective parameter as data accumulates.

An alternative view is the subjective view of probability, according to which there is no such thing as an objective probability model. For a subjective Bayesian, probabilities are nothing but a representation of degrees of belief. Here, we provide an interpretation of the model that is compatible with this view.

Usually, an individual has some information about a statistical problem—perhaps the order of magnitude of some parameter or some qualitative aspect of sampling. However, there is little reason to believe that an individual should have much confidence in a sharply

¹¹The absolutely continuous condition was suggested by Blackwell and Dubins [1962]. Kalai and Lehrer [1993] formalized this argument in a game theoretic setting, in which learning leads to Nash equilibrium.

¹²See, for example, Kurz [1994].

defined prior distribution, or that individuals with different backgrounds should agree on all minute details of the model that explains the data. In particular, it is much easier for subjective Bayesians to reach an agreement about qualitative features of the process, such as conditional exchangeability, or hyper-parameters, such as the concentration parameter α , than have a consensus about the whole prior distribution. However, especially in high-dimensional problems, there is no guarantee that the opinions of individuals with different subjective priors would eventually merge, no matter how much data they have. From this point of view, consistency represents asymptotic interpersonal agreement. Indeed, a model with nice frequentist properties is robust in the sense that small variations in the specification of the prior will not lead to large disagreements.

7 Discussion

The analysis presented in this paper has a number of limitations and poses questions that are yet to be addressed. A central question has to do with whether or not the learning process in the model can be extended to a model of preference updating over time, in line with the axiomatization proposed by Kopylov [2016].

Our analysis assumes conditional exchangeability of the underlying stochastic process. This guarantees asymptotic convergence of the learning process. The question not addressed in the paper is whether or not conditional exchangeability can be weakened to the case in which observed data are Cesàro summable. The relationship between Cesàro summable sequences and exchangeable processes is well understood – see for example Kingman [1978, Section 3-(c)]. Such an approach will have the added advantage of providing a non-probabilistic perspective on learning.

One of the open questions has to do with how the decision maker foresees future sample paths, having observed past samples. How do we incorporate to the model a decision maker that has the ability to theorize about plausible future sample paths? The philosophical approach we have in mind for this is Zabell’s conception of “unanticipated knowledge”

[Zabell, 1992, 2005]. In our model, the decision maker can only conceive counterfactual future experiments based on past observations. The limitation of this approach is that the decision maker has no theory in mind for the data generating process.

Finally, the model presented here allows for “awareness of unawareness” of consequences. However, it does not incorporate the possibility of “awareness of unawareness” regarding future actions. We believe that our model can be generalized to allow for the possibility of learning “unanticipated actions” that are revealed over time. The main difficulty is the interpretation of asymptotic results with such generalization. In particular, how does one define and interpret a limiting process of observed available actions? We think this is a potentially interesting avenue for future research.

Appendix

A Proof of Subsection 3.2

- Proof of Theorem 3.3.*
1. From Theorem 3.2 (1), the set $\widehat{\Lambda} \mid \widehat{H}_t$ is the set of all conditional distributions of Θ that follows a $DP(\alpha, \pi)$ distribution, for every $\pi \in \Delta(\widehat{S}_t)$ and every finite sample history consistent with the observations \widehat{x}_t .
 2. From Theorem 3.2 (2), the set $\mathbb{P}[X_{t+1}|_{A_t} \mid \widehat{H}_t]$ is the set of all predictive distributions of the next observation, computed for every finite sample history consistent with the observations \widehat{x}_t .

□

B Proof of Subsection 4.1

Proof of Theorem 4.1. [1 \Rightarrow 2]

Fix some $S \in \widehat{S}_\infty$ for which $|S| \geq 3$. Denote by s_1, s_2 and s_3 three distinct states in S .

Applying Axiom 1 (responsiveness), we have

$$\{s_2\} \subseteq \text{supp } \rho_t^S(\delta_{s_1}; s_2) \subseteq \{s_1, s_2\}$$

Thus it must be the case that for some $\beta_t^1 \in [0, 1)$,

$$\rho_t^S(\delta_{s_1}; s_2) = \beta_t^1 \delta_{s_1} + (1 - \beta_t^1) \delta_{s_2}.$$

From repeated applications of Axiom 2 (symmetric treatment) we obtain:

$$\rho_t^S(\delta_{s'}; s) = \beta_t^1 \delta_{s'} + (1 - \beta_t^1) \delta_s, \text{ for all pairs of states } s', s \in S.$$

Next by repeatedly applying Axiom 4 (linearity in beliefs) it follows that for all $p \in \Delta(S)$ and all $s \in S$

$$\begin{aligned} \rho_t^S(\pi; s) &= \rho_t^S\left(\sum_{s' \in S} \pi(s') \delta_{s'}; s\right) \\ &= \sum_{s' \in S} \pi(s') \rho_t^S(\delta_{s'}; s) \\ &= \sum_{s' \in S} \pi(s') [\beta_t^1 \delta_{s'} + (1 - \beta_t^1) \delta_s] \\ &= \beta_t^1 \sum_{s' \in S} \pi(s') \delta_{s'} + (1 - \beta_t^1) \delta_s \sum_{s' \in S} \pi(s') \\ &= \beta_t^1 \pi + (1 - \beta_t^1) \delta_s. \end{aligned}$$

By invoking Axiom 5 (intertemporal coherence) and utilizing what we have established above, we see that

$$\begin{aligned} \rho_t^S(\delta_{s_1}; s_2, s_3) &= \rho_{t+1}^S\left(\rho_t^S(\delta_{s_1}; s_2); s_3\right) \\ &= \beta_{t+1}^1 [\beta_t^1 \delta_{s_1} + (1 - \beta_t^1) \delta_{s_2}] + (1 - \beta_{t+1}^1) \delta_{s_3} \\ &= \beta_{t+1}^1 \beta_t^1 \delta_{s_1} + \beta_{t+1}^1 (1 - \beta_t^1) \delta_{s_2} + (1 - \beta_{t+1}^1) \delta_{s_3} \end{aligned}$$

Similarly,

$$\rho_t^S(\delta_{s_1}; s_3, s_2) = \beta_{t+1}^1 \beta_t^1 \delta_{s_1} + \beta_{t+1}^1 (1 - \beta_t^1) \delta_{s_3} + (1 - \beta_{t+1}^1) \delta_{s_2}$$

Since by Axiom 3 (invariance), $\rho_t^S(\delta_{s_1}; s_2, s_3) = \rho_t^S(\delta_{s_1}; s_3, s_2)$, and $\{\delta_s : s \in S\}$ forms a basis for $\Delta(S)$, it follows that

$$\beta_{t+1}^1(1 - \beta_t^1) = (1 - \beta_{t+1}^1) \Rightarrow \beta_{t+1}^1 = \frac{1}{2 - \beta_t^1} \quad (1)$$

So, by setting $\alpha := \beta_0^1/(1 - \beta_0^1) \geq 0$, and applying expression (1) repeatedly, we obtain

$$\beta_0^1 = \frac{\alpha}{\alpha + 1}, \beta_1^1 = \frac{\alpha + 1}{\alpha + 2}, \dots, \beta_t^1 = \frac{\alpha + t}{\alpha + t + 1}$$

Thus,

$$1 - \beta_{t+1}^1 = 1 - \frac{\alpha + t + 1}{\alpha + t + 2} = \frac{1}{\alpha + t + 2},$$

which plugging back into the expression for $\rho_t^S(\delta_{s_1}; s_2, s_3)$, yields

$$\rho_t^S(\delta_{s_1}; s_2, s_3) = \frac{\alpha + t}{\alpha + t + 2} \delta_{s_1} + \frac{2}{\alpha + t + 2} \left(\frac{1}{2} \delta_{s_2} + \frac{1}{2} \delta_{s_3} \right).$$

Furthermore, by repeated applications of Axiom 4 (linearity in beliefs) we have

$$\rho_t^S(\pi; s, s') = \frac{\alpha + t}{\alpha + t + 2} \pi + \frac{2}{\alpha + t + 2} \left(\frac{1}{2} \delta_s + \frac{1}{2} \delta_{s'} \right),$$

for all $\pi \in \Delta(S)$ and all $s, s' \in S$.

We extend this result by induction. We assume that the representation holds for any conceivable future sample path up to length $n - 1$. So in particular,

$$\rho_t^S(\pi; f_1, \dots, f_{n-1}) = \beta_t^{n-1} \pi + (1 - \beta_t^{n-1}) \sum_{k=1}^{n-1} \delta_{f_k},$$

where $\beta_t^{n-1} = (\alpha + t)/(\alpha + t + n - 1)$.

Applying Axiom 5 (intertemporal coherence), we obtain

$$\begin{aligned} \rho_t^S(\delta_s; f_1, \dots, f_n) &= \rho_{t+n}^S \left(\rho_t^S(\delta_s; f_1, \dots, f_{n-1}); f_n \right) \\ &= \beta_{t+n-1}^1 \left(\beta_t^{n-1} \delta_s + (1 - \beta_t^{n-1}) \sum_{k=1}^{n-1} \delta_{f_k} \right) + (1 - \beta_{t+n-1}^1) \delta_{f_n} \end{aligned}$$

By equating coefficients,

$$\begin{aligned} \beta_t^n &= \beta_{t+n-1}^1 \times \beta_t^{n-1} \\ &= \frac{\alpha + t + n - 1}{\alpha + t + n} \times \frac{\alpha + t}{\alpha + t + n - 1} \\ &= \frac{\alpha + t}{\alpha + t + n} \end{aligned}$$

and

$$1 - \beta_{t+n-1}^1 = \beta_{t+n-1}^1 (1 - \beta_t^{n-1}) = \frac{1}{\alpha + t + n}.$$

Finally, repeated applications of Axiom 4 (linearity in beliefs) yields

$$\rho_t^S(\pi; f_1, \dots, f_n) = \frac{\alpha + t}{\alpha + t + n} \pi + \frac{n}{\alpha + t + n} \sum_{k=1}^n \frac{1}{n} \delta_{f_k},$$

as required.

The feature that the parameter α does not depend on the conceivable state space S readily follows from Axiom 2 (symmetric treatment).

[2 \Rightarrow 1] The proof is straightforward and is omitted. □

C Proofs of Subsection 4.2

To prove the main result in Subsection 4.2, we need two lemmas. The first essentially shows that the empirical average of any bounded and measurable function computed from some fixed period t' onwards gets closer to the conditional expectation of that function. The second lemma establishes that the decision maker's assessment of that function converges to the limiting conditional expectation.

Lemma C.1. *Take some $t' \geq 1$ and let $g: \mathbb{C}^A \rightarrow \mathbb{R}$ be a bounded, \mathcal{B} -measurable function. If $(X_t)_{t \geq 1}$ is conditionally exchangeable, then there exists a random variable V_g such that*

$$\mathbb{E}_\mu[g(X_{t+1}) \mid \mathcal{G}_t] \rightarrow V_g \quad \mu\text{-a.s.}$$

and for $t \geq t'$

$$\frac{1}{t - t' + 1} \sum_{k=t'}^t \mathbb{E}_\mu[g(X_{k+1}) \mid \mathcal{G}_k] \rightarrow V_g \quad \mu\text{-a.s.}$$

Proof of Lemma C.1. Fix $t' \geq 1$ and let $g: \mathbb{C}^A \rightarrow \mathbb{R}$ be a \mathcal{B} -measurable function. By Berti, Pratelli, and Rigo [2004, Lemma 2.1], there exists a random variable V_g such that $\mathbb{E}_\mu[g(X_{t+1}) \mid \mathcal{G}_t] \rightarrow V_g$ μ -almost surely. For $t \geq t'$, define

$$Y_t = \sum_{k=t'}^t \left\{ \frac{V_g - \mathbb{E}_\mu[V_g \mid \mathcal{G}_k]}{k - t' + 1} \right\}.$$

The sequence $(Y_t)_{t \geq t'}$ is a uniformly integrable martingale with respect to $(\mathcal{G}_t)_{t \geq t'}$ and thus converges μ -almost surely. Taking into account Berti, Pratelli, and Rigo [2004, Lemma 2.1], an application of Kronecker's lemma gives

$$\begin{aligned} & \left| \mathbb{E}_\mu[g(X_{t+1}) \mid \mathcal{G}_t] - \frac{1}{t - t' + 1} \sum_{k=t'}^t \mathbb{E}_\mu[g(X_{k+1}) \mid \mathcal{G}_k] \right| \\ &= \frac{1}{t - t' + 1} \left| \sum_{k=t'}^t \{V_g - \mathbb{E}_\mu[V_g \mid \mathcal{G}_k]\} \right| \rightarrow 0 \quad \mu\text{-a.s.} \end{aligned}$$

This concludes the proof of the lemma. \square

Lemma C.2. *Take some $t' \geq 1$ and let $g: \widetilde{\mathcal{C}}^{A_{t'}} \rightarrow \mathbb{R}$ be a bounded, $\mathcal{B}_{t'}$ -measurable function. If $(X_t)_{t \geq 1}$ is conditionally exchangeable, then there exists a random variable V_g such that $\mathbb{E}_\mu[g(X_{t+1}|_{A_{t'}}) \mid \mathcal{G}_t] \rightarrow V_g$ μ -a.s. and*

1. $\mathbb{E}_\mu \left[\max_{\rho \in \mathbb{P}|\widehat{H}_t} \sum g(\psi_t(X_{t+1}|_{A_t})|_{A_{t'}}) \rho[\psi_t(X_{t+1}|_{A_t})] \mid \mathcal{G}_t \right] \rightarrow V_g \quad \mu\text{-a.s.}$
2. $\mathbb{E}_\mu \left[\min_{\rho \in \mathbb{P}|\widehat{H}_t} \sum g(\psi_t(X_{t+1}|_{A_t})|_{A_{t'}}) \rho[\psi_t(X_{t+1}|_{A_t})] \mid \mathcal{G}_t \right] \rightarrow V_g \quad \mu\text{-a.s.}$

Proof. For every $t \geq t'$, define

$$Z_t = \sum_{k=t'}^t \left\{ \frac{g(X_k|_{A_{t'}}) - \mathbb{E}_\mu[g(X_{k+1}|_{A_{t'}}) \mid \mathcal{G}_k]}{k - t' + 1} \right\}.$$

The sequence $(Z_t)_{t \geq t'}$ is a uniformly integrable martingale with respect to $(\mathcal{G}_t)_{t \geq t'}$ and, thus, converges μ -almost surely to a random variable Z . Furthermore, if $t \geq t'$, then $h_k|_{A_{t'}} = h'_k|_{A_{t'}} = X_k|_{A_{t'}}$ for every $k \geq t'$ and $h, h' \in \widehat{H}_t$. Thus,

$$\begin{aligned} & \max_{\rho \in \mathbb{P}|\widehat{H}_t} \mathbb{E}_\rho[g(\psi_t(X_{t+1}|_{A_t})|_{A_{t'}})] = \max_{\rho \in \mathbb{P}|\widehat{H}_t} \sum g(\psi_t(X_{t+1}|_{A_t})|_{A_{t'}}) \rho[\psi_t(X_{t+1}|_{A_t})] \\ &= \frac{\alpha + t' - 1}{\alpha + t} \frac{\max_{h \in \widehat{H}_t} \left\{ \sum_{k=1}^{t'-1} g(h_k|_{A_{t'}}) \right\}}{t' - 1} + \frac{t - t' + 1}{\alpha + t} \frac{\sum_{k=t'}^t g(h_k|_{A_{t'}})}{t - t' + 1} \\ &= \frac{\alpha + t' - 1}{\alpha + t} \frac{\max_{h \in \widehat{H}_t} \left\{ \sum_{k=1}^{t'-1} g(h_k|_{A_{t'}}) \right\}}{t' - 1} \\ &\quad - \frac{\alpha + t' - 1}{\alpha + t} \frac{\sum_{k=t'}^t g(h_k|_{A_{t'}})}{t - t' + 1} + \frac{\sum_{k=t'}^t g(h_k|_{A_{t'}})}{t - t' + 1}. \end{aligned}$$

Therefore, if $M_g > 0$ is the upper bound of $|g|$, then

$$\begin{aligned} & \left| \mathbb{E}_\mu \left[\max_{\rho \in \widehat{\mathcal{H}}_t} \mathbb{E}_\rho [g(\psi_t(X_{t+1}|A_t)|A_{t'})] \mid \mathcal{G}_t \right] - \frac{1}{t-t'+1} \sum_{k=t'}^t \mathbb{E}_\mu [g(X_{k+1}|A_{t'}) \mid \mathcal{G}_k] \right| \\ & \leq \frac{2(\alpha+t'-1)}{\alpha+t} M_g + \left| \frac{1}{t-t'+1} \sum_{k=t'}^t (k-t'+1) Z_t \right|. \end{aligned}$$

An application of the Kronecker's lemma yields that

$$\left| \frac{1}{t-t'+1} \sum_{k=t'}^t (k-t'+1) Z_t \right| \rightarrow 0 \quad \mu\text{-a.s.}$$

and, because $\frac{2(\alpha+t'-1)}{\alpha+t} \rightarrow 0$ surely, the desired convergence for g follows from Lemma C.1.

The same argument applied to case (2) completes the proof. \square

We are ready to prove Theorem 4.2.

Proof of Theorem 4.2. For each bounded, $\mathcal{B}_{t'}$ -measurable function $g: \mathbb{C}^{A_{t'}} \rightarrow \mathbb{R}$, let $N_g \in \mathcal{F}$ denote the set with $\mu(N_g) = 0$ outside of which the convergence shown in Lemma C.2 does not happen. Because \mathbb{G} is countably generated, there exists a countable subset \mathbb{G}^0 of bounded, measurable functions on $\widetilde{\mathbb{C}}^{A_{t'}}$ such that the convergence is uniform for every $g \in \mathbb{G}$ if and only if it happens for every $g \in \mathbb{G}^0$. Noting that $\mu(\cup_{g \in \mathbb{G}} N_g) = 0$ completes the proof. \square

References

- Benavoli, A., F. Mangili, F. Ruggeri, and M. Zaffalon (2015). Imprecise Dirichlet process with application to the hypothesis test on the probability that $X \leq Y$. *Journal of Statistical Theory and Practice* 9(3), 658–684.
- Berti, P., L. Pratelli, and P. Rigo (2004). Limit theorems for a class of identically distributed random variables. *The Annals of Probability* 32(3), 2029–2052.
- Bewley, T. F. (2002). Knightian decision theory. Part I. *Decisions in Economics and Finance* 25(2), 79–110.

- Billot, A., I. Gilboa, D. Samet, and D. Schmeidler (2005). Probabilities as similarity-weighted frequencies. *Econometrica* 73(4), 1125–1136.
- Blackwell, D. and L. Dubins (1962). Merging of opinions with increasing information. *The Annals of Mathematical Statistics* 33(3), 882–886.
- Boole, G. (1854). *An Investigation of the Laws of Thought: On which are Founded the Mathematical Theories of Logic and Probabilities*. George Boole’s Collected Logical Works. Walton and Maberly.
- Crès, H., I. Gilboa, and N. Vieille (2011). Aggregation of multiple prior opinions. *Journal of Economic Theory* 146(6), 2563–2582.
- Dominiak, A. and G. Tserenjigmid (2017). Does growing awareness increase ambiguity? *Mimeo*.
- Epstein, L. G. and M. Schneider (2007). Learning under ambiguity. *The Review of Economic Studies* 74(4), 1275–1303.
- Escobar, M. D. and M. West (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association* 90(430), 577–588.
- Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *The Annals of Statistics* 1(2), 209–230.
- Ferguson, T. S. (1974). Prior distributions on spaces of probability measures. *The Annals of Statistics* 2(4), 615–629.
- Genest, C. and J. V. Zidek (1986). Combining probability distributions: a critique and an annotated bibliography. *Statistical Science* 1(1), 114–135.
- Ghosh, J. K. and R. V. Ramamoorthi (2003). *Bayesian Nonparametrics*. Secaucus, US: Springer.

- Gilboa, I., F. Mcheroni, M. Marinacci, and D. Schmeidler (2010). Objective and subjective rationality in a multiple prior model. *Econometrica* 78(2), 755–770.
- Gilboa, I., D. Samet, and D. Schmeidler (2004). Utilitarian aggregation of beliefs and tastes. *Journal of Political Economy* 112(4), 932–938.
- Halpern, J. Y., N. Rong, and A. Saxena (2010). Mdps with unawareness. *Proceedings of the Twenty-Sixth Conference on Uncertainty in AI*, 228–235.
- Hylland, A. and R. Zeckhauser (1979). The impossibility of Bayesian group decision making with separate aggregation of beliefs and values. *Econometrica* 47(6), 1321–1336.
- Kalai, E. and E. Lehrer (1993). Rational learning leads to nash equilibrium. *Econometrica* 61(5), 1019–1045.
- Kallenberg, O. (1988). Spreading and predictable sampling in exchangeable sequences and processes. *The Annals of Probability* 16(2), 508–534.
- Karni, E. and D. Schmeidler (1991). Utility theory with uncertainty. In W. Hildenbrand and H. Sonnenschein (Eds.), *Handbook of Mathematical Economics*, pp. 1763–1831. New York, NY: Elsevier Science.
- Karni, E. and M.-L. Vierø (2013). “Reverse Bayesianism”: a choice-based theory of growing awareness. *American Economic Review* 103(7), 2790–2810.
- Karni, E. and M.-L. Vierø (2017). Awareness of unawareness: a theory of decision making in the face of ignorance. *Journal of Economic Theory* 168, 301–328.
- Kingman, J. F. C. (1978). Uses of exchangeability. *The Annals of Probability* 6(2), 183–197.
- Kochov, A. (2016). A behavioral definition of unforeseen contingencies. *Mimeo*.
- Kopylov, I. (2016). Subjective probability, confidence, and Bayesian updating. *Economic Theory* 62(4), 635–658.

- Kurz, M. (1994). On the structure and diversity of rational beliefs. *Economic Theory* 4(6), 877–900.
- Marinacci, M. (2002). Learning from ambiguous urns. *Statistical Papers* 43(1), 143–151.
- Mongin, P. (1995). Consistent Bayesian Aggregation. *Journal of Economic Theory* 66, 313–351.
- Morris, S. (1995). The common prior assumption in economic theory. *Economics and Philosophy* 11(02), 227–253.
- Nehring, K. (2009). "imprecise probabilistic belief as a context for decision-making under ambiguity". *Journal of Economic Theory* 144, 1054–1091.
- Raiffa, H. and R. Schlaifer (1961). *Applied Statistical Decision Theory*. Division of Research, Graduate School of Business Administration, Harvard University.
- Shafer, G. (1986). *A Mathematical Theory of Evidence*. Princeton University Press.
- Walker, S. G., P. Damien, P. W. Laud, and A. F. Smith (1999). Bayesian nonparametric inference for random distributions and related functions. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 61(3), 485–527.
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Number 42 in Monographs on Statistics and Applied Probability. Chapman and Hall.
- Walley, P. (1996). Inferences from multinomial data: learning about a bag of marbles. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 58(1), 3–57.
- Weerahandi, S. and J. V. Zidek (1981). Multi-Bayesian statistical decision theory. *Journal of the Royal Statistical Society. Series A (General)* 144(1), 85–93.
- Weerahandi, S. and J. V. Zidek (1983). Elements of multi-Bayesian decision theory. *The Annals of Statistics* 11(4), 1032–1046.

Zabell, S. L. (1992). Predicting the unpredictable. *Synthese* 90(2), 205–232.

Zabell, S. L. (2005). *Symmetry and its discontents: essays on the history of inductive probability*. Cambridge University Press.

Zidek, J. V. (1988). Group decision analysis and its application to combining opinions. *Journal of Statistical Planning and Inference* 20(3), 307–325.