Games and Economic Behavior 109 (2018) 616–624

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

# Learning the fundamentals in a stationary environment ☆

Nabil I. Al-Najjar *, Eran Shmaya

*Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston, IL 60208, United States*

ARTICLE INFO

*Article history:*
Received 16 December 2016
Available online 16 March 2018

*JEL classification:*
C61
D83

*Keywords:*
Learning
Merging
Stationarity

ABSTRACT

A Bayesian agent relies on past observations to learn the structure of a stationary process. We show that the agent's predictions about near-horizon events become arbitrarily close to those he would have made if he knew the long-run empirical frequencies of the process.



## 1. Introduction

Consider a Bayesian decision maker who observes a stationary stochastic process with a finite set of outcomes. From the ergodic theorem, we know that, in the limit, this observer will learn all the long-run empirical frequencies of the process. Learning from infinite histories, however, is of little value when making decisions based on finite past observations. This paper relates these two perspectives: frequentist learning over infinite horizons and Bayesian predictions based on finite data.

We consider the long-run properties of the *predictive distribution*, defined as the distribution on next period's outcome given the finite number of past observations. We show that, almost surely, the predictive distribution in most periods becomes arbitrarily close to the predictive distribution had the true data generating process been known. Thus, as data accumulates, an observer's predictive distribution based on a finite history becomes nearly as good as what it would have been given knowledge of the objective empirical frequencies over infinite histories. We demonstrate that the various qualifications we impose cannot be dropped.

Our results connect several literatures on learning and predictions in stochastic environments. First, there is the literature on merging of opinions, pioneered by Blackwell and Dubins (1962). They prove that the beliefs of Bayesian observers with mutually absolutely continuous priors will *strongly merge*, in the sense that their posteriors about the infinite future will become arbitrarily close. Motivated by applications for decision making under uncertainty and game theory, Kalai and Lehrer (1994) and Lehrer and Smorodinsky (1996) introduced weaker notions of merging, which focus on closeness of near-horizon predictive distributions. While Blackwell and Dubins' strong merging obtains only under stringent assumptions,

☆ We thank Ehud Kalai, Ehud Lehrer, Rann Smorodinsky and Benjamin Weiss for helpful discussions.
* Corresponding author.
*E-mail address:* al-najjar@northwestern.edu (N.I. Al-Najjar).

https://doi.org/10.1016/j.geb.2018.02.007
0899-8256/

weak merging can be more easily satisfied. In our setting the posteriors typically do not strongly merge with the true parameter, no matter how much data accumulates.

Another line of inquiry focuses on representations of the form $\mu = \int \mu_\theta \, \mathrm{d}\lambda(\theta)$, where the law of the stochastic process $\mu$ is expressed as a convex combination of distributions $\{\mu_\theta\}_{\theta\in\Theta}$ that may be viewed as "simple," or "elementary." Such representations, also called *decompositions*, are useful in models of learning where the set of parameters $\Theta$ may be viewed as the object of learning. Two seminal theorems of this type are de Finetti's representation of exchangeable distributions and the ergodic representation of stationary processes. The ergodic decomposition is the finest decomposition possible using parameters that are themselves stationary. Our main theorem states that a Bayesian decision maker's predictions about a stationary process become arbitrarily close to those he would have made given knowledge of the true ergodic component.

Our result should also be contrasted with Doob's consistency theorem which states that Bayesian posteriors weakly converge to the true parameter. When the focus is on decision making, what matters is not the agents' beliefs about the true parameter but the quality of their predictions. Although the two concepts are related, they are not the same. The difference is seen in the following example (Jackson et al., 1999, Example 5): Assume that the outcomes Heads and Tails are generated by tossing a fair coin. If we take the set of all Dirac measures on infinite sequences of Heads and Tails as "parameters," then the posterior about the parameter converges weakly to a belief that is concentrated on the true realization. On the other hand, the agent's predictions about next period's outcome is constant and never approach the predictions given the true "parameter." This example highlights that convergence of posterior beliefs to the true parameter may have little relevance to an agent's predictions and behavior.

A third related literature, which traces to Cover (1975), concerns the non-Bayesian estimation of stationary processes. See Morvai and Weiss (2005) and the references therein. This literature seeks an algorithm for making predictions about near-horizon events that are accurate for every stationary process. Our proof of Theorem 1 and Example 4 rely on techniques that were developed in this literature. There is however a major difference between that literature and our work: we are interested in a specific algorithm, namely predictions derived from Bayesian updating.

## 2. Decompositions and learning

In this section we recall the notion of ergodic decomposition of a stationary process, interpreted as an agent's belief, and introduce the notion of weak merging which formalizes the idea of learning used in this paper. We also contrast this notion of learning with the more familiar notion of Bayesian consistency which concerns learning the underlying parameter.

### 2.1. Preliminaries

An agent (a decision maker, a player, or a statistician) observes a stochastic process $(\zeta_0, \zeta_1, \zeta_2, \dots)$ where the outcome in each period belongs to some fixed finite set $A$. Time is indexed by $n$ and the agent starts observing the process at $n = 0$. Let $\Omega = A^{\mathbb{N}}$ be the space of *realizations* of the process, with generic element denoted $\omega = (a_0, a_1, \dots)$. Endow $\Omega$ with the product topology and the induced Borel structure $\mathcal{F}$. Let $\Delta(\Omega)$ be the set of probability distributions over $\Omega$. A standard way to represent uncertainty about the process is in terms of an index set of "parameters:"

**Definition 1.** A *decomposition* of $\mu \in \Delta(\Omega)$ is a quadruple $\left(\Theta, \mathcal{B}, \lambda, (\mu_\theta)_{\theta\in\Theta}\right)$ where

- $(\Theta, \mathcal{B}, \lambda)$ is a standard probability space of *parameters*;
- $\mu_\theta \in \Delta(\Omega)$ for every $\theta \in \Theta$;
- for every $S \in \mathcal{F}$, the map $\theta \mapsto \mu_\theta(S)$ is $\mathcal{B}$-measurable and

$$\mu(S) = \int_\Theta \mu_\theta(S)\, \lambda(\mathrm{d}\theta). \tag{1}$$

A decomposition may be thought of as a way for a Bayesian agent to arrange his beliefs. The agent views the process as a two stages randomization: in the first stage, a parameter $\theta$ is chosen according to $\lambda$, and in the second the outcomes are generated according to $\mu_\theta$.

Beliefs can be represented in many ways. The two extreme decompositions are: (1) the *trivial decomposition* with $\Theta = \{\bar\theta\}$, $\mathcal{B}$ is trivial, and $\mu_{\bar\theta} = \mu$; and (2) the *Dirac decomposition* with $\Theta = A^{\mathbb{N}}$, $\mathcal{B} = \mathcal{F}$, and $\lambda = \mu$. A "parameter" in this case is just a Dirac measure $\delta_\omega$ that assigns probability 1 to the realization $\omega$. Note that parameters are not required to be elements of some finite dimensional vector space, as is often assumed in statistical models.

Stationary beliefs admit a well-known decomposition with natural properties. Recall that a stochastic process $(\zeta_0, \zeta_1, \dots)$ is *stationary* if, for every natural number $k$, the joint distribution of the $k$-tuple $\left(\zeta_n, \zeta_{n+1}, \dots, \zeta_{n+k-1}\right)$ does not depend on $n$. A *stationary distribution* is the distribution of a stationary process. The set of stationary distributions over $\Omega$ is convex and compact in the weak*-topology. Its extreme points are called *ergodic distributions*. We denote the set of ergodic distributions by $\mathcal{E}$. Every stationary belief $\mu \in \Delta(\Omega)$ admits a unique decomposition in which the parameter set is the set of ergodic distributions: $\mu = \int \nu \; \lambda(\mathrm{d}\nu)$ for some belief $\lambda \in \Delta(\mathcal{E})$. This decomposition is called *the ergodic decomposition*.