

Noguchi, Yuichi

Article

Merging with a set of probability measures: a characterization

Theoretical Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Noguchi, Yuichi (2015) : Merging with a set of probability measures: a characterization, Theoretical Economics, ISSN 1555-7561, The Econometric Society, New Haven, CT, Vol. 10, Iss. 2, pp. 411-444, <https://doi.org/10.3982/TE1360>

This Version is available at:

<https://hdl.handle.net/10419/150254>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/3.0/>

Merging with a set of probability measures: A characterization

YUICHI NOGUCHI

Department of Economics, Kanto Gakuin University

In this paper, I provide a characterization of a *set* of probability measures with which a prior weakly merges. In this regard, I introduce the concept of *conditioning rules* that represent the regularities of probability measures and define the *eventual generation* of probability measures by a family of conditioning rules. I then show that a set of probability measures is learnable (i.e., all probability measures in the set are weakly merged by a prior) if and only if all probability measures in the set are eventually generated by a *countable* family of conditioning rules. I also demonstrate that quite similar results are obtained with almost weak merging. In addition, I argue that my characterization result can be extended to the case of infinitely repeated games and has some interesting applications with regard to the impossibility result in Nachbar (1997, 2005).

KEYWORDS. Bayesian learning, weak merging, conditioning rules, eventual generation, frequency-based prior.

JEL CLASSIFICATION. C72, C73, D83.

1. INTRODUCTION

Bayesian learning is a learning procedure that has been widely studied in game theory. In particular, Kalai and Lehrer (1993) introduced a learning concept called *merging* to repeated games and showed that if every player's prior belief merges with the probability measure induced by the players' true strategies, then the path of play converges to Nash equilibrium. Merging requires that the updated forecast (i.e., the posterior) about any future events be eventually accurate; the future events include *infinite future* ones, such as tail events. However, when players discount future payoffs in a repeated game, the merging property is more than enough for obtaining convergence to Nash equilibrium. In other words, any information about the infinite future is *not* useful for the discounting players. Accordingly, Kalai and Lehrer (1994) propose a weaker concept of merging called *weak merging*. Weak merging means that the updated forecast about any

Yuichi Noguchi: ynoguchi@kanto-gakuin.ac.jp

I am very grateful to Drew Fudenberg and John H. Nachbar for their encouragement and helpful comments. This paper also benefited from a conversation with H. Peyton Young. Furthermore, I thank Kenichi Amaya and the seminar participants at the 16th International Conference on Game Theory at Stony Brook, the Game Theory Workshop at Hitotsubashi, the Microeconomics Workshop at the University of Tokyo, and the Game Theory Workshop at the University of Tsukuba for their helpful comments. I also wish to thank the co-editor and three anonymous referees for their valuable comments and suggestions that have led to many substantial improvements in this paper. Of course, all errors are mine.

Copyright © 2015 Yuichi Noguchi. Licensed under the Creative Commons Attribution-NonCommercial License 3.0. Available at <http://econtheory.org>.
DOI: 10.3982/TE1360

finite period future event is eventually accurate. Furthermore, the weak merging property is sufficient to deal with learning to play Nash equilibrium (see, e.g., Lehrer and Smorodinsky 1997 and Sandroni 1998). Since then, the literature has mainly focused on weak merging.

In order to give a *general* argument, we consider (a set of) probability measures that represent the evolutions of discrete-time finite-state stochastic processes. That is, given a *finite* set S of states, we focus on (a set of) probability measures over the set $\prod_{t=1}^{\infty} S$ of infinite sequences of states. In this setting, I provide a characterization of a *set* of probability measures with which a prior (belief) “weakly merges.” For that purpose, I introduce the concept of conditioning rules that represent the regularities of probability measures and define the eventual generation of probability measures by a family of conditioning rules. I then show that a set of probability measures is learnable; that is, all probability measures in the set are weakly merged by a prior, if and only if the set is included in a set of probability measures eventually generated by a *countable* family of conditioning rules. In other words, Bayesian learning can eventually make accurate predictions regarding all probability measures eventually generated by a countable family of conditioning rules, but it cannot do so regarding more than those.

The basic ideas of the key concepts, that is, conditioning rules and eventual generation, are easily explained by a simple example. For instance, a second-order Markov probability measure has the regularity that the current probabilities of states are always determined by the states realized in the last two periods. In other words, the current probabilities are conditioned on the states of the last two periods. A conditioning rule captures such a conditioning property of a probability measure so that the regularity of any probability measure is arbitrarily approximated by some conditioning rule. Furthermore, by “a probability measure *eventually generated* by a family of conditioning rules,” I mean that the regularity of the probability measure is (arbitrarily) approximated by one in the family of conditioning rules *from some period on*. As for the above Markov case, a probability measure eventually generated by the second-order Markov conditioning rule means that the current probabilities of states (with respect to the measure) are determined by the states of the last two periods from some period on.

My characterization is particularly important in the context of repeated games. In a repeated game, players sequentially interact with each other and thus tend to be quite uncertain about their opponents’ strategies at the beginning of the game. Therefore, it is natural to start with the assumption that no player knows her opponents’ characteristics except that they play behavior strategies (independently). In this situation, a player would want to use a prior belief that weakly merges with as many strategies of her opponents as possible. Nonetheless, it is not difficult to show that there is *no* prior belief that weakly merges with *all* the opponents’ strategies. Then a fundamental issue arises: the identification of a *learnable set* of the opponents’ strategies, that is, a set of the opponents’ strategies with which a prior belief could weakly merge. Characterizing a learnable set is certainly helpful in clarifying the possibilities of Bayesian learning in repeated games. For example, as Nachbar (1997, 2005) shows, some *diversity* property of learnable sets may be related to the impossibility of learning to play Nash equilibrium. In the last section, I remark that my characterization is related to Nachbar’s impossibility result; see Noguchi (2015) for details.

I provide two results to obtain my characterization. First, I show that *no* prior can weakly merge with more probability measures than those eventually generated by a countable family of conditioning rules. Therefore, *any* learnable set must be included in a set of probability measures eventually generated by a countable family of conditioning rules. Second, and more importantly, I show that for *any* countable family of conditioning rules, there exists a prior such that the prior weakly merges with *all* probability measures eventually generated by the family. This means that if a set of probability measures is included in a set of those eventually generated by a countable family of conditioning rules, then the set is learnable. Therefore, I conclude that a learnable set is characterized by a *countable* family of conditioning rules. Furthermore, I demonstrate that quite similar results are obtained with “almost weak merging.”

The objective of this paper is to demonstrate how to form or construct a prior for obtaining the characterization result. As Gilboa et al. (2004) point out, “Bayesian learning means nothing more than the updating of a given prior. It does not offer any theory, explanation, or insight into the process by which prior beliefs are formed.” Indeed, not many ideas have been provided for the construction of nontrivial priors. As such, I make use of an insight obtained from a study of another learning procedure called *conditional smooth fictitious play* (CSFP): I construct a prior on the basis of conditional empirical frequencies. To be specific, the prior that I construct is a modification of the belief formation process for CSFP presented in Noguchi (2000, 2009). Although the process is based on a quite simple intuitive story of individual learning behavior, it is powerful enough to eventually make accurate predictions regarding as many probability measures as possible. This will be shown later. Furthermore, to prove that the prior works, I use a different mathematical tool than those used in previous works: the theory of *large deviations*, which gives precise probability evaluations about rare events.

Previous work in game theory has mainly explored the conditions on relations between a prior and a (true) probability measure for the prior to (weakly) merge with the measure (e.g., Blackwell and Dubins 1962, Kalai and Lehrer 1993, 1994)¹; Sorin (1999) explains the links to reputation models. The main reason is that if we interpret priors as players’ prior beliefs and a probability measure as that induced by players’ true strategies, then such merging conditions are also conditions for convergence to Nash equilibrium. Clearly, these merging conditions may be helpful in considering learnable sets. For example, the absolute continuity condition implies that any countable set of probability measures is merged by any prior that puts positive probability on each of the probability measures.

In other research fields (e.g., computer sciences, information theory, statistical learning theory, and so on), various issues about (Bayesian) learning have been studied; see, for example, Solomonoff (1978), Ryabko (1988), Cesa-Bianchi and Lugosi (2006), Ryabko and Hutter (2008), and Ryabko (2011) for issues related to this paper.² In particular, Ryabko (2010) provides a general result of learnability with respect to merging:

¹As for other related studies, Jackson et al. (1999) investigate a natural (convex combination) representation of a given prior belief from the viewpoint of learning, and Kalai et al. (1999) explore the relationships between calibration tests and merging.

²The author thanks an anonymous referee for introducing him to works in other research fields.

(under a mild assumption) a set of probability measures with which a prior merges is arbitrarily approximated by its countable subset in a strong sense. I will discuss a merged set of probability measures (by comparing merging with weak merging) in the last section.

This paper is organized as follows. Section 2 gives a description of the basic model and several concepts. Section 3, the main part of the paper, provides a characterization of a set of probability measures with which a prior weakly merges. In Section 4, the characterization result is applied to repeated games. Section 5 provides quite similar results with respect to almost weak merging. Section 6 concludes with several remarks.

2. MODEL

2.1 Basic model and notations

Let S be a *finite* set of states and let a state in S be generically denoted by s . Further, let $\Delta(S)$ designate the set of probability distributions over S . A history is a sequence of states realized each time. I write H_T for the set of all finite histories with length T : $H_T := \prod_{t=1}^T S$. Let H denote the set of all finite histories, including the null history h_0 , that is, $H := \bigcup_{t=0}^{\infty} H_t$, where $h_0 := \emptyset$ and $H_0 := \{h_0\}$. A finite history is denoted by h . When the length of a finite history is emphasized, I write h_T for a finite history up to time T : $h_T := (s_1, \dots, s_T)$. Let \mathbf{H}_∞ designate the set of all infinite histories: $\mathbf{H}_\infty := \prod_{t=1}^{\infty} S$. An infinite history is denoted by $h_\infty := (s_1, s_2, \dots)$. If a finite history h is an initial segment of a (finite or infinite) history h' , then it is denoted by $h \leq h'$. When $h \leq h'$ and $h \neq h'$, it is designated by $h < h'$. I assume the standard measurable structure \mathcal{F} on \mathbf{H}_∞ .³ Let $\{\mathcal{F}_T\}_T$ be the standard filtration: $\mathcal{F}_T \subseteq \mathcal{F}_{T+1} \subseteq \mathcal{F}$ for all T .⁴ Let μ denote a probability measure on $(\mathbf{H}_\infty, \mathcal{F})$. When a probability measure is considered as a prior, it is denoted by $\tilde{\mu}$.

2.2 Weak merging

The focus in this paper is primarily on *weak merging*. Weak merging requires that the updated forecast about any *finite-period* future event be eventually accurate.

DEFINITION 1. A prior $\tilde{\mu}$ weakly merges with a probability measure μ if, for all $k \geq 1$,

$$\lim_{T \rightarrow \infty} \sup_{\mathbf{A} \in \mathcal{F}_{T+k}} |\tilde{\mu}(\mathbf{A} | \mathcal{F}_T) - \mu(\mathbf{A} | \mathcal{F}_T)| = 0, \quad \mu\text{-a.s.}$$

Let $\mu(s|h_T)$ denote the probability of s at time $T + 1$ conditional on a realized past history h_T up to time T , and let $\mu(h)$ denote the probability of h . It is then important to note that $\tilde{\mu}$ weakly merges with μ if and only if the *one period ahead* forecast is eventually correct (see Lehrer and Smorodinsky 1996a): for μ -almost all h_∞ and all $s \in S$,

$$\lim_{T \rightarrow \infty} |\tilde{\mu}(s|h_T) - \mu(s|h_T)| = 0. \quad (1)$$

³The set \mathcal{F} is the minimum σ -algebra including all cylinder sets based on finite histories: $\mathcal{F} := \sigma(\bigcup_{t=0}^{\infty} \{\mathbf{C}_h \mid h \in H_t\})$, where $\mathbf{C}_h := \{h_\infty \mid h < h_\infty\}$.

⁴The set \mathcal{F}_T is the minimum σ -algebra including all cylinder sets based on finite histories with length T : $\mathcal{F}_T := \sigma(\{\mathbf{C}_h \mid h \in H_T\})$.

The purpose of this paper is to characterize a set of probability measures with which a prior weakly merges. Accordingly, I define the phrase “weak merging with a set of probability measures.”

DEFINITION 2. A prior $\tilde{\mu}$ weakly merges with a set M of probability measures if $\tilde{\mu}$ weakly merges with all probability measures in M . Further, a set M of probability measures is weakly merged if there exists a prior $\tilde{\mu}$ such that $\tilde{\mu}$ weakly merges with M .

2.3 Conditional probability systems

I make use of *conditional probability systems* (CPSs). A CPS represents the probability distribution of the current state after each past history (up to the previous time). Formally, a CPS is a mapping from the set H of finite histories to the set $\Delta(S)$ of probability distributions over S . It is denoted by $f : H \rightarrow \Delta(S)$. Then it follows from Kolmogorov’s extension theorem (see, e.g., Shiryaev 1996) that for all f , there exists a *unique* probability measure μ_f such that $\mu_f(s|h) = f(h)[s]$ for all $s \in S$ and all $h \in H$. Conversely, it is easy to see that for all μ , there exists a CPS f_μ such that $f_\mu(h)[s] = \mu(s|h)$ for all $s \in S$ and all $h \in H$.⁵ The correspondence makes it possible to focus on CPSs instead of probability measures. Indeed, from (1) in the previous subsection, $\tilde{\mu}$ weakly merges with μ if and only if for μ -almost all h_∞ ,

$$\lim_{T \rightarrow \infty} \|f_{\tilde{\mu}}(h_T) - f_\mu(h_T)\| = 0,$$

where $\|\cdot\|$ is the maximum norm: $\|x\| := \max_s |x[s]|$.

REMARK 1. If $\tilde{\mu}(\mathbf{A}) = \tilde{\mu}'(\mathbf{A})$ for all $\mathbf{A} \in \mathcal{F}$, then $\tilde{\mu}$ and $\tilde{\mu}'$ are identical as probability measures. However, it is possible that $\tilde{\mu}(s|h) \neq \tilde{\mu}'(s|h)$ for some $s \in S$ and some $h \in H$ with $\tilde{\mu}(h) (= \tilde{\mu}'(h)) = 0$. In this paper, I assume that $\tilde{\mu}$ and $\tilde{\mu}'$ are different as priors in such a case. Therefore, each prior has its unique corresponding CPS.

2.4 Conditioning rules and classes

I introduce a key concept to characterize a learnable set: *conditioning rules* (CRs). A CR represents a (approximate) regularity of a CPS or a probability measure. Formally, a CR is a *finite* partition of H and is denoted by \mathcal{P} . An element of a CR \mathcal{P} is called a *class* in \mathcal{P} and is denoted by α . Note that a class (in \mathcal{P}) is considered a subset of H because it is an element of a partition of H . In the following, I will often define a subset of H and call it a *class* (although the subset may not be an element of any given CR). If a realized history $h_{T-1} \in \alpha$, we say that α is *active at time T* or that time T is an α -*active period*.

For any CPS f , I define its ε -*approximate conditioning rule* (ε -ACR). The definition states that probability distributions (on S) after finite histories in each class α are almost the same.

⁵A CPS f_μ corresponding to μ is not necessarily unique, because if $\mu(h) = 0$, then $\mu(s|h)$ is arbitrary. However, $\mu(s|h)$ is uniquely determined for all $s \in S$ and all $h \in H$ with $\mu(h) > 0$. Thus, for any two CPSs f_μ and f'_μ corresponding to μ , $f_\mu(h)[s] = f'_\mu(h)[s]$ for all $s \in S$ and all $h \in H$ with $\mu(h) > 0$.

DEFINITION 3. Given $\varepsilon \geq 0$, a finite partition $\mathcal{P}_\varepsilon^f$ of H is called an ε -approximate conditioning rule (ε -ACR) of f if for all $\alpha \in \mathcal{P}_\varepsilon^f$ and all $h, h' \in \alpha$, $\|f(h) - f(h')\| \leq \varepsilon$.

If $\varepsilon = 0$, $\mathcal{P}_\varepsilon^f$ is simply called a conditioning rule (CR) of f . Note that any CPS f has its ε -ACR for all $\varepsilon > 0$.⁶

EXAMPLE 1. Let $S := \{L, R\}$ and let f be a first-order Markov CPS such that $f(h_T) := (\frac{1}{3}, \frac{2}{3})$ when $s_T = L$ and $f(h_T) := (\frac{2}{3}, \frac{1}{3})$ when $s_T = R$. Then let $\mathcal{P}^f := \{\alpha_L, \alpha_R\}$, where $\alpha_L := \{h_T \in H \mid s_T = L\}$ and $\alpha_R := \{h_T \in H \mid s_T = R\}$.⁷ Hence, \mathcal{P}^f is an ε -ACR of f for all $\varepsilon > 0$; that is, \mathcal{P}^f is a CR of f . \diamond

Conversely, CRs generate CPSs.

DEFINITION 4. A CPS $f : H \rightarrow \Delta(S)$ is generated by a family \mathbb{P} of CRs if for all $\varepsilon > 0$, there exists $\mathcal{P} \in \mathbb{P}$ such that \mathcal{P} is an ε -ACR of f .

The definition states that for all $\varepsilon > 0$, the regularity of f is ε -approximated by some CR in \mathbb{P} .

EXAMPLE 2. Let $S := \{L, R\}$ and $\mathcal{P} := \{\alpha_L, \alpha_R\}$, where $\alpha_L := \{h_T \in H \mid s_T = L\}$ and $\alpha_R := \{h_T \in H \mid s_T = R\}$. Furthermore, let $\mathcal{Q} := \{\alpha_E, \alpha_O\}$, where $\alpha_E := \{h_T \in H \mid T \text{ is odd}\}$ and $\alpha_O := \{h_T \in H \mid T \text{ is even}\}$. Then $f : H \rightarrow \Delta(S)$ is generated by $\{\mathcal{P}, \mathcal{Q}\}$ if and only if either there exist $0 \leq p, q \leq 1$ such that for all $h \in \alpha_L$, $f(h) = (p, 1 - p)$, and such that for all $h \in \alpha_R$, $f(h) = (q, 1 - q)$, or there exist $0 \leq p', q' \leq 1$ such that for all $h \in \alpha_E$, $f(h) = (p', 1 - p')$, and such that for all $h \in \alpha_O$, $f(h) = (q', 1 - q')$.⁸ \diamond

Note that all independent and identically distributed (i.i.d.) CPSs are generated by any (nonempty) family of CRs.⁹ Further, note that any CPS f is generated by any family $\{\mathcal{P}_{1/n}^f\}_n$ of its $1/n$ -ACRs.

Similarly, CRs also generate probability measures.

DEFINITION 5. A probability measure μ is generated by a family \mathbb{P} of CRs if there exists a CPS f_μ corresponding to μ such that f_μ is generated by \mathbb{P} .

⁶By the compactness of $\Delta(S)$, for all $\varepsilon > 0$, we may take a finite family $\{\Delta_j\}_{j=1}^{j=m}$ of subsets in $\Delta(S)$ such that (i) $\{\Delta_j\}_j$ covers $\Delta(S)$, that is, $\bigcup_{j=1}^{j=m} \Delta_j = \Delta(S)$, and (ii) those diameters are no more than ε , that is, $\sup_{\pi, \pi' \in \Delta_j} \|\pi - \pi'\| \leq \varepsilon$ for all j . Thus, for all CPSs f and all $\varepsilon > 0$, an ε -ACR $\mathcal{P}_\varepsilon^f$ of f is defined by the following equivalence relation on H : for all $h, h' \in H$,

$$h \sim_{\mathcal{P}_\varepsilon^f} h' \quad \text{if and only if} \quad \text{there exists } j \text{ such that } f(h), f(h') \in \Delta_j \text{ and } f(h), f(h') \notin \Delta_k \text{ for all } k < j.$$

⁷The null history h_0 ($:= \emptyset$) may belong to either class.

⁸More generally, for any finite family $\{\mathcal{P}_n\}_{n=1}^{n=N}$ of CRs, f is generated by $\{\mathcal{P}_n\}_{n=1}^{n=N}$ if and only if there exists n such that \mathcal{P}_n is a CR of f .

⁹A CPS f is i.i.d. if $f(h) = f(h')$ for all $h, h' \in H$.

The set of all probability measures generated by \mathbb{P} is denoted by $G(\mathbb{P})$. As in the CPS case, all i.i.d. probability measures are generated by any (nonempty) family of CRs.¹⁰ In addition, any probability measure is generated by a countable family of CRs.

REMARK 2. A useful fact is that if a set \mathcal{S}_m is countable for all $m = 1, 2, \dots$, the union $\bigcup_m \mathcal{S}_m$ is also countable. Hence, any countable set of probability measures $\{\mu_m\}_m$ is generated by a countable family of CRs because each μ_m is generated by a countable family $\{\mathcal{P}_{1/n}^{f_{\mu_m}}\}_n$ of CRs.

The CRs are ordered in *fineness*: if for all $\alpha \in \mathcal{P}$, there (uniquely) exists $\beta \in \mathcal{Q}$ such that $\alpha \subseteq \beta$, \mathcal{P} is finer than \mathcal{Q} (or \mathcal{Q} is coarser than \mathcal{P}), denoted by $\mathcal{Q} \leq \mathcal{P}$. Furthermore, if $\alpha \subseteq \beta$, α is finer than β (or β is coarser than α). It is important to note that a finer CR generates more probability measures; for example, if $\mathcal{Q} \leq \mathcal{P}$, $G(\mathcal{Q}) \subseteq G(\mathcal{P})$. Moreover, note that the partition consisting only of $\{H\}$, that is, $\mathcal{P}_{\text{id}} := \{\{H\}\}$, is the coarsest partition: $\mathcal{P}_{\text{id}} \leq \mathcal{P}$ for any partition \mathcal{P} (of H). For convenience, \mathcal{P}_{id} will be called *the identity partition*. It is easy to see that a CPS (or a probability measure) is i.i.d. if and only if it is generated by \mathcal{P}_{id} . Finally, note the following useful joint property of finite partitions. Given any two finite partitions \mathcal{P} and \mathcal{Q} , let $\mathcal{P} \wedge \mathcal{Q}$ denote the *joint* of \mathcal{P} and \mathcal{Q} : $\mathcal{P} \wedge \mathcal{Q} := \{\alpha \cap \beta \mid \alpha \in \mathcal{P}, \beta \in \mathcal{Q}\}$. Then $\mathcal{P} \wedge \mathcal{Q}$ is also a finite partition, and is finer than \mathcal{P} and \mathcal{Q} .

3. CHARACTERIZATION OF A LEARNABLE SET

3.1 Main result

To characterize a learnable set of probability measures, I slightly extend the generation of probability measures: *eventual generation*. The following definition is rather complicated, but it simply states that for any $\varepsilon > 0$, the regularity of f is (almost surely) ε -approximated by one of the CRs $\{\mathcal{P}_i\}_i$ from some period on.

DEFINITION 6. A CPS $f : H \rightarrow \Delta(S)$ is eventually generated by $\{\mathcal{P}_i\}_i$ if for all $\varepsilon > 0$, there exist an index i_0 , a μ_f -probability 1 set \mathbf{Z}_0 , and a time function $T_0 : \mathbf{Z}_0 \rightarrow \mathbb{N}$ such that for all $\alpha \in \mathcal{P}_{i_0}$ and all $h_T, h_{T'} \in \alpha$, if there exist $h_\infty, h'_\infty \in \mathbf{Z}_0$ wherein $h_T < h_\infty$, $T \geq T_0(h_\infty)$, $h_{T'} < h'_\infty$, and $T' \geq T_0(h'_\infty)$, then $\|f(h_T) - f(h_{T'})\| \leq \varepsilon$.¹¹

Clearly, any CPS generated by $\{\mathcal{P}_i\}_i$ is eventually generated by $\{\mathcal{P}_i\}_i$. As in the case of generation, I define the eventual generation of probability measures. A probability measure μ is eventually generated by $\{\mathcal{P}_i\}_i$ if there exists a CPS f_μ corresponding to μ such that f_μ is eventually generated by $\{\mathcal{P}_i\}_i$. The set of all probability measures eventually generated by $\{\mathcal{P}_i\}_i$ is denoted by $\text{EG}(\{\mathcal{P}_i\}_i)$. As in the CPS case, any probability measure generated by $\{\mathcal{P}_i\}_i$ is eventually generated by $\{\mathcal{P}_i\}_i$: $G(\{\mathcal{P}_i\}_i) \subseteq \text{EG}(\{\mathcal{P}_i\}_i)$. More precisely, $\text{EG}(\{\mathcal{P}_i\}_i)$ is strictly larger than $G(\{\mathcal{P}_i\}_i)$, that is, $G(\{\mathcal{P}_i\}_i) \subsetneq \text{EG}(\{\mathcal{P}_i\}_i)$. Now we can state the main result.

¹⁰A probability measure μ is i.i.d. if a CPS f_μ corresponding to μ is i.i.d.

¹¹Let \mathbb{N} denote the set of natural numbers.

THEOREM 1. *A set M of probability measures is weakly merged if and only if all probability measures in M are eventually generated by some countable family $\{\mathcal{P}_i\}_i$ of CRs, that is, if and only if there exists a countable family $\{\mathcal{P}_i\}_i$ of CRs such that $M \subseteq \text{EG}(\{\mathcal{P}_i\}_i)$.*

For example, the following corollaries are immediate from [Theorem 1](#).

COROLLARY 1. *The set of all i.i.d. probability measures is weakly merged.*

PROOF. As noted in [Section 2.4](#), all i.i.d. probability measures are generated by the identity partition \mathcal{P}_{Id} . From this and [Theorem 1](#), it follows that the set of all i.i.d. probability measures is weakly merged. \square

COROLLARY 2. *The set of all Markov probability measures is weakly merged.*

PROOF. For all $k = 1, 2, \dots$, let \mathcal{M}_k be the k th-order Markov CR.¹² Then $\{\mathcal{M}_k\}_k$ is a countable family of CRs. Since all Markov probability measures are generated by $\{\mathcal{M}_k\}_k$, it follows from [Theorem 1](#) that the set of all Markov probability measures (of all orders) is weakly merged. \square

In [Section 4](#), I provide more examples of the weakly merged set particularly related to repeated games.

3.2 Bounds of weak merging

First, I show that the weak merging property is always bounded by a countable family of CRs. In other words, *no* prior can weakly merge with more probability measures than those eventually generated by a *countable* family of CRs.

PROPOSITION 1. *For any prior $\tilde{\mu}$, there exists a countable family $\{\mathcal{P}_i\}_i$ of CRs such that $\tilde{\mu}$ does not weakly merge with any $\mu \notin \text{EG}(\{\mathcal{P}_i\}_i)$.*

PROOF. Fix any $\tilde{\mu}$. Let $f_{\tilde{\mu}}$ be the CPS corresponding to $\tilde{\mu}$. As noted in [Section 2.4](#), for each n , we may take a $1/n$ -ACR $\mathcal{P}_{1/n}^{f_{\tilde{\mu}}}$ of $f_{\tilde{\mu}}$. I show that $\tilde{\mu}$ does not weakly merge with any $\mu \notin \text{EG}(\{\mathcal{P}_{1/n}^{f_{\tilde{\mu}}}\}_n)$. Take any $\mu \notin \text{EG}(\{\mathcal{P}_{1/n}^{f_{\tilde{\mu}}}\}_n)$. Then there exists $\varepsilon_0 > 0$ such that for all n , all μ -probability 1 sets \mathbf{Z} , and all time functions $T : \mathbf{Z} \rightarrow \mathbb{N}$, there exist $\alpha \in \mathcal{P}_{1/n}^{f_{\tilde{\mu}}}$ and $h_T, h'_{T'} \in \alpha$ wherein $h_T < h_\infty$, $T \geq T(h_\infty)$, $h'_{T'} < h'_\infty$, and $T' \geq T(h'_\infty)$ for some $h_\infty, h'_\infty \in \mathbf{Z}$, and $\|f_\mu(h_T) - f_\mu(h'_{T'})\| > \varepsilon_0$.

Suppose that $\tilde{\mu}$ weakly merges with μ . Then there exists a μ -probability 1 set \mathbf{Z}_0 such that for all $h_\infty \in \mathbf{Z}_0$, there exists $T_0(h_\infty)$ wherein for all $T \geq T_0(h_\infty)$, $\|f_{\tilde{\mu}}(h_T) - f_\mu(h_T)\| \leq \varepsilon_0/4$. On the other hand, letting $n_0 \geq 4/\varepsilon_0$, it follows from the previous paragraph that

¹²The CR \mathcal{M}_k is defined by the following equivalence relation on H : for all $h, h' \in H$,

$$h \sim_{\mathcal{M}_k} h' \quad \text{if and only if} \quad \text{the states of the last } k \text{ periods in } h \text{ are the same as those in } h'.$$

for n_0, \mathbf{Z}_0 , and $T_0 : \mathbf{Z}_0 \rightarrow \mathbb{N}$, there exist $\alpha \in \mathcal{P}_{1/n_0}^{f_{\tilde{\mu}}}$ and $h_T, h'_{T'} \in \alpha$ such that $h_T < h_\infty, T \geq T_0(h_\infty), h'_{T'} < h'_\infty$, and $T' \geq T_0(h'_\infty)$ for some $h_\infty, h'_\infty \in \mathbf{Z}_0$, and $\|f_\mu(h_T) - f_\mu(h'_{T'})\| > \varepsilon_0$. These imply that $\|f_{\tilde{\mu}}(h_T) - f_{\tilde{\mu}}(h'_{T'})\| \geq \varepsilon_0/2$. Then, however, since $\alpha \in \mathcal{P}_{1/n_0}^{f_{\tilde{\mu}}}$ and $h_T, h'_{T'} \in \alpha$, $\|f_{\tilde{\mu}}(h_T) - f_{\tilde{\mu}}(h'_{T'})\| \leq 1/n_0 \leq \varepsilon_0/4$. This is a contradiction. Thus, $\tilde{\mu}$ does not weakly merge with μ . \square

3.3 Frequency-based prior

For any countable family of CRs, I construct a prior that weakly merges with all probability measures eventually generated by the family. To form the prior, I use the method for constructing a belief formation process for “conditional smooth fictitious play” provided in [Noguchi \(2000, 2009\)](#): a prior is defined on the basis of *conditional empirical frequencies*; such a prior will be called a *frequency-based prior*.

PROPOSITION 2. *For any countable family $\{\mathcal{P}_i\}_i$ of CRs, there exists a frequency-based prior $\tilde{\mu}_F$ such that $\tilde{\mu}_F$ weakly merges with all $\mu \in \text{EG}(\{\mathcal{P}_i\}_i)$.*

In the remainder of this subsection, I specify the construction of a frequency-based prior $\tilde{\mu}_F$, which includes a brief explanation of why $\tilde{\mu}_F$ works (for any $\mu \in \text{EG}(\{\mathcal{P}_i\}_i)$). The proof of [Proposition 2](#) is given in [Appendix B](#).

Without loss of generality, we may assume (in the remainder of this subsection) that $\{\mathcal{P}_i\}_i$ is ordered in fineness: $\mathcal{P}_i \leq \mathcal{P}_{i+1}$ for all i ; see [Section 2.4](#).¹³ To construct a frequency-based prior $\tilde{\mu}_F$, I start by determining the *prior sample size* n_0^α for each class $\alpha \in \bigcup_i \mathcal{P}_i$, where $\bigcup_i \mathcal{P}_i$ is the set of all classes in $\{\mathcal{P}_i\}_i$. (This is based on a result of large deviations given in [Appendix A](#).) Suppose that $\alpha \in \mathcal{P}_i$. Then choose any positive integer n_0^α such that

$$n_0^\alpha \geq \frac{i^2}{2} \left(i + \log \left(\frac{\#\mathcal{P}_i}{1 - \exp(-2i^{-2})} \right) \right), \tag{2}$$

where $\#$ denotes the cardinality of a set. Note that this inequality is equivalent to $\#\mathcal{P}_i \sum_{n=n_0^\alpha}^{n=\infty} \exp(-2ni^{-2}) \leq \exp(-i)$, which will be used in the proof of [Proposition 2](#). Taking a larger n_0^α means obtaining more prior samples for class α so as to make the probability of wrong prediction *exponentially* smaller.

Next, I introduce a categorizing rule that classifies observed samples into categories. A *category* is represented by a pair made of an index and a class, that is, (i, α) such that $\alpha \in \mathcal{P}_i$, and a *categorizing rule* is defined by two mappings $i : H \rightarrow \mathbb{N}$ and $\alpha : H \rightarrow \bigcup_i \mathcal{P}_i$. Specifically, $i(\cdot)$ and $\alpha(\cdot)$ represent the following forecaster’s learning behavior. A forecaster first employs all classes in \mathcal{P}_1 , for example, $\{\alpha_1^k\}_k$, to categorize past samples (i.e., realized states): each α_1^k is used as a (temporary) category. That is, for each α_1^k , the empirical frequencies of the samples obtained in α_1^k -active periods are used as the forecasts

¹³For any $\{\mathcal{P}_i\}_i$, let $\mathcal{Q}_i := \bigwedge_{j=1}^{j=i} \mathcal{P}_j$ for all i . Then $\{\mathcal{Q}_i\}_i$ has the following property: $\mathcal{P}_i \leq \mathcal{Q}_i$ for all i and $\mathcal{Q}_i \leq \mathcal{Q}_{i+1}$ for all i .

in α_1^k -active periods. However, if any α_1^k has been active many times so that enough samples have been obtained as prior samples for classes (in \mathcal{P}_2) finer than α_1^k , then the forecaster switches α_1^k to the finer classes (in \mathcal{P}_2); that is, $\{\alpha_2^l\}_l$ such that $\alpha_2^l \subseteq \alpha_1^k$. Otherwise, the forecaster continues using α_1^k . In other words, the forecaster starts to employ each α_2^l as a (temporary) category and uses the empirical frequencies consisting of the prior samples (obtained in the past α_1^k -active periods) and the samples observed in α_2^l -active periods after switching to α_2^l . Again, if the forecaster has obtained enough samples in α_2^l -active periods, then she switches α_2^l to classes (in \mathcal{P}_3) finer than α_2^l ; that is, $\{\alpha_3^m\}_m$ such that $\alpha_3^m \subseteq \alpha_2^l$, and so on.

In this behavior, the forecaster continues switching to finer classes because it enables the forecaster to learn more probability measures. Furthermore, enough prior samples for each class are obtained *before* employing the class as a (temporary) category. It eventually enables the forecaster to make accurate predictions from the first active period (to the last active period) of employing each class (as a category), which means that the forecaster eventually makes accurate predictions (regarding any $\mu \in EG(\{\mathcal{P}_i\}_i)$).

I provide the formal definitions of $i(\cdot)$, $\alpha(\cdot)$, and $\tilde{\mu}_F$. To define $i(\cdot)$ and $\alpha(\cdot)$, I introduce three other functions $m : H \rightarrow \mathbb{N}$, $n : H \rightarrow \mathbb{N}$, and $\beta : H \rightarrow \bigcup_i \mathcal{P}_i$. Roughly, $m(h_T)$ is the maximum index of past employed classes that are coarser than currently employed class $\alpha(h_T)$. Let $n(h_T)$ be the number of times that a currently employed class has been active and let $\beta(h_T)$ be a finer class that will be employed next. I recursively define $i(\cdot)$, $\alpha(\cdot)$, $m(\cdot)$, $n(\cdot)$, and $\beta(\cdot)$ as follows:

- First, let $i(h_0) := 1$ and $\alpha(h_0) := \alpha$, where $h_0 \in \alpha$ and $\alpha \in \mathcal{P}_1$. Furthermore, let $m(h_0) := 1$, $n(h_0) := 0$, and $\beta(h_0) := \alpha$, where $h_0 \in \alpha$ and $\alpha \in \mathcal{P}_{m(h_0)+1}$.
- Suppose that $i(h_t)$, $\alpha(h_t)$, $m(h_t)$, $n(h_t)$, and $\beta(h_t)$ are defined for $0 \leq t \leq T - 1$. Let $m(h_T) := \max\{i(h_t) \mid h_t < h_T, h_T \in \alpha(h_t)\}$,¹⁴ and let $\beta(h_T) := \alpha$, where $h_T \in \alpha$ and $\alpha \in \mathcal{P}_{m(h_T)+1}$. Furthermore, let $n(h_T) := \#\{h_t \mid h_t < h_T, h_T \in \alpha(h_t), i(h_t) = m(h_T)\}$. Then define

$$i(h_T) := \begin{cases} m(h_T) + 1 & \text{if } n(h_T) \geq n_0^{\beta(h_T)} \\ m(h_T) & \text{otherwise,} \end{cases} \tag{3}$$

where $n_0^{\beta(h_T)}$ is the prior sample size for class $\beta(h_T)$. Finally, let $\alpha(h_T) := \alpha$, where $h_T \in \alpha$ and $\alpha \in \mathcal{P}_{i(h_T)}$.

It is important to note that the inequality in the definition of $i(h_T)$ is the switching criterion such that if $n_0^{\beta(h_T)}$ samples are obtained as prior samples for $\beta(h_T)$, then a forecaster switches to a finer class $\beta(h_T)$; otherwise, the forecaster continues using a currently employed class.

Given h_{T-1} , if $i(h_{T-1}) = i$ and $\alpha(h_{T-1}) = \alpha$, time T is an *effective period of category* (i, α) (or category (i, α) is effective at time T). Note that given any h_∞ , each period has exactly one effective category. Further note that any category (i, α) with $i \geq 2$ has its (unique) *predecessor* (i_p, α_p) such that $i_p = i - 1$ and $\alpha \subseteq \alpha_p \in \mathcal{P}_{i_p}$; indeed, (i, α) can be effective only after (i_p, α_p) has been effective n_0^α times. Next, I define the *prior samples*

¹⁴If $\{i(h_t) \mid h_t < h_T, h_T \in \alpha(h_t)\} = \emptyset$, then let $m(h_T) := 1$.

$d_0^{(i,\alpha)}$ and the *prior sample size* $n_0^{(i,\alpha)}$ for each (i, α) . For all (i, α) (with $i \geq 2$), let $d_0^{(i,\alpha)}$ consist of samples observed in the effective periods of its predecessor (i_p, α_p) : each component $d_0^{(i,\alpha)}[s]$ is the number of times that s has occurred in the first n_0^α effective periods of (i_p, α_p) .¹⁵ Thus, $d_0^{(i,\alpha)}$ is *history-dependent*: $d_0^{(i,\alpha)}$ may change according to h_∞ . Let $n_0^{(i,\alpha)} := \sum_s d_0^{(i,\alpha)}[s]$ ($= n_0^\alpha$).

Finally, I define a prior $\tilde{\mu}_F$. Given a realized past history h_T , suppose that a category (i, α) is effective at time $T + 1$: $i = i(h_T)$ and $\alpha = \alpha(h_T)$. Then collect observed states in the past effective periods of (i, α) , which are represented by a vector $d_T^{(i,\alpha)}$: each component $d_T^{(i,\alpha)}[s]$ is the number of times that s has occurred in the past effective periods of (i, α) . Let $n_T^{(i,\alpha)}$ denote the sample size for (i, α) up to time T : $n_T^{(i,\alpha)} := \sum_s d_T^{(i,\alpha)}[s]$. Define the *conditional empirical distribution* $D_T^{(i,\alpha)}$ on (i, α) up to time T as

$$D_T^{(i,\alpha)} := \frac{d_T^{(i,\alpha)} + d_0^{(i,\alpha)}}{n_T^{(i,\alpha)} + n_0^{(i,\alpha)}}.$$

Then use $D_T^{(i,\alpha)}$ as the forecast at time $T + 1$. Accordingly, define *frequency-based CPS* f_F as follows: for all $h_T \in H$, $f_F(h_T) := D_T^{(i,\alpha)}$, where $i = i(h_T)$ and $\alpha = \alpha(h_T)$. Then let $\tilde{\mu}_F := \mu_{f_F}$.¹⁶

For the proof of [Proposition 2](#), see [Appendix B](#).

[Propositions 1](#) and [2](#) induce a characterization of weak merging with a set of probability measures, which is the main result of this paper.

PROOF OF [THEOREM 1](#). Suppose that a set M of probability measures is weakly merged. Then there exists a prior $\tilde{\mu}$ such that $\tilde{\mu}$ weakly merges with all μ in M . However, [Proposition 1](#) then states that there exists a countable family $\{\mathcal{P}_i\}_i$ of CRs such that $\tilde{\mu}$ does not weakly merge with any $\mu \notin \text{EG}(\{\mathcal{P}_i\}_i)$. Therefore, $M \subseteq \text{EG}(\{\mathcal{P}_i\}_i)$. Conversely, suppose that there exists a countable family $\{\mathcal{P}_i\}_i$ of CRs such that $M \subseteq \text{EG}(\{\mathcal{P}_i\}_i)$. [Proposition 2](#) then states that there exists a prior $\tilde{\mu}_F$ such that $\tilde{\mu}_F$ weakly merges with all $\mu \in \text{EG}(\{\mathcal{P}_i\}_i)$. Thus, $\tilde{\mu}_F$ weakly merges with all $\mu \in M$ because $M \subseteq \text{EG}(\{\mathcal{P}_i\}_i)$. Therefore, M is weakly merged. \square

4. APPLICATION TO REPEATED GAMES

4.1 Basic observation

The characterization result in the previous section is applied to an infinitely repeated game with perfect monitoring. Player $i = 1, \dots, I$ takes a pure action a_i from a finite set A_i each time. Specifically, each time, every player observes the history of all the players' (pure) actions taken (up to the previous time) and then chooses her mixed action (i.e., a probability distribution over A_i) independently. That is, each player i plays her behavior strategy, denoted by σ_i . A profile of all players' actions is denoted by $a := (a_i)_i$.

¹⁵For any $\alpha \in \mathcal{P}_1$, we may take $d_0^{(1,\alpha)}$ arbitrarily such that $\sum_s d_0^{(1,\alpha)}[s] = n_0^\alpha$. Then let $n_0^{(1,\alpha)} := \sum_s d_0^{(1,\alpha)}[s]$ ($= n_0^\alpha$).

¹⁶That is, for any $h_T = (s_1, \dots, s_T)$, $\tilde{\mu}_F(h_T) = \mu_{f_F}(h_T) = f_F(h_0)[s_1] \cdot f_F(h_1)[s_2] \cdots f_F(h_{T-1})[s_T]$.

Let A designate the set of all action profiles, that is, $A := \prod_i A_i$. Further, a history of the repeated game is a sequence of all players' actions realized each time. Notations about histories are the same as in Section 2.1; that is, $H_T := \prod_{t=1}^T A$, $H := \bigcup_{t=0}^{\infty} H_t$, and $\mathbf{H}_\infty := \prod_{t=1}^{\infty} A$. Hence, a behavior strategy (for player i) σ_i is formally represented by a mapping from H to $\Delta(A_i)$, where $\Delta(A_i)$ is the set of player i 's mixed actions. Let Σ_i denote the set of all behavior strategies of player i . Moreover, a strategy profile of player i 's opponents is denoted by $\sigma_{-i} := (\sigma_j)_{j \neq i}$. Let Σ_{-i} designate the set of all the opponents' strategy profiles; that is, $\Sigma_{-i} := \prod_{j \neq i} \Sigma_j$. Given a strategy profile $\sigma := (\sigma_1, \dots, \sigma_I)$, I write $\mu(\sigma)$ for the probability measure (on \mathbf{H}_∞) induced by playing σ . Kuhn's theorem for repeated games assures that each player i 's prior belief about her opponents' behavior strategies is identified with a profile of the opponents' behavior strategies, denoted by $\tilde{\rho}^i := (\tilde{\rho}_j^i)_{j \neq i}$, where $\tilde{\rho}_j^i$ is a behavior strategy of player j for each $j \neq i$ (see Aumann 1964, Kalai and Lehrer 1993, and Nachbar 2005, 2009). Note that given a player i 's strategy σ_i , $\mu(\sigma_i, \tilde{\rho}^i)$ weakly merges with $\mu(\sigma_i, \sigma_{-i})$ if and only if for $\mu(\sigma_i, \sigma_{-i})$ -almost all h_∞ ,

$$\lim_{T \rightarrow \infty} \|\tilde{\rho}_j^i(h_T) - \sigma_j(h_T)\| = 0 \quad \text{for all } j \neq i.$$

Moreover, it is important to note that merging with the opponents' strategies depends on the player's own behavior in a repeated game. Accordingly, I explicitly write the player's own strategy in the following definition of learning.

DEFINITION 7. A prior belief $\tilde{\rho}^i$ of player i leads the player to learn to predict her opponents' strategies σ_{-i} with her own strategy σ_i if $\mu(\sigma_i, \tilde{\rho}^i)$ weakly merges with $\mu(\sigma_i, \sigma_{-i})$.

Indeed, the following example shows that, in general, whether or not a set of the opponents' strategies is learnable depends on the player's own strategy.

EXAMPLE 3. Let $A_i := \{L, R\}$ for $i = 1, 2$. For $m = 1, 2, \dots$, let $\sigma_1^{(R,m)}$ be player 1's strategy, in which player 1 takes R for certain from time 1 to time m and then takes L for certain from time $m + 1$ onward. Further, for $n = 1, 2, \dots$, let $\hat{\Sigma}_2^{(R,n)} := \{\sigma_2 \mid \sigma_2(h^{(R,n)}) = (\frac{1}{2}, \frac{1}{2}) \text{ for any } h^{(R,n)}\}$, where $h^{(R,n)}$ denotes any finite history in which player 1 has taken R from time 1 to time n . Then, if $m \geq n$, there are many prior beliefs (of player 1) such that each of them leads the player to learn to predict all $\sigma_2 \in \hat{\Sigma}_2^{(R,n)}$ with $\sigma_1^{(R,m)}$. Otherwise (i.e., if $m < n$), no prior belief leads the player to learn to predict all $\sigma_2 \in \hat{\Sigma}_2^{(R,n)}$ with $\sigma_1^{(R,m)}$. \diamond

Let $\text{LS}(\tilde{\rho}^i, \sigma_i)$ denote the set of all the opponents' strategies that player i 's prior belief $\tilde{\rho}^i$ leads her to learn to predict with her own strategy σ_i :

$$\text{LS}(\tilde{\rho}^i, \sigma_i) := \{\sigma_{-i} \mid \tilde{\rho}^i \text{ leads player } i \text{ to learn to predict } \sigma_{-i} \text{ with } \sigma_i\}.$$

Notice that $\text{LS}(\tilde{\rho}^i, \cdot)$ is considered a *set correspondence* from the set Σ_i of player i 's strategies to the power set $2^{\Sigma_{-i}}$ of her opponents' strategies. This, along with Example 3, means that a set correspondence (rather than a set) is appropriate to capture the learning performance of a prior belief in a repeated game.

DEFINITION 8. A prior belief $\tilde{\rho}^i$ of player i leads the player to learn to predict a set correspondence $\Psi^i : \Sigma_i \rightarrow 2^{\Sigma_{-i}}$ if $\Psi^i(\sigma_i) \subseteq \text{LS}(\tilde{\rho}^i, \sigma_i)$ for all $\sigma_i \in \Sigma_i$.

DEFINITION 9. A set correspondence $\Psi^i : \Sigma_i \rightarrow 2^{\Sigma_{-i}}$ is learnable (by player i) if there exists a prior belief $\tilde{\rho}^i$ of player i such that $\tilde{\rho}^i$ leads her to learn to predict Ψ^i .

As in the CPS case, for any σ_{-i} , I define its ε -approximate conditioning rule (ε -ACR). Given $\varepsilon \geq 0$, a finite partition $\mathcal{P}_\varepsilon^{\sigma_{-i}}$ of H is said to be an ε -approximate conditioning rule (ε -ACR) of σ_{-i} if for all $\alpha \in \mathcal{P}_\varepsilon^{\sigma_{-i}}$ and all $h, h' \in \alpha$, $\|\sigma_j(h) - \sigma_j(h')\| \leq \varepsilon$ for all $j \neq i$. If $\varepsilon = 0$, $\mathcal{P}_\varepsilon^{\sigma_{-i}}$ is simply called a conditioning rule (CR) of σ_{-i} . It is not difficult to show that any σ_{-i} has its ε -ACR for all $\varepsilon > 0$; see footnote 6 in Section 2.4.

Conversely, CRs generate (the opponents') strategies. A strategy profile σ_{-i} of player i 's opponents is generated by a family \mathbb{P} of CRs if for all $\varepsilon > 0$, there exists $\mathcal{P} \in \mathbb{P}$ such that \mathcal{P} is an ε -ACR of σ_{-i} . The set of all the opponents' strategy profiles generated by \mathbb{P} is denoted by $G_{-i}(\mathbb{P})$.

Furthermore, noting that the path of play in a repeated game depends on player i 's own behavior, I define the eventual generation of strategies.

DEFINITION 10. A strategy profile σ_{-i} of player i 's opponents is eventually generated by a family $\{\mathcal{P}_n\}_n$ of CRs with player i 's strategy σ_i if for all $\varepsilon > 0$, there exist an index n_0 , a $\mu(\sigma_i, \sigma_{-i})$ -probability 1 set \mathbf{Z}_0 , and a time function $T_0 : \mathbf{Z}_0 \rightarrow \mathbb{N}$ such that for all $\alpha \in \mathcal{P}_{n_0}$ and all $h_T, h_{T'} \in \alpha$, if there exist $h_\infty, h'_\infty \in \mathbf{Z}_0$ wherein $h_T < h_\infty$, $T \geq T_0(h_\infty)$, $h_{T'} < h'_\infty$, and $T' \geq T_0(h'_\infty)$, then $\|\sigma_j(h_T) - \sigma_j(h_{T'})\| \leq \varepsilon$ for all $j \neq i$.

The set of all the opponents' strategy profiles eventually generated by $\{\mathcal{P}_n\}_n$ with σ_i is denoted by $\text{EG}(\{\mathcal{P}_n\}_n, \sigma_i)$. Obviously, $G_{-i}(\{\mathcal{P}_n\}_n) \subseteq \bigcap_{\sigma_i} \text{EG}(\{\mathcal{P}_n\}_n, \sigma_i)$. As in the CPS case, all i.i.d. strategy profiles (of the opponents)¹⁷ are generated by any (nonempty) family of CRs¹⁸; thus, $G_{-i}(\{\mathcal{P}_n\}_n)$ is always uncountable. In general, even $G_{-i}(\{\mathcal{P}_n\}_n)$ is much larger than any previously known learnable set. In addition, any profile of the opponents' strategies is generated by a countable family of CRs; for example, any σ_{-i} is generated by any family $\{\mathcal{P}_{1/n}^{\sigma_{-i}}\}_n$ of its $1/n$ -ACRs. Thus, any countable set of the opponents' strategy profiles is generated by a countable family of CRs; see Remark 2 in Section 2.4.

Propositions 3 and 4 correspond to Propositions 1 and 2, respectively. In other words, we can take a countable family of CRs to characterize a learnable set correspondence.

PROPOSITION 3. For any prior belief $\tilde{\rho}^i$ of player i , there exists a countable family $\{\mathcal{P}_n\}_n$ of CRs such that for all $\sigma_i \in \Sigma_i$ and all $\sigma_{-i} \notin \text{EG}(\{\mathcal{P}_n\}_n, \sigma_i)$, $\tilde{\rho}^i$ does not lead her to learn to predict σ_{-i} with σ_i .

PROOF. As in the CPS case, $\tilde{\rho}^i$ has a $1/n$ -ACR $\mathcal{P}_{1/n}^{\tilde{\rho}^i}$ for all n . The rest of the argument is the same as in the proof of Proposition 1. \square

¹⁷We say that σ_{-i} is i.i.d. if $\sigma_{-i}(h) = \sigma_{-i}(h')$ for all $h, h' \in H$.

¹⁸As in the CPS case, σ_{-i} is i.i.d. if and only if σ_{-i} is generated by the identity partition \mathcal{P}_{id} .

PROPOSITION 4. For any countable family $\{\mathcal{P}_n\}_n$ of CRs, there exists a frequency-based prior belief $\tilde{\rho}_F^i$ of player i such that for all $\sigma_i \in \Sigma_i$ and all $\sigma_{-i} \in \text{EG}(\{\mathcal{P}_n\}_n, \sigma_i)$, $\tilde{\rho}_F^i$ leads her to learn to predict σ_{-i} with σ_i .

PROOF. The construction of $\tilde{\rho}_F^i$ is just the same as that of the frequency-based CPS f_F except that $D_T^{(i,\alpha)}$ ($= (D_{T,j}^{(i,\alpha)})_{j \neq i}$) are the conditional empirical distributions of the opponents' realized actions. The rest of the argument is similar to the proof of Proposition 2. \square

Propositions 3 and 4 entail the characterization of a learnable set correspondence.

THEOREM 2. A set correspondence $\Psi^i : \Sigma_i \rightarrow 2^{\Sigma_{-i}}$ is learnable (by player i) if and only if there exists a countable family $\{\mathcal{P}_n\}_n$ of CRs such that for all $\sigma_i \in \Sigma_i$, $\Psi^i(\sigma_i) \subseteq \text{EG}(\{\mathcal{P}_n\}_n, \sigma_i)$.

The proof of Theorem 2 is similar to that of Theorem 1.

4.2 Frequency-based belief formation

Frequency-based prior belief $\tilde{\rho}_F^i$ is a belief for conditional smooth fictitious play (CSFP): taking a smooth approximate *myopic* best response to prior belief $\tilde{\rho}_F^i$ is CSFP (see Fudenberg and Levine 1998, 1999 and Noguchi 2000, 2009). Thus, CSFP is interpreted as a Bayesian learning procedure (in the myopic case). Interestingly, Propositions 3 and 4 imply that the class of (prior) beliefs for CSFP weakly merges with *any* learnable set correspondence: the class of beliefs for CSFP has a kind of dominance property of weak merging.

COROLLARY 3. For any prior belief $\tilde{\rho}^i$ of player i , there exists a frequency-based prior belief $\tilde{\rho}_F^i$ of player i such that $\tilde{\rho}_F^i$ leads her to learn to predict $\text{LS}(\tilde{\rho}^i, \cdot)$: for all $\sigma_i \in \Sigma_i$ and all $\sigma_{-i} \in \Sigma_{-i}$, if $\mu(\sigma_i, \tilde{\rho}^i)$ weakly merges with $\mu(\sigma_i, \sigma_{-i})$, then $\mu(\sigma_i, \tilde{\rho}_F^i)$ also weakly merges with $\mu(\sigma_i, \sigma_{-i})$.

By slightly modifying the switching criterion (3) in Section 3.3, Corollary 3 is generalized to have a similar dominance property of *almost weak merging* as well; the modification also enables CSFP to have a sophisticated *no-regret* property. See Section 5 for details.

4.3 Some noteworthy examples

4.3.1 *Uniformly learnable set* I have argued that, in general, learnability crucially depends on the player's own behavior in a repeated game. However, many well known classes of (the opponents') strategies are learnable *regardless of* the player's strategy: they are *uniformly learnable*.

DEFINITION 11. A set S_{-i} of the opponents' strategies is uniformly learnable (by player i) if there exists a prior belief $\tilde{\rho}^i$ of player i such that for all $\sigma_i \in \Sigma_i$ and all $\sigma_{-i} \in S_{-i}$, $\tilde{\rho}^i$ leads the player to learn to predict σ_{-i} with σ_i .

From the argument in Section 4.1, it is evident that the set of i.i.d. strategies is uniformly learnable because it is generated by the identity partition \mathcal{P}_{Id} . Similarly, the set of all Markov strategies (of all orders) is uniformly learnable because it is generated by the (countable) family $\{\mathcal{M}_k\}_k$ of all Markov CRs. More generally, for any countable family $\{\mathcal{P}_n\}_n$ of CRs, the set $G_{-i}(\{\mathcal{P}_n\}_n)$ of the opponents' strategy profiles generated by $\{\mathcal{P}_n\}_n$ is uniformly learnable.

COROLLARY 4. For any countable family $\{\mathcal{P}_n\}_n$ of CRs, $G_{-i}(\{\mathcal{P}_n\}_n)$ is uniformly learnable (by player i).

PROOF. For $\{\mathcal{P}_n\}_n$, construct a frequency-based prior belief $\tilde{\rho}_F^i$, as in the proof of Proposition 4. As noted in Section 4.1, $G_{-i}(\{\mathcal{P}_n\}_n) \subseteq \bigcap_{\sigma_i} \text{EG}(\{\mathcal{P}_n\}_n, \sigma_i)$. This, along with Proposition 4, implies that for all $\sigma_i \in \Sigma_i$ and all $\sigma_{-i} \in G_{-i}(\{\mathcal{P}_n\}_n)$, $\tilde{\rho}_F^i$ leads player i to learn to predict σ_{-i} with σ_i . \square

A typical example of Corollary 4 is the set of *finite automaton strategies*¹⁹; strategies implemented by finite automata have been studied much in the repeated game literature (see, e.g., Osborne and Rubinstein 1994 for details).²⁰ A CR \mathcal{P} is said to be a *finite automaton conditioning rule* (FACR) if for any class $\alpha \in \mathcal{P}$ and any action profile $a \in A$, there exists a (unique) class $\beta \in \mathcal{P}$ such that for any h , if $h \in \alpha$, then $h * a \in \beta$, where $h * a$ is the concatenation of h with a .²¹ (The identity partition \mathcal{P}_{Id} and any Markov CR \mathcal{M}_k are examples of FACRs.) Note that the opponents' strategy profile σ_{-i} is generated by an FACR if and only if σ_{-i} is implemented by a finite automaton.²² Further, the family of all FACRs, denoted by \mathbb{P}_F , is countable.²³ Therefore, we can apply Corollary 4 to $G_{-i}(\mathbb{P}_F)$,

¹⁹The author thanks an anonymous referee and the co-editor for suggesting this important example.

²⁰This paper follows Mailath and Samuelson (2006). A finite automaton consists of a *finite* set of states W , an initial state $w_0 \in W$, an output function $g: W \rightarrow \prod_i \Delta(A_i)$, and a transition function $q: W \times A \rightarrow W$. Then the finite automaton (W, w_0, g, q) implements a (unique) strategy profile σ as follows. First, for the null history h_0 , let $\sigma(h_0) := g(w_0)$. Next, for any history h_1 with length 1, define $\sigma(h_1) := g(w_1)$, where $w_1 = q(w_0, a)$ and $h_1 = a$. Similarly, for any history h_2 with length 2, let $\sigma(h_2) := g(w_2)$, where $w_2 = q(w_1, a^2)$, $w_1 = q(w_0, a^1)$, and $h_2 = (a^1, a^2)$, and so on. Accordingly, we say that the opponents' strategy profile σ_{-i} is implemented by a finite automaton if for some player i 's strategy σ_i , (σ_i, σ_{-i}) is implemented by a finite automaton.

²¹That is, $h * a$ is the finite history such that h occurs first, and then a is realized in the period after h .

²²To any FACR \mathcal{P} , there corresponds a pair (w_0, q) made of an initial state w_0 and a transition function q . Indeed, each class in \mathcal{P} is interpreted as a state, and the initial state w_0 is the class to which the null history belongs. Then, for any $\alpha \in \mathcal{P}$ and any $a \in A$, let $q(\alpha, a) := \beta$, where $h * a \in \beta$ ($\in \mathcal{P}$) for all $h \in \alpha$. (Evidently, q is well defined.) Hence, any strategy profile σ_{-i} generated by \mathcal{P} is implemented by a finite automaton corresponding to \mathcal{P} . Conversely, it is easy to see that any strategy profile σ_{-i} implemented by a finite automaton is also implemented by some finite automaton corresponding to an FACR. Thus, σ_{-i} is generated by the FACR.

²³For each $k = 1, 2, \dots$, the number of transition functions in all automata with k states (i.e., $\#W = k$) is $k^{k\#A}$, and the number of (possible initial) states is k . Hence, the number of all (w_0, q) s is $k \cdot k^{k\#A}$ and

and the set of all (the opponents') strategy profiles implemented by finite automata is uniformly learnable.

As a more general example of $G_{-i}(\{\mathcal{P}_n\}_n)$, we may consider the set of (the opponents') *computably regular strategies*. Computably regular strategies have strong regularities in the sense that their CRs are determined by computer algorithms. A function $\Lambda : H \times H \rightarrow \{0, 1\}$ is called the characteristic function of a partition \mathcal{P} if for all $h, h' \in H$, $h \sim_{\mathcal{P}} h' \Leftrightarrow \Lambda(h, h') = 1$. A CR \mathcal{P} is said to be *computably regular* if its characteristic function is (Turing machine) computable. Let \mathbb{P}_C denote the family of all computably regular CRs. A strategy profile σ_{-i} of the opponents is *computably regular* if σ_{-i} is generated by \mathbb{P}_C . Most practical strategies, including all i.i.d. strategies, all computable pure strategies,²⁴ all Markov strategies of all orders, all strategies implemented by finite automata, and equilibrium strategies in folk theorems,²⁵ are computably regular. In addition, computably regular strategies are interpreted as a generalization of computable pure (behavior) strategies to *mixed* (behavior) strategies. Indeed, any computable pure strategy (profile) is generated by some computably regular CR. Finally, we may apply [Proposition 4](#) (and [Corollary 4](#)) to the set Σ_{-i}^C of (the opponents') computably regular strategies because \mathbb{P}_C is countable and $\Sigma_{-i}^C := G_{-i}(\mathbb{P}_C)$.

COROLLARY 5. *There exists a frequency-based prior belief $\tilde{\rho}_F^i$ of player i such that for all $\sigma_i \in \Sigma_i$ and all $\sigma_{-i} \in \Sigma_{-i}^C$, $\tilde{\rho}_F^i$ leads her to learn to predict σ_{-i} with σ_i .*

REMARK 3. The union of \mathbb{P}_C and any countable family $\{\mathcal{P}_n\}_n$ of CRs is also countable. Thus, we may assume that Σ_{-i}^C is included in a learnable set (correspondence).

4.3.2 Examples for Nachbar's impossibility theorem Nachbar (1997, 2005) provides a celebrated impossibility theorem about Bayesian learning in repeated games. Roughly, the impossibility theorem states that if each player's prior belief leads her to learn to predict her opponents' *diverse* strategies and the learning performances of the players' prior beliefs are *symmetric* to each other, then some player's prior belief *cannot* lead her to learn to predict her opponents' *true* strategies (i.e., her opponents' (approximate) optimal strategies with respect to their prior beliefs) with *all* of her own *diverse* strategies, including her *true* one; recall that the path of play depends on the player's own behavior. In this subsection, I show that for *any* countable family $\{\mathcal{P}_n\}_n$ of CRs, the pair $(G_1(\{\mathcal{P}_n\}_n), G_2(\{\mathcal{P}_n\}_n))$ is an (general) example for Nachbar's impossibility theorem. To demonstrate this, I provide several formal definitions and a version of the impossibility theorem. Following Nachbar (2005), consider a two-player infinitely repeated game. Let

is thus finite. Therefore, the set of (w_0, q) s in all finite automata is countable. Since there is a one-to-one correspondence between FACRs and their corresponding (w_0, q) s, \mathbb{P}_F is also countable.

²⁴For example, see Nachbar and Zame (1996) for computable pure strategies.

²⁵For example, equilibrium strategies in Fudenberg and Maskin (1991) are computably regular provided that players' discount factors, their payoffs in a stage game, and the target values of their averaged discounted payoff sums are computable numbers. In addition, equilibrium strategies (in the perfect monitoring case) in Hörner and Olszewski (2006) are also computably regular if the same computability condition holds. Moreover, a simple strategy profile (and an optimal simple penal code) in Abreu (1988) is computably regular provided that the initial path and the punishments for players are computable sequences.

$\hat{\Sigma}_i$ generically denote a set of player i 's strategies for $i = 1, 2$. First, I impose an appropriate *learnability* condition on a pair $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ of players' strategy sets.

LEARNABILITY. A pair $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ is learnable by a pair $(\tilde{\rho}^1, \tilde{\rho}^2)$ of players' prior beliefs if (a) for any $i \neq j$, any $\sigma_i \in \hat{\Sigma}_i$, and any $\sigma_j \in \hat{\Sigma}_j$, the prior belief $\tilde{\rho}^i$ of player i leads the player to learn to predict σ_j with σ_i , and (b) for any $i \neq j$, any $\sigma_i \in \hat{\Sigma}_i$, any $\sigma_j \in \hat{\Sigma}_j$, and any $h \in H$, if $\mu(\sigma_i, \tilde{\rho}^i)(h) = 0$, then $\mu(\sigma_i, \sigma_j)(h) = 0$.²⁶

Next, the following condition requires $\hat{\Sigma}_1$ and $\hat{\Sigma}_2$ to be diverse, and symmetric to each other.

CSP. A pair $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ satisfies CS (caution and symmetry) if for any $i \neq j$, any pure strategy $s_i \in \hat{\Sigma}_i$, and any (relabeling) function $\gamma_{ij} : A_i \rightarrow A_j$, there exists a pure strategy $s_j \in \hat{\Sigma}_j$ such that for the infinite history h_∞ generated by s_i and s_j , there exists a dense sequence \mathbb{D} of periods wherein for any $T \in \mathbb{D}$, $s_j(h_T) = \gamma_{ij}(s_i(h_T))$.²⁷ Furthermore, $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ satisfies P (pure strategies) if there exists $\xi > 0$ such that for any $i = 1, 2$ and any $\sigma_i \in \hat{\Sigma}_i$, there exists a pure strategy $s_i \in \hat{\Sigma}_i$ wherein for any $h \in H$, if $s_i(h) = a_i$, then $\sigma_i(h)[a_i] > \xi$.

Finally, I define the condition that requires every player to learn to predict her opponent's true strategy (i.e., her opponent's *uniform approximate best response* to the opponent's prior belief) with *all* of her own diverse strategies, including her true one.

CONSISTENCY. Given $\varepsilon \geq 0$, $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ is ε -consistent with respect to $(\tilde{\rho}^1, \tilde{\rho}^2)$ if for any $i = 1, 2$, player i has a uniform ε -best response σ_i^ρ to $\tilde{\rho}^i$ in $\hat{\Sigma}_i$; that is, $(\sigma_i^\rho, \sigma_j^\rho) \in \hat{\Sigma}_1 \times \hat{\Sigma}_2$.²⁸

The following version of Nachbar's impossibility theorem states that if a learnable pair of players' strategy sets satisfies CSP, then the pair *cannot* satisfy the consistency condition. Let δ denote the (common) discount factor between the players in the repeated game.

IMPOSSIBILITY THEOREM (Nachbar 2005²⁹). (i) *Suppose that neither player has a weakly dominant action in the stage game (NWD). Then there exists $\bar{\delta} > 0$ such that for any $0 \leq \delta < \bar{\delta}$, there exists $\varepsilon > 0$ wherein for any $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ and any $(\tilde{\rho}^1, \tilde{\rho}^2)$,*

²⁶Nachbar (2005) uses a slightly weaker learnability condition: weak learnability. The weak learnability condition is equivalent to almost weak merging (instead of weak merging) plus (b). See Section 5 for details on almost weak merging.

²⁷A set \mathbb{D} of nonnegative integers is a dense sequence of periods if $\lim_{T \rightarrow \infty} \#(\mathbb{D} \cap \{0, 1, \dots, T-1\})/T = 1$.

²⁸Player i 's strategy σ_i is a uniform ε -best response to her prior belief $\tilde{\rho}^i$ if (a) σ_i is an ε -best response to $\tilde{\rho}^i$ and (b) for any h such that $\mu(\sigma_i, \tilde{\rho}^i)(h) > 0$, the continuation strategy $\sigma_{i,h}$ (following h) is an ε -best response to the posterior $\tilde{\rho}_h^i$ ($:= (\tilde{\rho}_{j,h}^i)_{j \neq i}$) (after h) in the continuation game following h . See Nachbar (1997, 2005) for details.

²⁹The original version in Nachbar (2005) is slightly more general because Nachbar (2005) only imposes the weak learnability condition on pure strategies in $(\hat{\Sigma}_1, \hat{\Sigma}_2)$; that is, "pure weak learnability," as mentioned in footnote 26. Pure weak learnability means that the weak learnability condition is satisfied by the pair (\hat{S}_1, \hat{S}_2) of pure strategy sets, where \hat{S}_i is the set of all (player i 's) pure strategies in $\hat{\Sigma}_i$ for $i = 1, 2$.

if $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ is learnable by $(\tilde{\rho}^1, \tilde{\rho}^2)$ and satisfies CSP, then $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ is not ε -consistent with respect to $(\tilde{\rho}^1, \tilde{\rho}^2)$.

- (ii) Suppose that each player's pure action maxmin payoff is strictly less than her min-max payoff in the stage game (MM). Then, for any $0 \leq \delta < 1$, there exists $\varepsilon > 0$ such that for any $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ and any $(\tilde{\rho}^1, \tilde{\rho}^2)$, if $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ is learnable by $(\tilde{\rho}^1, \tilde{\rho}^2)$ and satisfies CSP, then $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ is not ε -consistent with respect to $(\tilde{\rho}^1, \tilde{\rho}^2)$.³⁰

Note that from the viewpoint of convergence (to approximate Nash equilibrium), it is pointless to apply Nachbar's impossibility theorem to $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ that is not learnable (by any pair of players' prior beliefs). In other words, Nachbar's impossibility theorem must be applied to $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ that is learnable (and satisfies CSP). The following lemma shows that for any $\{\mathcal{P}_n\}_n$, the pair $(G_1(\{\mathcal{P}_n\}_n), G_2(\{\mathcal{P}_n\}_n))$ is learnable (by some pair of players' prior beliefs) and satisfies CSP.

LEMMA 1. For any countable family $\{\mathcal{P}_n\}_n$ of CRs, $(G_1(\{\mathcal{P}_n\}_n), G_2(\{\mathcal{P}_n\}_n))$ is learnable (by some pair of players' prior beliefs $(\tilde{\rho}^1, \tilde{\rho}^2)$) and satisfies CSP.

PROOF. Fix any $\{\mathcal{P}_n\}_n$. From Corollary 4 in Section 4.3.1, for any $i \neq j$, any $\sigma_i \in \Sigma_i$, and any $\sigma_j \in G_j(\{\mathcal{P}_n\}_n)$, $\tilde{\rho}_F^i$ leads player i to learn to predict σ_j with σ_i . Furthermore, we can easily modify $\tilde{\rho}_F^i$ to satisfy (b) in the learnability condition: for any $j \neq i$, $\tilde{\rho}_{F,j}^i(h)[a_j] > 0$ for any $a_j \in A_j$ and any $h \in H$. Therefore, $(G_1(\{\mathcal{P}_n\}_n), G_2(\{\mathcal{P}_n\}_n))$ is learnable by $(\tilde{\rho}_F^1, \tilde{\rho}_F^2)$.

Next, take any $i (\neq j)$, any pure strategy $s_i \in G_i(\{\mathcal{P}_n\}_n)$, and any (relabeling) function γ_{ij} . From the definition of $G_i(\{\mathcal{P}_n\}_n)$, there exists \bar{n} such that s_i is generated by $\mathcal{P}_{\bar{n}}$ (or, equivalently, $\mathcal{P}_{\bar{n}}$ is a CR of s_i). That is, for each $\alpha \in \mathcal{P}_{\bar{n}}$, there exists a pure action $a_i \in A_i$ such that $s_i(h) = a_i$ for all $h \in \alpha$. Then define player j 's pure strategy s_j as follows: $s_j(h) := \gamma_{ij}(s_i(h))$ for all $h \in H$. Evidently, $(s_j$ is well defined and) s_j is generated by $\mathcal{P}_{\bar{n}}$. Hence, $s_j \in G_j(\{\mathcal{P}_n\}_n)$. In addition, by the definition of s_j , it is obvious that for the infinite history h_∞ generated by s_i and s_j , $s_j(h_T) = \gamma_{ij}(s_i(h_T))$ for all T . Therefore, $(G_1(\{\mathcal{P}_n\}_n), G_2(\{\mathcal{P}_n\}_n))$ satisfies CS.

Finally, let $\xi := 1/(2(\#\bar{A} + 1))$, where $\#\bar{A} := \max_i \#A_i$. Take any i and any strategy $\sigma_i \in G_i(\{\mathcal{P}_n\}_n)$. Then, by the definition of $G_i(\{\mathcal{P}_n\}_n)$, there exists \hat{n} such that for any $\alpha \in \mathcal{P}_{\hat{n}}$ and any $h, h' \in \alpha$, $\|\sigma_i(h) - \sigma_i(h')\| \leq 1/(2\#\bar{A})$. For each $\alpha \in \mathcal{P}_{\hat{n}}$, take a finite history $\hat{h} \in \alpha$ and a pure action $\hat{a}_i \in A_i$ such that $\sigma_i(\hat{h})[\hat{a}_i] \geq 1/\#A_i$. Then define a pure strategy s_i as follows: for any $\alpha \in \mathcal{P}_{\hat{n}}$ and any $h \in \alpha$, let $s_i(h) := \hat{a}_i$, where $\hat{h} \in \alpha$ and $\sigma_i(\hat{h})[\hat{a}_i] \geq 1/\#A_i$. Evidently, $(s_i$ is well defined and) s_i is generated by $\mathcal{P}_{\hat{n}}$. Therefore, $s_i \in G_i(\{\mathcal{P}_n\}_n)$. Furthermore, take any $h \in H$. Then there exists a unique class $\alpha \in \mathcal{P}_{\hat{n}}$ such that $h \in \alpha$. By the definition of s_i , $s_i(h) = \hat{a}_i$, where $\hat{h} \in \alpha$ and $\sigma_i(\hat{h})[\hat{a}_i] \geq 1/\#A_i$. Since $\|\sigma_i(h) - \sigma_i(\hat{h})\| \leq 1/(2\#\bar{A})$, this implies that $\sigma_i(h)[\hat{a}_i] \geq \sigma_i(\hat{h})[\hat{a}_i] - 1/(2\#\bar{A}) \geq 1/\#A_i - 1/(2\#\bar{A}) \geq 1/\#\bar{A} - 1/(2\#\bar{A}) = 1/(2\#\bar{A}) > 1/(2(\#\bar{A} + 1)) = \xi$. Therefore, $(G_1(\{\mathcal{P}_n\}_n), G_2(\{\mathcal{P}_n\}_n))$ satisfies P. \square

³⁰Suppose that $(\hat{\Sigma}_1, \hat{\Sigma}_2)$ is ε -consistent with respect to $(\tilde{\rho}^1, \tilde{\rho}^2)$ (and is learnable by $(\tilde{\rho}^1, \tilde{\rho}^2)$). Then each player i 's prior belief $\tilde{\rho}^i$ leads her to learn to predict her opponent's true strategy (i.e., her opponent's uniform ε -best response to $\tilde{\rho}^j$) σ_j^ρ with her own true strategy σ_i^ρ . This implies that $(\sigma_1^\rho, \sigma_2^\rho)$ (almost surely) converges to approximate Nash equilibrium. Therefore, consistency is a stronger condition than convergence (to approximate Nash equilibrium).

From Nachbar’s impossibility theorem and Lemma 1, it immediately follows that Nachbar’s impossibility always holds for $(G_1(\{\mathcal{P}_n\}_n), G_2(\{\mathcal{P}_n\}_n))$; from this and the previous subsection, Nachbar’s impossibility holds for the pair of (players’) i.i.d. strategy sets, that of Markov strategy sets, that of finite automaton strategy sets, and that of computably regular strategy sets.

5. ALMOST WEAK MERGING

We also obtain quite similar characterization results with respect to almost weak merging. The learning criterion is introduced as a weaker concept of merging in Lehrer and Smorodinsky (1996b). Almost weak merging requires that the updated forecast about any finite-period future event be accurate *in almost all periods*. A set \mathbb{D} of nonnegative integers is a *dense* sequence of periods if $\lim_{T \rightarrow \infty} \#(\mathbb{D} \cap \{0, 1, \dots, T - 1\})/T = 1$.

DEFINITION 12. A prior $\tilde{\mu}$ almost weakly merges with a probability measure μ if for all $\varepsilon > 0$, all $k \geq 1$, and μ -almost all h_∞ , there exists a dense sequence \mathbb{D} of periods such that for all $T \in \mathbb{D}$,

$$\sup_{\mathbf{A} \in \mathcal{F}_{T+k}} |\tilde{\mu}(\mathbf{A}|\mathcal{F}_T) - \mu(\mathbf{A}|\mathcal{F}_T)| \leq \varepsilon.$$

Notice that a prior $\tilde{\mu}$ almost weakly merges with a probability measure μ if and only if for all $\varepsilon > 0$ and μ -almost all h_∞ , there exists a dense sequence \mathbb{D} of periods such that for all $T \in \mathbb{D}$, $\|f_{\tilde{\mu}}(h_T) - f_\mu(h_T)\| \leq \varepsilon$. This is also equivalent to the following: for all $\varepsilon > 0$, there exist a μ -probability 1 set \mathbf{Z} and an ordered family $\{T_m\}_{m=1}^{m=\infty}$ of time functions³¹; that is, $T_m : \mathbf{Z} \rightarrow \mathbb{N}$, such that (i) for all $h_\infty \in \mathbf{Z}$ and all m , $\|f_{\tilde{\mu}}(h_{T_m}) - f_\mu(h_{T_m})\| \leq \varepsilon$, and (ii) $\lim_{T \rightarrow \infty} N_T(h_\infty)/T = 1$ for all $h_\infty \in \mathbf{Z}$, where $N_T(h_\infty) := \#\{m \mid T_m(h_\infty) + 1 \leq T\}$.

As in the weak merging case, I define the phrase “almost weak merging with a set of probability measures.” A prior $\tilde{\mu}$ almost weakly merges with a set M of probability measures if $\tilde{\mu}$ almost weakly merges with all probability measures in M . Moreover, a set M of probability measures is almost weakly merged if there exists a prior $\tilde{\mu}$ such that $\tilde{\mu}$ almost weakly merges with M .

Now, I define the *almost generation* of CPSs. The definition simply states that for any $\varepsilon > 0$, the regularity of f is (almost surely) ε -approximated by one in a given family of CRs in almost all periods.

DEFINITION 13. A CPS $f : H \rightarrow \Delta(S)$ is almost generated by a family $\{\mathcal{P}_i\}_i$ of CRs if for all $\varepsilon > 0$, there exist an index i_0 , a μ_f -probability 1 set \mathbf{Z}_0 , and an ordered family $\{T_m\}_m$ of time functions such that (a) for all $\alpha \in \mathcal{P}_{i_0}$ and all $h_T, h'_{T'} \in \alpha$, if there exist $h_\infty, h'_\infty \in \mathbf{Z}_0$ and m, m' such that $h_T < h_\infty$, $T = T_m(h_\infty)$, $h'_{T'} < h'_\infty$, and $T' = T_{m'}(h'_\infty)$, then $\|f(h_T) - f(h'_{T'})\| \leq \varepsilon$; and (b) $\lim_{T \rightarrow \infty} N_T(h_\infty)/T = 1$ for all $h_\infty \in \mathbf{Z}_0$, where $N_T(h_\infty) := \#\{m \mid T_m(h_\infty) + 1 \leq T\}$.

³¹By an ordered family $\{T_m\}_m$ of time functions, I mean that $T_m(h_\infty) < T_{m+1}(h_\infty)$ for all m and all $h_\infty \in \mathbf{Z}$.

Further, a probability measure μ is almost generated by a family $\{\mathcal{P}_i\}_i$ of CRs if there exists a CPS f_μ corresponding to μ such that f_μ is almost generated by $\{\mathcal{P}_i\}_i$. The set of all probability measures almost generated by $\{\mathcal{P}_i\}_i$ is denoted by $\text{AG}(\{\mathcal{P}_i\}_i)$. Obviously, $\text{EG}(\{\mathcal{P}_i\}_i) \subsetneq \text{AG}(\{\mathcal{P}_i\}_i)$.

As in the weak merging case, no prior almost weakly merges with more probability measures than those almost generated by a countable family of CRs.

PROPOSITION 5. *For any prior $\tilde{\mu}$, there exists a countable family $\{\mathcal{P}_i\}_i$ of CRs such that $\tilde{\mu}$ does not almost weakly merge with any $\mu \notin \text{AG}(\{\mathcal{P}_i\}_i)$.*

See [Appendix C](#) for the proof.

For any $\{\mathcal{P}_i\}_i$, I construct a prior $\tilde{\mu}_M$ that almost weakly merges with all probability measures almost generated by $\{\mathcal{P}_i\}_i$. For that purpose, I only have to modify the switching criterion (3) in the definition of $i(\cdot)$ as follows:

$$i(h_T) := \begin{cases} m(h_T) + 1 & \text{if } n(h_T) \geq n_0^{\beta(h_T)} \text{ and } \frac{\sum_{i=1}^{i=m(h_T)+1} \sum_{\alpha \in \mathcal{P}_i} n_0^\alpha}{n(h_T)} < \frac{1}{m(h_T)} \\ m(h_T) & \text{otherwise.} \end{cases} \quad (4)$$

All other things are exactly the same as in the weak merging case. Let f_M denote the modified frequency-based CPS and let $\tilde{\mu}_M := \tilde{\mu}_{f_M}$. In the definition of $i(\cdot)$, a new inequality is added to the switching criterion so that for almost all categories, the prior sample size is negligible relative to the number of effective periods; this fact will be used for proving [Proposition 6](#).

PROPOSITION 6. *For any countable family $\{\mathcal{P}_i\}_i$ of CRs, there exists a frequency-based prior $\tilde{\mu}_M$ such that $\tilde{\mu}_M$ almost weakly merges with all $\mu \in \text{AG}(\{\mathcal{P}_i\}_i)$.*

See [Appendix C](#) for the proof.

Note that $\tilde{\mu}_M$ also weakly merges with $\text{EG}(\{\mathcal{P}_i\}_i)$; the proof is quite similar to that of [Proposition 2](#) in [Section 3.3](#). Therefore, $\tilde{\mu}_M$ not only almost weakly merges with $\text{AG}(\{\mathcal{P}_i\}_i)$ but also weakly merges with $\text{EG}(\{\mathcal{P}_i\}_i)$.

COROLLARY 6. *For any countable family $\{\mathcal{P}_i\}_i$ of CRs, there exists a prior $\tilde{\mu}_M$ such that $\tilde{\mu}_M$ not only almost weakly merges with $\text{AG}(\{\mathcal{P}_i\}_i)$ but also weakly merges with $\text{EG}(\{\mathcal{P}_i\}_i)$.*

Finally, [Propositions 5](#) and [6](#) give us the characterization of almost weak merging with a set of probability measures, as in the weak merging case.

THEOREM 3. *A set M of probability measures is almost weakly merged if and only if there exists a countable family $\{\mathcal{P}_i\}_i$ of CRs such that $M \subseteq \text{AG}(\{\mathcal{P}_i\}_i)$.*

The proof is similar to that of [Theorem 1](#).

REMARK 4. As noted in [Section 4.2](#), let $\tilde{\rho}_M^i$ be the modified frequency-based prior belief (in the repeated game) that is the same as f_M except that $D_T^{(i,\alpha)} (= (D_{T,j}^{(i,\alpha)})_{j \neq i})$ are

the conditional empirical distributions of the opponents' realized actions. Then, by a similar argument to [Corollary 6](#), both [Corollary 3](#) and the almost weak merging version of [Corollary 3](#) hold for $\tilde{\rho}_M^i$: $\tilde{\rho}_M^i$ has the dominance property of both weak merging and almost weak merging.

REMARK 5. As a belief for CSFP, $\tilde{\rho}_M^i$ has a sophisticated no-regret property (i.e., universal (classwise) conditional consistency) as well. See [Noguchi \(2000, 2003, 2009\)](#) for details on the (no-regret) properties of CSFP. Further, this fact suggests that there may be some important relations between (almost) weak merging and no-regret, but this is an open issue. See [Cesa-Bianchi and Lugosi \(2006\)](#) for a comprehensive survey on associated issues and learning procedures in the no-regret literature.

6. CONCLUDING REMARKS

I have provided a characterization of a learnable set with respect to (almost) weak merging by using the conditioning rule and eventual generation. That is, I have shown that a set of probability measures is learnable if and only if the set is contained in the set of probability measures eventually generated by some countable family of conditioning rules. Furthermore, I have demonstrated that this characterization result can be extended to the case of infinitely repeated games by introducing the concept of a learnable set correspondence. I conclude by stating several issues pertaining to my characterization.

Application to convergence to approximate Nash equilibrium

My characterization may have interesting applications. In particular, it may enable us to find out various types of “smart” prior beliefs.³² Indeed, making use of the characterization result in this paper, [Noguchi \(2015\)](#) constructs smart prior beliefs that lead players to learn to play approximate Nash equilibrium in any repeated game with perfect monitoring, combined with *smooth approximate* optimal behavior. This positive result has a significant implication for the impossibility result in [Nachbar \(1997, 2005\)](#); see [Section 4.3.2](#) for the impossibility result and its examples. Although Nachbar's impossibility result does not necessarily imply the impossibility of convergence to (approximate) Nash equilibrium, it, along with other impossibility results, has led to skepticism regarding whether a general result of convergence could be obtained for Bayesian learning in repeated games. However, the positive result in [Noguchi \(2015\)](#) induces the following *possibility* theorem: there exist prior beliefs $\tilde{\rho}_* := (\tilde{\rho}_*^i)_i$ such that although each player i 's prior belief $\tilde{\rho}_*^i$ leads her to approximately learn to predict her opponents' diverse strategies and the learning performances of $\tilde{\rho}_*$ are symmetric to each other, the prior beliefs $\tilde{\rho}_*$ (almost surely) lead the players to learn to play approximate Nash equilibrium for any stage game payoff and any discount factor (combined with smooth approximate optimal behavior). In other words, the possibility result clarifies that Nachbar's impossibility is *different* from the impossibility of learning to play approximate Nash equilibrium in a general sense.

³²By a smart prior belief, I mean a prior belief that leads the player to learn to predict as many strategies of her opponents as possible.

Size of the learnable set: Large or small

I have characterized a learnable set. Then a natural (open) question arises as to whether a learnable set is generally large or small.³³ There are many different criteria to measure the size of a set.³⁴ Here, I briefly discuss several criteria; in the following argument, a probability measure is identified with its corresponding CPSs. First, in the supremum norm topology, a canonical learnable set $EG(\{\mathcal{P}_i\}_i)$ is not open. However, given $\varepsilon > 0$, we may consider the set $EG_\varepsilon(\{\mathcal{P}_i\}_i)$ of probability measures ε -eventually generated by $\{\mathcal{P}_i\}_i$,³⁵ and we can show that for any $\varepsilon > 0$, frequency-based prior $\tilde{\mu}_F$ ε -weakly merges with $EG_\varepsilon(\{\mathcal{P}_i\}_i)$ ³⁶; $EG_\varepsilon(\{\mathcal{P}_i\}_i)$ is considered an approximately learnable set. Evidently, the interior of $EG_\varepsilon(\{\mathcal{P}_i\}_i)$ is nonempty in this topology. From this point of view, $EG(\{\mathcal{P}_i\}_i)$ may not be small, at least in an approximate sense. However, even $G_\varepsilon(\mathcal{P}_{\text{Id}})$ (i.e., the set of ε -i.i.d. probability measures) has a nonempty interior (for any $\varepsilon > 0$) although $G_\varepsilon(\mathcal{P}_{\text{Id}})$ may be supposed to be small.

Next, in the weak (-star) topology, a *first category* set is often considered as a “small” one (e.g., Dekel and Feinberg 2006).³⁷ In this topology, for any small $\varepsilon > 0$ and any CR \mathcal{P} , $G_\varepsilon(\mathcal{P})$ is a closed set with its interior being empty. Since $G_\varepsilon(\{\mathcal{P}_i\}_i) = \bigcup_i G_\varepsilon(\mathcal{P}_i)$,³⁸ $G_\varepsilon(\{\mathcal{P}_i\}_i)$ is a countable union of closed sets with their interiors being empty, that is, a first category set. As for $EG(\{\mathcal{P}_i\}_i)$, however, we can easily show that both $EG(\{\mathcal{P}_i\}_i)$ and its complement are dense (for any $\{\mathcal{P}_i\}_i$). This reflects the following two facts. On the one hand, the definition of $EG(\{\mathcal{P}_i\}_i)$ only imposes a restriction on the *limit* (or long-run) property of a probability measure (i.e., a CPS): (for any $\varepsilon > 0$) the regularity of a probability measure (in $EG(\{\mathcal{P}_i\}_i)$) is (almost surely) eventually ε -approximated by one in $\{\mathcal{P}_i\}_i$. On the other hand, given any two probability measures, the weak (-star) topology is only effective to measure the differences in the probabilities of finite-period events (i.e., short- and medium-run events)³⁹; that is, it *cannot* capture the differences in the limit (or long-run) property. These facts lead us to doubt the validity of the weak (-star) topology as a criterion for measuring the size of a learnable set; that is, in general, neither $EG(\{\mathcal{P}_i\}_i)$ nor its complement may be of first category. (This is in contrast to the case of merging. See the next remark on merging.)

Finally, instead of the above topological criteria, we may consider a “measure theoretic” one, *shyness*, which is an infinite dimensional version of Lebesgue measure zero;

³³The author thanks Drew Fudenberg for suggesting this issue and introducing him to the relevant literature. Further, the author is grateful to an anonymous referee and the co-editor for introducing him to the literature on prevalence and shyness.

³⁴Miller and Sanchirico (1997) investigate an interesting related problem.

³⁵The set $EG_\varepsilon(\{\mathcal{P}_i\}_i)$ is defined as follows: $\mu \in EG_\varepsilon(\{\mathcal{P}_i\}_i)$ if and only if there exist an index i_0 , a μ -probability 1 set \mathbf{Z}_0 , and a time function $T_0: \mathbf{Z}_0 \rightarrow \mathbb{N}$ such that for all $\alpha \in \mathcal{P}_{i_0}$ and all $h_T, h_{T'} \in \alpha$, if there exist $h_\infty, h'_\infty \in \mathbf{Z}_0$ wherein $h_T < h_\infty$, $T \geq T_0(h_\infty)$, $h_{T'} < h'_\infty$, and $T' \geq T_0(h'_\infty)$, then $\|f_\mu(h_T) - f_\mu(h_{T'})\| \leq \varepsilon$.

³⁶We say that $\tilde{\mu}$ ε -weakly merges with μ if for μ -almost all h_∞ , $\limsup_{T \rightarrow \infty} \|f_{\tilde{\mu}}(h_T) - f_\mu(h_T)\| \leq \varepsilon$.

³⁷A set is of first category if the set is contained in a countable union of closed sets with their interiors being empty.

³⁸Given $\varepsilon > 0$, $G_\varepsilon(\{\mathcal{P}_i\}_i)$ is defined as follows: $\mu \in G_\varepsilon(\{\mathcal{P}_i\}_i)$ if and only if there exists i such that \mathcal{P}_i is an ε -ACR of a CPS f_μ corresponding to μ . Evidently, $G(\{\mathcal{P}_i\}_i) \subseteq G_\varepsilon(\{\mathcal{P}_i\}_i) \subseteq EG_\varepsilon(\{\mathcal{P}_i\}_i)$ for any $\varepsilon > 0$.

³⁹Indeed, a sequence $\{\mu_n\}_n$ weakly converges to μ if and only if for any finite history h , $\mu_n(h) \rightarrow \mu(h)$ as $n \rightarrow \infty$. See Parthasarathy (1967) for the properties of the weak (-star) topology.

see [Anderson and Zame \(2001\)](#). As an attempt to check whether this criterion is valid, I consider a tractable example, that is, the set $\text{SEG}(\{\mathcal{P}_i\}_i)$ of probability measures that are *surely* eventually approximated by $\{\mathcal{P}_i\}_i$: a probability 1 set \mathbf{Z}_0 (in the definition of $\text{EG}(\{\mathcal{P}_i\}_i)$) is replaced by the set \mathbf{H}_∞ of all infinite histories.⁴⁰ (Hence, for any $\{\mathcal{P}_i\}_i$, $\text{G}(\{\mathcal{P}_i\}_i) \subsetneq \text{SEG}(\{\mathcal{P}_i\}_i) \subsetneq \text{EG}(\{\mathcal{P}_i\}_i)$.) On the one hand, as in the case of $\text{EG}(\{\mathcal{P}_i\}_i)$, both $\text{SEG}(\{\mathcal{P}_i\}_i)$ and its complement are dense in the weak (-star) topology. On the other hand, we can show that $\text{SEG}(\{\mathcal{P}_i\}_i)$ is shy.⁴¹

From the above argument, shyness may be more appropriate than the other two criteria. Nonetheless, it does not seem easy to verify the shyness of $\text{EG}(\{\mathcal{P}_i\}_i)$, because a probability 1 set \mathbf{Z}_0 is quite different according to a probability measure.⁴² Therefore, it is unknown whether, in general, $\text{EG}(\{\mathcal{P}_i\}_i)$ is shy.

Merging

This paper only explores (almost) weak merging. The same issue can be explored about merging; that is, the identification of a set of probability measures with which a prior merges. It has been known that no prior can merge with even the set of i.i.d. probability measures. Further, [Sandroni and Smorodinsky \(1999\)](#) investigate the relationships between merging and weak merging, and show that merging requires weak merging with fast speed in almost all periods. These results lead us to conjecture that, in general, a merged set (of probability measures) may be much smaller than a weakly merged set. Accordingly, [Ryabko \(2010\)](#) provides a general result of the merged set: under a standard assumption (i.e., local absolute continuity),⁴³ any merged set M has a countable dense subset in the *total variation metric* topology.⁴⁴ As a corollary, any strictly convex combination of the probability measures in the countable subset merges with M . Interestingly, combining this corollary with a mathematical proposition in [Dekel and Feinberg \(2006\)](#), we immediately obtain that a merged set is of first category (in the weak (-star) topology). See [Appendix D](#) for the proof.

I address several questions from the viewpoint of game theory. First, since the total variation metric topology is quite strong, we may expect a sharper evaluation of the size of a merged set. Second, it is desirable to provide a (complete) characterization of a merged set without the local absolute continuity assumption. This is important for game theory because the conditional probability (or the posterior) on an event of prob-

⁴⁰That is, $\text{SEG}(\{\mathcal{P}_i\}_i)$ is defined as follows: $\mu \in \text{SEG}(\{\mathcal{P}_i\}_i)$ if and only if for any $\varepsilon > 0$, there exist an index i_0 and a time function $T_0 : \mathbf{H}_\infty \rightarrow \mathbb{N}$ such that for all $\alpha \in \mathcal{P}_{i_0}$ and all $h_T, h_{T'} \in \alpha$, if there exist h_∞, h'_∞ wherein $h_T < h_\infty, T \geq T_0(h_\infty), h_{T'} < h'_\infty$, and $T' \geq T_0(h'_\infty)$, then $\|f_\mu(h_T) - f_\mu(h_{T'})\| \leq \varepsilon$.

⁴¹To be precise, the set of CPSs corresponding to $\text{SEG}(\{\mathcal{P}_i\}_i)$ is *finitely shy* in $\prod_{h \in H} \Delta(S)$ equipped with the supremum norm topology. (Finite shyness implies shyness; see [Anderson and Zame 2001](#).)

⁴²Note that the regularity of any probability measure μ in $\text{SEG}(\{\mathcal{P}_i\}_i)$ is eventually approximated by $\{\mathcal{P}_i\}_i$ on the same set; that is, the set of all infinite histories. This property enables us to show the (finite) shyness of $\text{SEG}(\{\mathcal{P}_i\}_i)$.

⁴³The literature on merging usually assumes the local absolute continuity condition on a prior $\tilde{\mu}$ and a probability measure μ : for any finite history h , if $\tilde{\mu}(h) = 0$, then $\mu(h) = 0$. For example, see [Kalai and Lehrer \(1994\)](#), [Sandroni \(1998\)](#), and [Sandroni and Smorodinsky \(1999\)](#). Note that in order to provide a characterization of a weakly merged set, this paper does not assume any condition on the relationships between a prior and a probability measure.

⁴⁴The total variation metric d is defined as follows: for any μ, μ' , $d(\mu, \mu') := \sup_{\mathbf{A} \in \mathcal{F}} |\mu(\mathbf{A}) - \mu'(\mathbf{A})|$.

ability zero is often considered in many games, including repeated ones.⁴⁵ Finally, it is worthwhile to explore how the above results of merging can be applied to the study of Bayesian learning in repeated games, as in Section 4 of this paper.

APPENDIX A

I prepare a mathematical proposition to prove Propositions 2 and 6. Let $\mathcal{T}_n^\alpha(h_\infty)$ denote the calendar time of the n th α -active period in h_∞ ; $\mathcal{T}_n^\alpha(h_\infty) < \infty$ means that α is active at least n times in h_∞ . Let $\mathbf{d}_n^\alpha(h_\infty)$ designate the vector in which each coordinate $\mathbf{d}_n^\alpha(h_\infty)[s]$ is the number of times that s has occurred in the first n α -active periods along h_∞ . The next proposition extends a basic fact of large deviations to a conditional case. It states that if the probabilities of a state s have common upper and lower bounds in active periods of a given class, then the probability that the frequency of s in the first n active periods of that class is not between those bounds, decreases exponentially in the sample size n .

PROPOSITION A. *Let α be any subset of H . Take any probability measure μ and any state s that satisfy the following condition: for all $h \in \alpha$ with $\mu(h) > 0$, $l \leq \mu(s|h) \leq L$, where l and L are nonnegative constants. Then, for all $\varepsilon > 0$ and all $n = 1, 2, \dots$,*

$$\mu\left(\mathcal{T}_n^\alpha < \infty, \frac{\mathbf{d}_n^\alpha[s]}{n} \leq l - \varepsilon \text{ or } \frac{\mathbf{d}_n^\alpha[s]}{n} \geq L + \varepsilon\right) \leq 2 \exp(-2n\varepsilon^2),$$

where $\mathbf{d}_n^\alpha[s]$ is the s -coordinate of $\mathbf{d}_n^\alpha(h_\infty)$.

I first show the following claim in the case that $\mu(s|h) \geq l$. The case that $\mu(s|h) \leq L$ is proved similarly. Let P_l denote the probability measure in the coin-flipping process: in each independent trial, a coin is used that generates heads with probability l . Let $d_n[\mathbb{H}]$ denote the number of times that heads has come out in the first n trials. By $h * h'$, I mean the concatenation of h with h' , that is, h occurs first and then h' happens.

CLAIM A. *Let α be any subset of H . Take any probability measure μ and any state s that satisfy the following condition: for all $h \in \alpha$ with $\mu(h) > 0$, $\mu(s|h) \geq l$. Then, for all $n \geq 1$ and all $m = 0, 1, \dots, n$,*

$$\mu\left(\mathcal{T}_n^\alpha < \infty, \frac{\mathbf{d}_n^\alpha[s]}{n} \leq \frac{m}{n}\right) \leq P_l\left(\frac{d_n[\mathbb{H}]}{n} \leq \frac{m}{n}\right).$$

PROOF. I prove Claim A inductively. Define a class $\alpha_F := \{h \mid h \in \alpha, h' \notin \alpha \text{ for all } h' < h\}$: α_F consists of histories whose next period is the first α -active period. Clearly, $\alpha_F \subseteq \alpha$. Notice that $\sum_{h \in \alpha_F} \mu(h) \leq 1$. For notational simplicity, let $\lambda(m, n) := P_l(d_n[\mathbb{H}]/n \leq m/n)$.

Step 1. I first show Claim A for the case that $n = 1$. Let $m = 0$. Then

$$\mu(\mathcal{T}_1^\alpha < \infty, \mathbf{d}_1^\alpha[s] = 0) = \sum_{h \in \alpha_F} (1 - \mu(s|h))\mu(h) \leq \sum_{h \in \alpha_F} (1 - l)\mu(h) \leq 1 - l.$$

⁴⁵The local absolute continuity implies that a forecaster does not need to consider the posterior $\tilde{\mu}(\cdot|h)$ after any history h that she believes never happens (i.e., $\tilde{\mu}(h) = 0$), because h actually does not occur (i.e., $\mu(h) = 0$).

Since $\lambda(0, 1) = 1 - l$, **Claim A** is true for the case that $n = 1$ and $m = 0$. When $m = 1$, $\lambda(1, 1) = 1$ so that **Claim A** holds in that case. Therefore, **Claim A** is true for the case that $n = 1$.

Step 2. Suppose that **Claim A** is true for n . Given the inductive hypothesis, I show that **Claim A** is also true for $n + 1$. In the case that $m = n + 1$, $\lambda(n + 1, n + 1) = 1$. Thus, **Claim A** is trivial in that case. Consider the case that $m < n + 1$. First, note that the following equality holds:

$$\mu\left(\mathcal{T}_{n+1}^\alpha < \infty, \frac{\mathbf{d}_{n+1}^\alpha[s]}{n+1} \leq \frac{m}{n+1}\right) = \sum_{h \in \alpha_F} \sum_{s' \in S} \mu\left(\mathcal{T}_{n+1}^\alpha < \infty, \frac{\mathbf{d}_{n+1}^\alpha[s]}{n+1} \leq \frac{m}{n+1} \mid h * s'\right) \mu(h * s').$$

Step 3. Let $\alpha_h := \{h' \mid h * h' \in \alpha\}$ and define a probability measure μ_h by $\mu_h(h') := \mu(h * h') / \mu(h)$ (when $\mu(h) > 0$). Note that for all $h \in H$ (with $\mu(h) > 0$), (α_h, μ_h, s, l) satisfies the condition in **Claim A**. Furthermore, when $h \in \alpha_F$ (with $\mu(h * s) > 0$),

$$\mu\left(\mathcal{T}_{n+1}^\alpha < \infty, \frac{\mathbf{d}_{n+1}^\alpha[s]}{n+1} \leq \frac{m}{n+1} \mid h * s\right) = \mu_{h*s}\left(\mathcal{T}_n^{\alpha_{h*s}} < \infty, \frac{\mathbf{d}_n^{\alpha_{h*s}}[s]}{n} \leq \frac{m-1}{n}\right).$$

(Note that when $m = 0$, the right hand side is zero.) When $h \in \alpha_F$ and $s' \neq s$ (with $\mu(h * s') > 0$),

$$\mu\left(\mathcal{T}_{n+1}^\alpha < \infty, \frac{\mathbf{d}_{n+1}^\alpha[s]}{n+1} \leq \frac{m}{n+1} \mid h * s'\right) = \mu_{h*s'}\left(\mathcal{T}_n^{\alpha_{h*s'}} < \infty, \frac{\mathbf{d}_n^{\alpha_{h*s'}}[s]}{n} \leq \frac{m}{n}\right).$$

Since $(\alpha_{h*s}, \mu_{h*s}, s, l)$ and $(\alpha_{h*s'}, \mu_{h*s'}, s, l)$ satisfy the condition in **Claim A**, the inductive hypothesis implies that $\mu_{h*s}(\mathcal{T}_n^{\alpha_{h*s}} < \infty, \mathbf{d}_n^{\alpha_{h*s}}[s]/n \leq (m-1)/n) \leq \lambda(m-1, n)$ and $\mu_{h*s'}(\mathcal{T}_n^{\alpha_{h*s'}} < \infty, \mathbf{d}_n^{\alpha_{h*s'}}[s]/n \leq m/n) \leq \lambda(m, n)$ for all $s' \neq s$. From this and Step 2, it is derived that

$$\begin{aligned} \mu\left(\mathcal{T}_{n+1}^\alpha < \infty, \frac{\mathbf{d}_{n+1}^\alpha[s]}{n+1} \leq \frac{m}{n+1}\right) &\leq \sum_{h \in \alpha_F} \left[\lambda(m-1, n) \mu(h * s) + \sum_{s' \neq s} \lambda(m, n) \mu(h * s') \right] \\ &= \sum_{h \in \alpha_F} \mu(h) [\lambda(m-1, n) \mu(s|h) + \lambda(m, n) (1 - \mu(s|h))]. \end{aligned}$$

Step 4. Note that $\lambda(m-1, n) \leq \lambda(m, n)$. Furthermore, $1 - \mu(s|h) \leq 1 - l$ for all $h \in \alpha$ (with $\mu(h) > 0$). Thus, for all $h \in \alpha_F$ (with $\mu(h) > 0$),

$$\mu(s|h) \lambda(m-1, n) + (1 - \mu(s|h)) \lambda(m, n) \leq l \lambda(m-1, n) + (1 - l) \lambda(m, n).$$

From this and the inequality in Step 3, it follows that

$$\begin{aligned} \mu\left(\mathcal{T}_{n+1}^\alpha < \infty, \frac{\mathbf{d}_{n+1}^\alpha[s]}{n+1} \leq \frac{m}{n+1}\right) &\leq \sum_{h \in \alpha_F} \mu(h) [l \lambda(m-1, n) + (1 - l) \lambda(m, n)] \\ &\leq l \lambda(m-1, n) + (1 - l) \lambda(m, n). \end{aligned}$$

Since $l \lambda(m-1, n) + (1 - l) \lambda(m, n) = \lambda(m, n + 1)$, the proof is completed. \square

PROOF OF PROPOSITION A. We have the following inequalities of large deviations (see, e.g., Shiryaev 1996): for all $\varepsilon > 0$ and all n , $P_l(d_n[\mathbb{H}]/n \leq l - \varepsilon) \leq \exp(-2n\varepsilon^2)$ and $P_L(d_n[\mathbb{H}]/n \geq L + \varepsilon) \leq \exp(-2n\varepsilon^2)$. Let $m_{l,\varepsilon}^n := \max\{m \mid m \leq (l - \varepsilon)n\}$ and $M_{L,\varepsilon}^n := \min\{m \mid m \geq (L + \varepsilon)n\}$. By Claim A and the inequalities above, for all $\varepsilon > 0$ and all n ,

$$\begin{aligned} \mu\left(\mathcal{T}_n^\alpha < \infty, \frac{\mathbf{d}_n^\alpha[s]}{n} \leq l - \varepsilon \text{ or } \frac{\mathbf{d}_n^\alpha[s]}{n} \geq L + \varepsilon\right) & \\ & \leq \mu\left(\mathcal{T}_n^\alpha < \infty, \frac{\mathbf{d}_n^\alpha[s]}{n} \leq \frac{m_{l,\varepsilon}^n}{n}\right) + \mu\left(\mathcal{T}_n^\alpha < \infty, \frac{\mathbf{d}_n^\alpha[s]}{n} \geq \frac{M_{L,\varepsilon}^n}{n}\right) \\ & \leq P_l\left(\frac{d_n[\mathbb{H}]}{n} \leq \frac{m_{l,\varepsilon}^n}{n}\right) + P_L\left(\frac{d_n[\mathbb{H}]}{n} \geq \frac{M_{L,\varepsilon}^n}{n}\right) \\ & = P_l\left(\frac{d_n[\mathbb{H}]}{n} \leq l - \varepsilon\right) + P_L\left(\frac{d_n[\mathbb{H}]}{n} \geq L + \varepsilon\right) \\ & \leq 2 \exp(-2n\varepsilon^2). \quad \square \end{aligned}$$

APPENDIX B

Without loss of generality, we may assume (throughout Appendix B) that $\{\mathcal{P}_i\}_i$ is ordered in fineness: $\mathcal{P}_i \leq \mathcal{P}_{i+1}$ for all i . The following lemma will be used to prove Proposition 2. It states that a forecaster employs finer and finer classes as (temporary) categories, as time proceeds.

LEMMA B. For all h_∞ , $\lim_{T \rightarrow \infty} i(h_T) = \infty$.

PROOF. Suppose that $\liminf_{T \rightarrow \infty} i(h_T) < \infty$ for some h_∞ . Thus, $i(h_{T_k}) = i_0$ for infinitely many T_k . Since \mathcal{P}_{i_0} and \mathcal{P}_{i_0+1} only have finite classes and $\mathcal{P}_{i_0} \leq \mathcal{P}_{i_0+1}$, there exist $\alpha_0 \in \mathcal{P}_{i_0}$ and $\beta_0 \in \mathcal{P}_{i_0+1}$ such that $\beta_0 \subseteq \alpha_0$, and $\alpha(h_{T_{k_l}}) = \alpha_0$ and $\beta(h_{T_{k_l}}) = \beta_0$ for some infinite subsequence $\{T_{k_l}\}_l$ of $\{T_k\}_k$; clearly, $i(h_{T_{k_l}}) = i_0$ for all T_{k_l} . This, along with the definition of $m(\cdot)$, implies that $m(h_{T_{k_l}}) = i_0$ for all T_{k_l} . However, then, by the definition of $n(\cdot)$, $n(h_{T_{k_l}}) \rightarrow \infty$ as $T_{k_l} \rightarrow \infty$. This means that for some T_{k_l} , $n(h_{T_{k_l}}) \geq n_0^{\beta_0} = n_0^{\beta(h_{T_{k_l}})}$, and, consequently, by (3) in Section 3.3, $i(h_{T_{k_l}}) = m(h_{T_{k_l}}) + 1 = i_0 + 1$. This is a contradiction. \square

PROOF OF PROPOSITION 2. Let f_F be the frequency-based CPS for $\{\mathcal{P}_i\}_i$. Fix any $\mu \in \text{EG}(\{\mathcal{P}_i\}_i)$ and let f_μ be a CPS corresponding to μ . Then it suffices to show that for μ -almost all h_∞ , $\|f_F(h_T) - f_\mu(h_T)\| \rightarrow 0$ as $T \rightarrow \infty$:

$$\mu\left(\bigcap_{m \geq 1} \bigcup_{T' \geq 0} \bigcap_{T \geq T'} \left\{h_\infty \mid \|f_F(h_T) - f_\mu(h_T)\| < \frac{1}{m}\right\}\right) = 1.$$

Equivalently, I only have to show that for all $m = 1, 2, \dots$,

$$\mu\left(\bigcap_{T' \geq 0} \bigcup_{T \geq T'} \left\{h_\infty \mid \|f_F(h_T) - f_\mu(h_T)\| \geq \frac{1}{m}\right\}\right) = 0.$$

Since $\mu \in \text{EG}(\{\mathcal{P}_i\}_i)$, we obtain a μ -probability 1 set \mathbf{Z}_0 (i.e., $\mu(\mathbf{Z}_0) = 1$) such that for any $\varepsilon > 0$, there exist i_0 and $T_0 : \mathbf{Z}_0 \rightarrow \mathbb{N}$ wherein for all $\alpha \in \mathcal{P}_{i_0}$ and all $h_T, h'_{T'} \in \alpha$, if there exist $h_\infty, h'_\infty \in \mathbf{Z}_0$ such that $h_T < h_\infty$, $T \geq T_0(h_\infty)$, $h'_{T'} < h'_\infty$, and $T' \geq T_0(h'_\infty)$, then $\|f_\mu(h_T) - f_\mu(h'_{T'})\| \leq \varepsilon$.

Step 1. Let $\varepsilon := 1/(3m)$. Since $\{\mathcal{P}_i\}_i$ is ordered in fineness, we may take $i_0 \geq 3m$. For each $\alpha \in \mathcal{P}_{i_0}$, define a class $\hat{\alpha}$ as follows: $h_T \in \hat{\alpha}$ if and only if (a) $h_T \in \alpha$ and (b) $h_T < h_\infty$ and $T \geq T_0(h_\infty)$ for some $h_\infty \in \mathbf{Z}_0$. Then, for all $\alpha \in \mathcal{P}_{i_0}$ and all s , let $L^\alpha[s] := \sup_{h \in \hat{\alpha}} f_\mu(h)[s]$ and $l^\alpha[s] := \inf_{h \in \hat{\alpha}} f_\mu(h)[s]$; note that $L^\alpha[s] - l^\alpha[s] \leq \varepsilon$ for all s . Furthermore, for all (i, α) , define a class $\gamma(i, \alpha)$ as follows: $h_T \in \gamma(i, \alpha)$ if and only if (a) time $T + 1$ is either an effective period of (i, α) or one of the first $n_0^{(i, \alpha)}$ effective periods of (i_p, α_p) , and (b) $h_T < h_\infty$ and $T \geq T_0(h_\infty)$ for some $h_\infty \in \mathbf{Z}_0$. Since $\{\mathcal{P}_i\}_i$ is ordered in fineness, for each (i, α) with $i \geq i_0 + 1$, there exists a unique class $\beta \in \mathcal{P}_{i_0}$ such that $\gamma(i, \alpha) \subseteq \hat{\beta}$; then let $L^{(i, \alpha)}[s] := L^\beta[s]$ and $l^{(i, \alpha)}[s] := l^\beta[s]$ for all s . Hence, for all $h \in \gamma(i, \alpha)$ and all s , $l^{(i, \alpha)}[s] \leq f_\mu(h)[s] (= \mu(s|h)) \leq L^{(i, \alpha)}[s]$. Moreover, it follows from the definition of f_F (in Section 3.3) and Lemma B that for all $h_\infty \in \mathbf{Z}_0$ and all $j \geq i_0 + 1$, there exists \hat{T} such that for all $T \geq \hat{T}$, $f_F(h_T) = D_T^{(i, \alpha)} = \mathbf{d}_n^{\gamma(i, \alpha)}(h_\infty)/n$ for some (i, α) with $i \geq j$ and some $n \geq n_0^{(i, \alpha)}$ ($= n_0^\alpha$).

Step 2. For all categories (i, α) with $i \geq i_0 + 1$, let

$$\mathbf{B}_n^{(i, \alpha)} := \left\{ h_\infty \mid \mathcal{T}_n^{\gamma(i, \alpha)} < \infty, \exists s \left(\frac{\mathbf{d}_n^{\gamma(i, \alpha)}[s]}{n} \leq l^{(i, \alpha)}[s] - \frac{1}{i} \text{ or } \frac{\mathbf{d}_n^{\gamma(i, \alpha)}[s]}{n} \geq L^{(i, \alpha)}[s] + \frac{1}{i} \right) \right\}.$$

Then, from Step 1, it follows that for all $j \geq i_0 + 1$,

$$\bigcap_{T' \geq 0} \bigcup_{T \geq T'} \left\{ h_\infty \mid \|f_F(h_T) - f_\mu(h_T)\| \geq \frac{1}{m} \right\} \cap \mathbf{Z}_0 \subseteq \bigcup_{i \geq j} \bigcup_{\alpha \in \mathcal{P}_i} \bigcup_{n \geq n_0^{(i, \alpha)}} \mathbf{B}_n^{(i, \alpha)}.$$

Step 3. From Step 1 and Proposition A (in Appendix A), it follows that for all (i, α) with $i \geq i_0 + 1$, $\mu(\mathbf{B}_n^{(i, \alpha)}) \leq 2\#S \exp(-2ni^{-2})$ for all n . Furthermore, by the definition of $n_0^{(i, \alpha)}$, $n_0^{(i, \alpha)} = n_0^\alpha$ for all $\alpha \in \mathcal{P}_i$ and all i . From this and (2) in Section 3.3, it follows that for all i and all $\alpha \in \mathcal{P}_i$, $\#\mathcal{P}_i \sum_{n=n_0^{(i, \alpha)}}^{n=\infty} \exp(-2ni^{-2}) \leq \exp(-i)$. These imply that for all $j \geq i_0 + 1$,

$$\begin{aligned} \mu \left(\bigcup_{i \geq j} \bigcup_{\alpha \in \mathcal{P}_i} \bigcup_{n \geq n_0^{(i, \alpha)}} \mathbf{B}_n^{(i, \alpha)} \right) &\leq \sum_{i \geq j} \sum_{\alpha \in \mathcal{P}_i} \sum_{n \geq n_0^{(i, \alpha)}} 2\#S \exp(-2ni^{-2}) \\ &\leq 2\#S \sum_{i \geq j} \#\mathcal{P}_i \sum_{n \geq n_0^{(i, \alpha)}} \exp(-2ni^{-2}) \\ &\leq 2\#S \sum_{i \geq j} \exp(-i) \\ &\leq 2\#S(1 - \exp(-1))^{-1} \exp(-j). \end{aligned}$$

From this inequality and the set inclusion in Step 2, it follows that for all $j \geq i_0 + 1$,

$$\mu \left(\bigcap_{T' \geq 0} \bigcup_{T \geq T'} \left\{ h_\infty \mid \|f_F(h_T) - f_\mu(h_T)\| \geq \frac{1}{m} \right\} \cap \mathbf{Z}_0 \right) \leq 2\#S(1 - \exp(-1))^{-1} \exp(-j).$$

Thus, letting $j \rightarrow \infty$, we have

$$\mu\left(\bigcap_{T' \geq 0} \bigcup_{T \geq T'} \left\{ h_\infty \mid \|f_F(h_T) - f_\mu(h_T)\| \geq \frac{1}{m} \right\} \cap \mathbf{Z}_0\right) = 0.$$

Note that the complement of \mathbf{Z}_0 is of μ -probability 0. Therefore, the above equality implies the desired result. \square

APPENDIX C

PROOF OF PROPOSITION 5. Fix any $\tilde{\mu}$ and let $f_{\tilde{\mu}}$ be the CPS corresponding to $\tilde{\mu}$. I show that $\tilde{\mu}$ does not almost weakly merge with any $\mu \notin \text{AG}(\{\mathcal{P}_{1/n}^{f_{\tilde{\mu}}}\}_n)$. Take any $\mu \notin \text{AG}(\{\mathcal{P}_{1/n}^{f_{\tilde{\mu}}}\}_n)$. Then there exists $\varepsilon_0 > 0$ such that for all n , all μ -probability 1 sets \mathbf{Z} , and all ordered families of time functions $\{T_m\}_m$, either there exist $\alpha \in \mathcal{P}_{1/n}^{f_{\tilde{\mu}}}$ and $h_T, h'_{T'} \in \alpha$ such that for some $h_\infty, h'_\infty \in \mathbf{Z}$ and some $m, m', h_T < h_\infty, T = T_m(h_\infty), h'_{T'} < h'_\infty, T' = T_{m'}(h'_\infty)$, and $\|f_\mu(h_T) - f_\mu(h'_{T'})\| > \varepsilon_0$; or $\liminf_{T \rightarrow \infty} N_T(h_\infty)/T < 1$ for some $h_\infty \in \mathbf{Z}$, where $N_T(h_\infty) := \#\{m \mid T_m(h_\infty) + 1 \leq T\}$.

Suppose that $\tilde{\mu}$ almost weakly merges with μ . Then, for $\varepsilon_0/4$, there exist a μ -probability 1 set \mathbf{Z}_0 and an ordered family $\{T_m^0\}_m$ of time functions such that for all $h_\infty \in \mathbf{Z}_0$ and all $m, \|f_{\tilde{\mu}}(h_{T_m^0}) - f_\mu(h_{T_m^0})\| \leq \varepsilon_0/4$, and that $\lim_{T \rightarrow \infty} N_T^0(h_\infty)/T = 1$ for all $h_\infty \in \mathbf{Z}_0$, where $N_T^0(h_\infty) := \#\{m \mid T_m^0(h_\infty) + 1 \leq T\}$. On the other hand, letting $n_0 \geq 4/\varepsilon_0$, it follows from the previous paragraph that for n_0, \mathbf{Z}_0 , and $\{T_m^0\}_m$, either there exist $\alpha \in \mathcal{P}_{1/n_0}^{f_{\tilde{\mu}}}$ and $h_T, h'_{T'} \in \alpha$ such that for some $h_\infty, h'_\infty \in \mathbf{Z}_0$ and some $m, m', h_T < h_\infty, T = T_m^0(h_\infty), h'_{T'} < h'_\infty, T' = T_{m'}^0(h'_\infty)$, and $\|f_\mu(h_T) - f_\mu(h'_{T'})\| > \varepsilon_0$; or $\liminf_{T \rightarrow \infty} N_T^0(h_\infty)/T < 1$ for some $h_\infty \in \mathbf{Z}_0$. Since $\lim_{T \rightarrow \infty} N_T^0(h_\infty)/T = 1$ for all $h_\infty \in \mathbf{Z}_0$, these imply that $\|f_{\tilde{\mu}}(h_T) - f_{\tilde{\mu}}(h'_{T'})\| \geq \varepsilon_0/2$. However, then, since $\alpha \in \mathcal{P}_{1/n_0}^{f_{\tilde{\mu}}}$ and $h_T, h'_{T'} \in \alpha, \|f_{\tilde{\mu}}(h_T) - f_{\tilde{\mu}}(h'_{T'})\| \leq 1/n_0 \leq \varepsilon_0/4$. This is a contradiction. Thus, $\tilde{\mu}$ does not almost weakly merge with μ . \square

Without loss of generality, we may assume (in the remainder of Appendix C) that $\{\mathcal{P}_i\}_i$ are ordered in fineness: $\mathcal{P}_i \leq \mathcal{P}_{i+1}$ for all i . Given h_T , let $n_T^{(i,\alpha)}$ be the number of times that category (i, α) has been effective up to time T , and let Γ_T denote the set of categories that have been effective (up to time T). Before proving Proposition 6, I show the following lemma. Lemma C(i) is the same as Lemma B, and Lemma C(ii) states that for almost all categories, the prior sample size is negligible relative to the number of effective periods.

LEMMA C. (i) $\lim_{T \rightarrow \infty} i(h_T) = \infty$ for all h_∞ and (ii) $\lim_{T \rightarrow \infty} \sum_{(i,\alpha) \in \Gamma_T} (n_T^{(i,\alpha)}/T) \times (n_0^{(i,\alpha)}/n_T^{(i,\alpha)}) = 0$ for all h_∞ .

PROOF. (i) Suppose that $\liminf_{T \rightarrow \infty} i(h_T) < \infty$ for some h_∞ . Then $i(h_{T_k}) = i_0$ for infinitely many T_k . Since \mathcal{P}_{i_0} and \mathcal{P}_{i_0+1} only have finite classes and $\mathcal{P}_{i_0} \leq \mathcal{P}_{i_0+1}$, there exist

$\alpha_0 \in \mathcal{P}_{i_0}$ and $\beta_0 \in \mathcal{P}_{i_0+1}$ such that $\beta_0 \subseteq \alpha_0$, and $\alpha(h_{T_{k_l}}) = \alpha_0$ and $\beta(h_{T_{k_l}}) = \beta_0$ for some infinite subsequence $\{T_{k_l}\}_l$ of $\{T_k\}_k$; $i(h_{T_{k_l}}) = i_0$ for all T_{k_l} . This, along with the definition of $m(\cdot)$, implies that $m(h_{T_{k_l}}) = i_0$ for all T_{k_l} . However, then, by the definition of $n(\cdot)$, $n(h_{T_{k_l}}) \rightarrow \infty$ as $T_{k_l} \rightarrow \infty$. This means that for some T_{k_l} , $n(h_{T_{k_l}}) \geq n_0^{\beta_0} = n_0^{\beta(h_{T_{k_l}})}$ and

$$\frac{\sum_{i=1}^{i=m(h_{T_{k_l}})+1} \sum_{\alpha \in \mathcal{P}_i} n_0^\alpha}{n(h_{T_{k_l}})} = \frac{\sum_{i=1}^{i=i_0+1} \sum_{\alpha \in \mathcal{P}_i} n_0^\alpha}{n(h_{T_{k_l}})} < \frac{1}{i_0} = \frac{1}{m(h_{T_{k_l}})}.$$

Therefore, by (4) in Section 5, $i(h_{T_{k_l}}) = m(h_{T_{k_l}}) + 1 = i_0 + 1$. This is a contradiction.

(ii) Let $i^*(h_T) := \max\{i(h_t) \mid t \leq T\}$, $t^*(h_T) := \min\{t \mid i(h_t) = i^*(h_T), t \leq T\}$, $n^*(h_T) := n(h_{t^*(h_T)})$, and $m^*(h_T) := m(h_{t^*(h_T)})$. Since $i(h_T) \rightarrow \infty$ as $T \rightarrow \infty$, $i^*(h_T) \rightarrow \infty$ as $T \rightarrow \infty$. Note that $i^*(h_T) = m^*(h_T) + 1$ and switching occurs at time $t^*(h_T) + 1$. Further note that if category (i, α) has been effective up to time $T + 1$, then $\alpha \in \bigcup_{j=1}^{j=i^*(h_T)} \mathcal{P}_j$. Thus, $\sum_{(i,\alpha) \in \Gamma_{T+1}} n_0^{(i,\alpha)} \leq \sum_{j=1}^{j=m^*(h_T)+1} \sum_{\alpha \in \mathcal{P}_j} n_0^\alpha$. Obviously, $n^*(h_T) \leq T + 1$. These imply that

$$\begin{aligned} \sum_{(i,\alpha) \in \Gamma_{T+1}} \frac{n_{T+1}^{(i,\alpha)} n_0^{(i,\alpha)}}{T+1 n_{T+1}^{(i,\alpha)}} &= \frac{\sum_{(i,\alpha) \in \Gamma_{T+1}} n_0^{(i,\alpha)}}{T+1} \leq \frac{\sum_{j=1}^{j=m^*(h_T)+1} \sum_{\alpha \in \mathcal{P}_j} n_0^\alpha}{n^*(h_T)} \\ &< \frac{1}{m^*(h_T)} = \frac{1}{i^*(h_T) - 1}. \end{aligned}$$

The first equality and the second inequality are obvious. The third inequality holds because switching occurs at time $t^*(h_T) + 1$; that is, the switching criterion (4) in Section 5 is passed. Since $i^*(h_T) \rightarrow \infty$ as $T \rightarrow \infty$, the desired result is obtained. \square

PROOF OF PROPOSITION 6. Let f_M be the modified frequency-based CPS for $\{\mathcal{P}_i\}_i$. Fix any $\mu \in \text{AG}(\{\mathcal{P}_i\}_i)$. Suppose that $\tilde{\mu}_M$ does not almost weakly merge with μ .

Step 1. On the one hand, since $\tilde{\mu}_M$ does not almost weakly merge with μ , there exists $\varepsilon_0 > 0$ such that for any μ -probability 1 set \mathbf{Z} , there exist $h_\infty \in \mathbf{Z}$ and a set \mathbb{I} of nonnegative integers such that for all $T \in \mathbb{I}$, $\|f_M(h_T) - f_\mu(h_T)\| > \varepsilon_0$, and that $\limsup_{T \rightarrow \infty} \#\{\mathbb{I} \cap \{0, 1, \dots, T-1\}\} / T > 0$.

Step 2. On the other hand, since $\mu \in \text{AG}(\{\mathcal{P}_i\}_i)$, for all $\varepsilon > 0$, there exist an index i_0 , a μ -probability 1 set \mathbf{Z}_0 , and an ordered family $\{T_m^0\}_m$ of time functions such that (i) for all $\alpha \in \mathcal{P}_{i_0}$ and all $h_T, h_{T'} \in \alpha$, if there exist $h_\infty, h'_\infty \in \mathbf{Z}_0$ and m, m' such that $h_T < h_\infty$, $T = T_m^0(h_\infty)$, $h_{T'} < h'_\infty$, and $T' = T_{m'}^0(h'_\infty)$, then $\|f_\mu(h_T) - f_\mu(h_{T'})\| \leq \varepsilon$; and that (ii) $\lim_{T \rightarrow \infty} N_T^0(h_\infty) / T = 1$ for all $h_\infty \in \mathbf{Z}_0$, where $N_T^0(h_\infty) := \#\{m \mid T_m^0(h_\infty) + 1 \leq T\}$.

Step 3. Let $\varepsilon := \varepsilon_0 / 4$. From Step 2, it follows that for ε , there exist an index i_0 , a μ -probability 1 set \mathbf{Z}_0 , and an ordered family $\{T_m^0\}_m$ of time functions such that (i) and (ii) hold. Since $\mathcal{P}_i \subseteq \mathcal{P}_{i+1}$ for all i , we may take $i_0 \geq 4/\varepsilon_0$. For all $\alpha \in \mathcal{P}_{i_0}$, define a class $\hat{\alpha}$ as follows: $h_T \in \hat{\alpha}$ if and only if (a) $h_T \in \alpha$ and (b) $h_T < h_\infty$ and $T = T_m^0(h_\infty)$ for some $h_\infty \in \mathbf{Z}_0$ and some m . Then, for all $\alpha \in \mathcal{P}_{i_0}$ and all s , let $L^\alpha[s] := \sup_{h \in \hat{\alpha}} f_\mu(h)[s]$ and $l^\alpha[s] := \inf_{h \in \hat{\alpha}} f_\mu(h)[s]$; note that $L^\alpha[s] - l^\alpha[s] \leq \varepsilon = \varepsilon_0 / 4$ for all s . Furthermore, for all (i, α) , define a class $\tilde{\gamma}(i, \alpha)$ as follows: $h_T \in \tilde{\gamma}(i, \alpha)$ if and only if (a) either time $T + 1$ is an

effective period of (i, α) , that is, $(i(h_T), \alpha(h_T)) = (i, \alpha)$, or time $T + 1$ is one of the first n_0^α effective periods of (i_p, α_p) ; and (b) $h_T < h_\infty$ and $T = T_m^0(h_\infty)$ for some $h_\infty \in \mathbf{Z}_0$ and some m . Since $\mathcal{P}_i \leq \mathcal{P}_{i+1}$ for all i , for each (i, α) with $i \geq i_0 + 1$, there exists a unique class $\beta \in \mathcal{P}_{i_0}$ such that $\bar{\gamma}(i, \alpha) \subseteq \hat{\beta}$; let $L^{(i,\alpha)}[s] := L^\beta[s]$ and $l^{(i,\alpha)}[s] := l^\beta[s]$ for all s . Thus, for all $h \in \bar{\gamma}(i, \alpha)$ and all s , $l^{(i,\alpha)}[s] \leq f_\mu(h)[s] (= \mu(s|h)) \leq L^{(i,\alpha)}[s]$.

Step 4. For all (i, α) with $i \geq i_0 + 1$, let

$$\mathbf{C}_n^{(i,\alpha)} := \left\{ h_\infty \mid \mathcal{T}_n^{\bar{\gamma}(i,\alpha)} < \infty, \exists s \left(\frac{\mathbf{d}_n^{\bar{\gamma}(i,\alpha)}[s]}{n} \leq l^{(i,\alpha)}[s] - \frac{1}{i} \text{ or } \frac{\mathbf{d}_n^{\bar{\gamma}(i,\alpha)}[s]}{n} \geq L^{(i,\alpha)}[s] + \frac{1}{i} \right) \right\}.$$

Then, from Step 3 and Proposition A (in Appendix A), it follows that for all (i, α) with $i \geq i_0 + 1$, $\mu(\mathbf{C}_n^{(i,\alpha)}) \leq 2\#S \exp(-2ni^{-2})$ for all n . In addition, by (2) in Section 3.3, $\#\mathcal{P}_i \sum_{n=n_0^\alpha}^{n=\infty} \exp(-2ni^{-2}) \leq \exp(-i)$ for all $\alpha \in \mathcal{P}_i$ and all i . These imply that for all $j \geq i_0 + 1$,

$$\mu\left(\bigcap_{j \geq i_0+1} \bigcup_{i \geq j} \bigcup_{\alpha \in \mathcal{P}_i} \bigcup_{n \geq n_0^\alpha} \mathbf{C}_n^{(i,\alpha)}\right) \leq \mu\left(\bigcup_{i \geq j} \bigcup_{\alpha \in \mathcal{P}_i} \bigcup_{n \geq n_0^\alpha} \mathbf{C}_n^{(i,\alpha)}\right) \leq 2\#S(1 - \exp(-1))^{-1} \exp(-j).$$

Thus, letting $j \rightarrow \infty$, we have $\mu(\bigcap_{j \geq i_0+1} \bigcup_{i \geq j} \bigcup_{\alpha \in \mathcal{P}_i} \bigcup_{n \geq n_0^\alpha} \mathbf{C}_n^{(i,\alpha)}) = 0$.

Step 5. Let $\mathbf{C} := \bigcup_{j \geq i_0+1} \bigcap_{i \geq j} \bigcap_{\alpha \in \mathcal{P}_i} \bigcap_{n \geq n_0^\alpha} (\mathbf{C}_n^{(i,\alpha)})^c$, where $(\mathbf{C}_n^{(i,\alpha)})^c$ is the complement of $\mathbf{C}_n^{(i,\alpha)}$. From Steps 3 and 4, $\mu(\mathbf{C} \cap \mathbf{Z}_0) = 1$. Thus, by Step 1, for $\mathbf{C} \cap \mathbf{Z}_0$, there exist $h_\infty \in \mathbf{C} \cap \mathbf{Z}_0$ and a set \mathbb{I} of nonnegative integers such that for all $T \in \mathbb{I}$, $\|f_M(h_T) - f_\mu(h_T)\| > \varepsilon_0$, and that $\limsup_{T \rightarrow \infty} \#\mathbb{I} \cap \{0, 1, \dots, T-1\} / T > 0$. Let $I_T^{(i,\alpha)}$ denote the number of effective periods of category (i, α) in which $\|f_M(h) - f_\mu(h)\| > \varepsilon_0$ (up to time T). Recall that Γ_T is the set of categories that have been effective up to time T . Then the above statement is equivalent to $\limsup_{T \rightarrow \infty} \sum_{(i,\alpha) \in \Gamma_T} (n_T^{(i,\alpha)} / T) (I_T^{(i,\alpha)} / n_T^{(i,\alpha)}) > 0$. Letting $\Gamma_T^l(\delta) := \{(i, \alpha) \mid I_T^{(i,\alpha)} / n_T^{(i,\alpha)} \geq \delta\}$, this, in turn, implies that there exists $\delta_0 > 0$ such that for infinitely many T_l , $\sum_{(i,\alpha) \in \Gamma_{T_l}^l(\delta_0)} n_{T_l}^{(i,\alpha)} / T_l > 2\delta_0$.

Step 6. Since $h_\infty \in \mathbf{Z}_0$, $\lim_{T \rightarrow \infty} N_T^0(h_\infty) / T = 1$ by Steps 2 and 3. Let $J_T^{(i,\alpha)}$ be the number of times that (for some m) time $T_m^0(h_\infty) + 1$ has been an effective period of (i, α) (up to time T). Note that (for all T), $\sum_{(i,\alpha) \in \Gamma_T} J_T^{(i,\alpha)} = N_T^0(h_\infty)$ by the definitions of $J_T^{(i,\alpha)}$ and $N_T^0(h_\infty)$. Therefore, $\lim_{T \rightarrow \infty} \sum_{(i,\alpha) \in \Gamma_T} (n_T^{(i,\alpha)} / T) (J_T^{(i,\alpha)} / n_T^{(i,\alpha)}) = 1$. This means that for any $\eta > 0$, there exists l_0 such that for all $l \geq l_0$, $\sum_{(i,\alpha) \in \Gamma_{T_l}^l(\eta)} n_{T_l}^{(i,\alpha)} / T_l \geq 1 - \eta$, where $\Gamma_{T_l}^l(\eta) := \{(i, \alpha) \mid J_{T_l}^{(i,\alpha)} / n_{T_l}^{(i,\alpha)} \geq 1 - \eta\}$. Moreover, from Lemma C(ii), it follows that $\lim_{T \rightarrow \infty} \sum_{(i,\alpha) \in \Gamma_T} (n_T^{(i,\alpha)} / T) (n_0^{(i,\alpha)} / n_T^{(i,\alpha)}) = 0$. This implies that for any $\eta > 0$, there exists l_1 such that for all $l \geq l_1$, $\sum_{(i,\alpha) \in \Gamma_{T_l}^K(\eta)} n_{T_l}^{(i,\alpha)} / T_l \geq 1 - \eta$, where $\Gamma_{T_l}^K(\eta) := \{(i, \alpha) \mid n_0^{(i,\alpha)} / n_{T_l}^{(i,\alpha)} \leq \eta\}$.

Step 7. It follows from Steps 5 and 6 that for any sufficiently small $\eta > 0$, there exists l_2 such that for all $l \geq l_2$, $\sum_{(i,\alpha) \in \Gamma_{T_l}^l(\delta_0) \cap \Gamma_{T_l}^l(\eta) \cap \Gamma_{T_l}^K(\eta)} n_{T_l}^{(i,\alpha)} / T_l > \delta_0$: for any sufficiently large T_l , $\Gamma_{T_l}^l(\delta_0) \cap \Gamma_{T_l}^l(\eta) \cap \Gamma_{T_l}^K(\eta) \neq \emptyset$. This, along with $\lim_{T \rightarrow \infty} i(h_T) = \infty$ from Lemma C(i), implies that for any sufficiently small $\eta > 0$, any T , and any $j \geq i_0 + 1$, there exist $\bar{T} \geq T$ and $(\bar{i}, \bar{\alpha})$ with $\bar{i} \geq j$ such that (i) time $\bar{T} + 1$ is an effective period of $(\bar{i}, \bar{\alpha})$, that

is, $i(h_{\bar{T}}) = \bar{i}$ and $\alpha(h_{\bar{T}}) = \bar{\alpha}$, (ii) $\bar{T} = T_m^0(h_\infty)$ for some m , (iii) $\|f_M(h_{\bar{T}}) - f_\mu(h_{\bar{T}})\| > \varepsilon_0$, (iv) $J_{\bar{T}}^{(\bar{i}, \bar{\alpha})}/n_{\bar{T}}^{(\bar{i}, \bar{\alpha})} \geq (\delta_0 - \eta)/\delta_0$, and (v) $n_0^{(\bar{i}, \bar{\alpha})}/n_{\bar{T}}^{(\bar{i}, \bar{\alpha})} \leq \eta/(\delta_0 - \eta)$. Furthermore, note that $f_M(h_{\bar{T}}) = D_{\bar{T}}^{(\bar{i}, \bar{\alpha})}$ and

$$D_{\bar{T}}^{(\bar{i}, \bar{\alpha})} = \frac{d_{\bar{T}}^{(\bar{i}, \bar{\alpha})} + d_0^{(\bar{i}, \bar{\alpha})}}{n_{\bar{T}}^{(\bar{i}, \bar{\alpha})} + n_0^{(\bar{i}, \bar{\alpha})}} = \frac{n}{n_{\bar{T}}^{(\bar{i}, \bar{\alpha})} + n_0^{(\bar{i}, \bar{\alpha})}} \frac{\mathbf{d}_n^{\bar{y}(\bar{i}, \bar{\alpha})}}{n} + \frac{n_{\bar{T}}^{(\bar{i}, \bar{\alpha})} + n_0^{(\bar{i}, \bar{\alpha})} - n}{n_{\bar{T}}^{(\bar{i}, \bar{\alpha})} + n_0^{(\bar{i}, \bar{\alpha})}} \frac{d_{\bar{T}}^{(\bar{i}, \bar{\alpha})} + d_0^{(\bar{i}, \bar{\alpha})} - \mathbf{d}_n^{\bar{y}(\bar{i}, \bar{\alpha})}}{n_{\bar{T}}^{(\bar{i}, \bar{\alpha})} + n_0^{(\bar{i}, \bar{\alpha})} - n},$$

where n is the number of times that (for some m), time $T_m^0(h_\infty) + 1$ has been either an effective period of $(\bar{i}, \bar{\alpha})$ or one of the first $n_0^{(\bar{i}, \bar{\alpha})}$ effective periods of its predecessor $(\bar{i}_p, \bar{\alpha}_p)$ (up to time \bar{T}). Hence, $n \geq J_{\bar{T}}^{(\bar{i}, \bar{\alpha})}$ by the definitions of n and $J_{\bar{T}}^{(\bar{i}, \bar{\alpha})}$. From this and (iv) and (v), it follows that $n \geq J_{\bar{T}}^{(\bar{i}, \bar{\alpha})} \geq (\delta_0 - \eta)n_{\bar{T}}^{(\bar{i}, \bar{\alpha})}/\delta_0 \geq (\delta_0 - \eta)^2 n_0^{(\bar{i}, \bar{\alpha})}/\delta_0 \eta$ and $n/(n_{\bar{T}}^{(\bar{i}, \bar{\alpha})} + n_0^{(\bar{i}, \bar{\alpha})}) \geq (\delta_0 - \eta)J_{\bar{T}}^{(\bar{i}, \bar{\alpha})}/\delta_0 n_{\bar{T}}^{(\bar{i}, \bar{\alpha})} \geq (\delta_0 - \eta)^2/(\delta_0)^2$. Therefore, for any sufficiently small $\eta > 0$, $n \geq n_0^{(\bar{i}, \bar{\alpha})}$ ($= n_0^{\bar{\alpha}}$) and $\|f_M(h_{\bar{T}}) - \mathbf{d}_n^{\bar{y}(\bar{i}, \bar{\alpha})}/n\| < \varepsilon_0/4$. Furthermore, recall that $i_0 \geq 4/\varepsilon_0$ and $\bar{i} \geq i_0 + 1$. These, along with (i), (ii), (iii), and Step 3, imply that for some s ,

$$\frac{\mathbf{d}_n^{\bar{y}(\bar{i}, \bar{\alpha})}[s]}{n} \leq l^{(\bar{i}, \bar{\alpha})}[s] - \frac{1}{\bar{i}} \quad \text{or} \quad \frac{\mathbf{d}_n^{\bar{y}(\bar{i}, \bar{\alpha})}[s]}{n} \geq L^{(\bar{i}, \bar{\alpha})}[s] + \frac{1}{\bar{i}}.$$

Thus, $h_\infty \in (\mathbf{C})^c$, where $(\mathbf{C})^c$ is the complement of \mathbf{C} . However, then, from Step 5, it follows that $h_\infty \in \mathbf{C} \cap \mathbf{Z}_0$. This is a contradiction to $\mathbf{C} \cap (\mathbf{C})^c = \emptyset$. Therefore, $\tilde{\mu}_M$ almost weakly merges with μ . \square

APPENDIX D

PROOF OF THE MERGED SET BEING OF FIRST CATEGORY. Let M denote any merged set and let $\sum_n w_n \mu_n$ denote any strictly convex combination of a countable dense subset $\{\mu_n\}_n$ of M , where $\sum_n w_n = 1$ and $w_n > 0$ for all n . Proposition 1(1) in [Dekel and Feinberg \(2006\)](#) states that any probability measure has a first category set of probability 1. Therefore, $\sum_n w_n \mu_n$ has a first category set \mathbf{Z}_I of probability 1: $\sum_n w_n \mu_n(\mathbf{Z}_I) = 1$. Hence, $\mu_n(\mathbf{Z}_I) = 1$ for all n . Since $\{\mu_n\}_n$ is dense in M , for any $\mu \in M$, there exists n such that $|\mu(\mathbf{Z}_I) - \mu_n(\mathbf{Z}_I)| \leq \frac{1}{2}$. Thus, $\mu(\mathbf{Z}_I) \geq \frac{1}{2}$ for all $\mu \in M$. However, Proposition 1(2) in [Dekel and Feinberg \(2006\)](#) states that for any first category set \mathbf{D} (of infinite histories), the set of probability measures that put *positive* probability on \mathbf{D} is of first category (in the weak (-star) topology). Hence, M is also a first category set. \square

REFERENCES

- Abreu, Dilip (1988), "On the theory of infinitely repeated games with discounting." *Econometrica*, 56, 383–396. [426]
- Anderson, Robert M. and William R. Zame (2001), "Genericity with infinitely many parameters." *Advances in Theoretical Economics*, 1, 1–62. [433]

Aumann, Robert J. (1964), "Mixed and behavior strategies in infinite extensive games." In *Advances in Game Theory* (Melvin Dresher, Lloyd S. Shapley, and Albert W. Tucker, eds.), 627–650, Princeton University Press, Princeton. [422]

Blackwell, David and Lester Dubins (1962), "Merging of opinions with increasing information." *The Annals of Mathematical Statistics*, 33, 882–886. [413]

Cesa-Bianchi, Nicolò and Gábor Lugosi (2006), *Prediction, Learning, and Games*. Cambridge University Press, New York. [413, 431]

Dekel, Eddie and Yossi Feinberg (2006), "Non-Bayesian testing of a stochastic prediction." *Review of Economic Studies*, 73, 893–906. [432, 433, 441]

Fudenberg, Drew and David K. Levine (1998), *The Theory of Learning in Games*. MIT Press, Cambridge, Massachusetts. [424]

Fudenberg, Drew and David K. Levine (1999), "Conditional universal consistency." *Games and Economic Behavior*, 29, 104–130. [424]

Fudenberg, Drew and Eric S. Maskin (1991), "On the dispensability of public randomization in discounted repeated games." *Journal of Economic Theory*, 53, 428–438. [426]

Gilboa, Itzhak, Andrew Postlewaite, and David Schmeidler (2004), "Rationality of belief or: Why Bayesianism is neither necessary nor sufficient for rationality." PIER Working Paper 04-011, Cowles Foundation Discussion Paper 1484. [413]

Hörner, Johannes and Wojciech Olszewski (2006), "The folk theorem for games with private almost-perfect monitoring." *Econometrica*, 74, 1499–1544. [426]

Jackson, Matthew O., Ehud Kalai, and Rann Smorodinsky (1999), "Bayesian representation of stochastic processes under learning: de Finetti revisited." *Econometrica*, 67, 875–893. [413]

Kalai, Ehud and Ehud Lehrer (1993), "Rational learning leads to Nash equilibrium." *Econometrica*, 61, 1019–1045. [411, 413, 422]

Kalai, Ehud and Ehud Lehrer (1994), "Weak and strong merging of opinions." *Journal of Mathematical Economics*, 23, 73–86. [411, 413, 433]

Kalai, Ehud, Ehud Lehrer, and Rann Smorodinsky (1999), "Calibrated forecasting and merging." *Games and Economic Behavior*, 29, 151–169. [413]

Lehrer, Ehud and Rann Smorodinsky (1996a), "Merging and learning." In *Statistics, Probability and Game Theory: Papers in Honor of David Blackwell* (Lloyd S. Shapley, Thomas S. Ferguson, and James B. McQueen, eds.), 147–168, Springer, New York. [414]

Lehrer, Ehud and Rann Smorodinsky (1996b), "Compatible measures and merging." *Mathematics of Operations Research*, 21, 697–706. [429]

Lehrer, Ehud and Rann Smorodinsky (1997), "Repeated large games with incomplete information." *Games and Economic Behavior*, 18, 116–134. [412]

- Mailath, George J. and Larry Samuelson (2006), *Repeated Games and Reputations*. Oxford University Press, New York. [425]
- Miller, Ronald I. and Chris W. Sanchirico (1997), “Almost everybody disagrees almost all the time: The genericity of weakly merging nowhere.” Discussion Paper 9697-25, Economics Department, Columbia University. [432]
- Nachbar, John H. (1997), “Prediction, optimization, and learning in repeated games.” *Econometrica*, 65, 275–309. [411, 412, 426, 427, 431]
- Nachbar, John H. (2005), “Beliefs in repeated games.” *Econometrica*, 73, 459–480. [411, 412, 422, 426, 427, 431]
- Nachbar, John H. (2009), “Learning in games.” In *Encyclopedia of Complexity and Systems Science* (Robert A. Meyers, ed.), 5177–5187, Springer, New York. [422]
- Nachbar, John H. and William R. Zame (1996), “Non-computable strategies and discounted repeated games.” *Economic Theory*, 8, 103–122. [426]
- Noguchi, Yuichi (2000), “Optimal properties of conditional fictitious play.” In *Three Essays on Economic Dynamics and Learning in Games*, Ph.D. thesis, Harvard University. [413, 419, 424, 431]
- Noguchi, Yuichi (2003), “Optimal properties of conditional adaptive strategies.” Unpublished paper, Hitotsubashi University. [431]
- Noguchi, Yuichi (2009), “Note on universal conditional consistency.” *International Journal of Game Theory*, 38, 193–207. [413, 419, 424, 431]
- Noguchi, Yuichi (2015), “Bayesian learning, smooth approximate optimal behavior, and convergence to ε -Nash equilibrium.” *Econometrica*, 83, 353–373. [412, 431]
- Osborne, Martin J. and Ariel Rubinstein (1994), *A Course in Game Theory*. MIT Press, Cambridge, Massachusetts. [425]
- Parthasarathy, Kalyanapuram R. (1967), *Probability Measures on Metric Spaces*. Academic Press, New York. [432]
- Ryabko, Boris Y. (1988), “Prediction of random sequences and universal coding.” *Problems of Information Transmission*, 24, 87–96. [413]
- Ryabko, Daniil (2010), “On finding predictors for arbitrary families of processes.” *Journal of Machine Learning Research*, 11, 581–602. [413, 433]
- Ryabko, Daniil (2011), “On the relation between realizable and nonrealizable cases of the sequence prediction problem.” *Journal of Machine Learning Research*, 12, 2161–2180. [413]
- Ryabko, Daniil and Marcus Hutter (2008), “Predicting non-stationary processes.” *Applied Mathematics Letters*, 21, 477–482. [413]
- Sandroni, Alvaro (1998), “Necessary and sufficient conditions for convergence to Nash equilibrium: The almost absolute continuity hypothesis.” *Games and Economic Behavior*, 22, 121–147. [412, 433]

Sandroni, Alvaro and Rann Smorodinsky (1999), “The speed of rational learning.” *International Journal of Game Theory*, 28, 199–210. [433]

Shiryaev, Albert N. (1996), *Probability*, second edition. Springer, New York. [415, 436]

Solomonoff, Ray J. (1978), “Complexity-based induction systems: Comparisons and convergence theorems.” *IEEE Transactions on Information Theory*, 24, 422–432. [413]

Sorin, Sylvain (1999), “Merging, reputation, and repeated games with incomplete information.” *Games and Economic Behavior*, 29, 274–308. [413]

Submitted 2012-10-8. Final version accepted 2014-5-10. Available online 2014-5-15.