

Article

Predictive Constructions Based on Measure-Valued Pólya Urn Processes

Sandra Fortini ^{1,*}, Sonia Petrone ¹ and Hristo Sariev ²¹ Department of Decision Sciences, Bocconi University, 20136 Milano, Italy; sonia.petrone@unibocconi.it² Institute of Mathematics and Informatics, Bulgarian Academy of Sciences, 1113 Sofia, Bulgaria; h.sariev@math.bas.bg

* Correspondence: sandra.fortini@unibocconi.it

Abstract: Measure-valued Pólya urn processes (MVPP) are Markov chains with an additive structure that serve as an extension of the generalized k -color Pólya urn model towards a continuum of possible colors. We prove that, for any MVPP $(\mu_n)_{n \geq 0}$ on a Polish space \mathbb{X} , the normalized sequence $(\mu_n / \mu_n(\mathbb{X}))_{n \geq 0}$ agrees with the marginal predictive distributions of some random process $(X_n)_{n \geq 1}$. Moreover, $\mu_n = \mu_{n-1} + R_{X_n}$, $n \geq 1$, where $x \mapsto R_x$ is a random transition kernel on \mathbb{X} ; thus, if μ_{n-1} represents the contents of an urn, then X_n denotes the color of the ball drawn with distribution $\mu_{n-1} / \mu_{n-1}(\mathbb{X})$ and R_{X_n} —the subsequent reinforcement. In the case $R_{X_n} = W_n \delta_{X_n}$, for some non-negative random weights W_1, W_2, \dots , the process $(X_n)_{n \geq 1}$ is better understood as a randomly reinforced extension of Blackwell and MacQueen's Pólya sequence. We study the asymptotic properties of the predictive distributions and the empirical frequencies of $(X_n)_{n \geq 1}$ under different assumptions on the weights. We also investigate a generalization of the above models via a randomization of the law of the reinforcement.



Citation: Fortini, S.; Petrone, S.; Sariev, H. Predictive Constructions Based on Measure-Valued Pólya Urn Processes. *Mathematics* **2021**, *9*, 2845. <https://doi.org/10.3390/math9222845>

Keywords: predictive distributions; random probability measures; reinforced processes; Pólya sequences; urn schemes; Bayesian inference; conditional identity in distribution; total variation distance

MSC: 60G57; 60B10; 60G25; 60F05; 60G09

Academic Editors: Emanuele Dolera and Federico Bassetti

Received: 4 October 2021
Accepted: 8 November 2021
Published: 10 November 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

1. Introduction

Let $(X_n)_{n \geq 1}$ be a sequence of homogeneous random observations, taking values in a Polish space \mathbb{X} . The central assumption in the Bayesian approach to inductive reasoning is that $(X_n)_{n \geq 1}$ is exchangeable, that is, its law is invariant under finite permutations. Then, by de Finetti's theorem, there exists a random probability measure \tilde{P} on \mathbb{X} such that, given \tilde{P} , the random variables X_1, X_2, \dots are conditionally independent and identically distributed with marginal distribution \tilde{P} (see [1], Section 3), denoted

$$X_n \mid \tilde{P} \stackrel{i.i.d.}{\sim} \tilde{P}. \quad (1)$$

Furthermore, \tilde{P} is the almost sure (a.s.) weak limit of the predictive distributions and the empirical frequencies,

$$\mathbb{P}(X_{n+1} \in \cdot \mid X_1, \dots, X_n) \xrightarrow{w} \tilde{P}(\cdot) \quad \text{a.s.} \quad \text{and} \quad \frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\cdot) \xrightarrow{w} \tilde{P}(\cdot) \quad \text{a.s.} \quad (2)$$

The model (1) is completed by choosing a prior distribution for \tilde{P} . Inference consists in computing the conditional (posterior) distribution of \tilde{P} given an observed sample (X_1, \dots, X_n) , with most inferential conclusions depending on some average with respect to the posterior distribution; for example, under squared loss, for any measurable set



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

$B \subseteq \mathbb{X}$, the best estimate of $\tilde{P}(B)$ is the posterior mean, $\mathbb{E}[\tilde{P}(B)|X_1, \dots, X_n]$. In addition, the posterior mean can be utilized for predictive inference since

$$\mathbb{P}(X_{n+1} \in B|X_1, \dots, X_n) = \mathbb{E}[\tilde{P}(B)|X_1, \dots, X_n]. \tag{3}$$

A different modeling strategy uses the Ionescu–Tulcea theorem to define the law of the process from the sequence of predictive distributions, $(\mathbb{P}(X_{n+1} \in \cdot|X_1, \dots, X_n))_{n \geq 1}$. In that case, one can refer to Theorem 3.1 in [2] for necessary and sufficient conditions on $(\mathbb{P}(X_{n+1} \in \cdot|X_1, \dots, X_n))_{n \geq 1}$ to be consistent with exchangeability. The predictive approach to model building is deeply rooted in Bayesian statistics, where the parameter \tilde{P} is assigned an auxiliary role and the focus is on observable “facts”, see [2–6]. Moreover, using the predictive distributions as primary objects allows one to make predictions instantly or helps ease computations. See [7] for a review on some well-known predictive constructions of priors for Bayesian inference.

In this work, we consider a class of predictive constructions based on measure-valued Pólya urn processes (MVPP). MVPPs have been introduced in the probabilistic literature [8,9] as an extension of k -color urn models, but their implications for (Bayesian) statistics have yet to be explored. A first aim of the paper is thus to show the potential use of MVPPs as predictive constructions in Bayesian inference. In fact, some popular models in Bayesian nonparametric inference can be framed in such a way, see Equation (8). A second aim of the paper is to suggest novel extensions of MVPPs that we believe can offer more flexibility in statistical applications.

MVPPs are essentially measure-valued Markov processes that have an additive structure, with the formal definition being postponed to Section 2.1 (Definition 1). Given an MVPP $(\mu_n)_{n \geq 0}$, we consider a sequence of random observations that are characterized by $\mathbb{P}(X_1 \in \cdot) = \mu_0(\cdot)/\mu_0(\mathbb{X})$ and, for $n \geq 1$,

$$\mathbb{P}(X_{n+1} \in \cdot | X_1, \mu_1, \dots, X_n, \mu_n) = \frac{\mu_n(\cdot)}{\mu_n(\mathbb{X})}. \tag{4}$$

The random measure μ_n is not necessarily measurable with respect to (X_1, \dots, X_n) , so the predictive construction (4) is more flexible than models based solely on the predictive distributions of $(X_n)_{n \geq 1}$; for example, $(\mu_n)_{n \geq 0}$ allows for the presence of latent variables or other sources of observable data (see also [10] for a covariate-based predictive construction). However, (4) can lead to an imbalanced design, which may break the symmetry imposed by exchangeability. Nevertheless, it is still possible that the sequence $(X_n)_{n \geq 1}$ satisfies (2) for some \tilde{P} , in which case Lemma 8.2 in [1] implies that $(X_n)_{n \geq 1}$ is asymptotically exchangeable with directing random measure \tilde{P} .

In Theorem 1, we show that, taking $(\mu_n)_{n \geq 0}$ as primary, the sequence $(X_n)_{n \geq 1}$ in (4) can be chosen such that

$$\mu_n = \mu_{n-1} + R_{X_n}, \tag{5}$$

where $x \mapsto R_x$ is a measurable map from \mathbb{X} to the space of finite measures on \mathbb{X} . Models of the kind (4)–(5) are computationally efficient. Indeed, as new observations become available, predictions can be updated at a constant computational cost and with limited storage of information. If, in addition, $(X_n)_{n \geq 1}$ is asymptotically exchangeable, then (4)–(5) can provide a computationally simple approximation of an exchangeable scheme for Bayesian inference, along the lines in [11].

The recursive formula (5) allows us to interpret the dynamics of MVPPs in terms of an urn sampling scheme, as the name suggests. Let μ_0 be a non-random finite measure on \mathbb{X} . Suppose we have an urn whose contents are described by μ_0 in the sense that $\mu_0(B)$ denotes the total mass of balls with colors in $B \subseteq \mathbb{X}$. At time $n = 1$, a ball is extracted at random from the urn, and we denote its color by X_1 . The urn is then reinforced according to a replacement rule $(R_x)_{x \in \mathbb{X}}$, so that the updated composition becomes $\mu_1 \equiv \mu_0 + R_{X_1}$. At any time $n > 1$, a ball of color X_n is picked with probability distribution $\mu_{n-1}/\mu_{n-1}(\mathbb{X})$, and the contents of the urn are subsequently reinforced by R_{X_n} . In the case the space of

colors is finite, $|\mathbb{X}| = k$, the above procedure is better known as a generalized k -color Pólya urn [12].

We focus our analysis on MVPPs for which R_x is concentrated on x ; thus, after each draw, we reinforce only the color of the observed ball. More formally, we consider MVPPs that have a reinforcement measure of the kind $R_{X_n} = W_n \delta_{X_n}$, $n \geq 1$, where W_n is some non-negative random variable. In that case, Equations (4) and (5) become

$$\mathbb{P}(X_{n+1} \in \cdot \mid X_1, W_1, \dots, X_n, W_n) = \sum_{i=1}^n \frac{W_i}{\mu_0(\mathbb{X}) + \sum_{j=1}^n W_j} \delta_{X_i}(\cdot) + \frac{\mu_0(\mathbb{X})}{\mu_0(\mathbb{X}) + \sum_{j=1}^n W_j} \mu'_0(\cdot), \tag{6}$$

and

$$\mu_n = \mu_{n-1} + W_n \delta_{X_n}. \tag{7}$$

A notable example is Blackwell and MacQueen’s em Pólya sequence [13], which is a random process $(X_n)_{n \geq 1}$ characterized by $\mathbb{P}(X_1 \in \cdot) = \nu(\cdot)$ and, for $n \geq 1$,

$$\mathbb{P}(X_{n+1} \in \cdot \mid X_1, \dots, X_n) = \sum_{i=1}^n \frac{1}{\theta + n} \delta_{X_i}(\cdot) + \frac{\theta}{\theta + n} \nu(\cdot), \tag{8}$$

for some probability measure ν on \mathbb{X} and a constant $\theta > 0$. By [13], $(X_n)_{n \geq 1}$ is exchangeable and corresponds to the model (1) with Dirichlet process prior with parameters (θ, ν) . It is easily seen that (8) is related to the MVPP $(\mu_n)_{n \geq 0}$ given by $\mu_0 = \theta \nu$ and, for $n \geq 1$,

$$\mu_n = \mu_{n-1} + \delta_{X_n}.$$

Therefore, we will call any MVPP a randomly reinforced Pólya process (RRPP) if it admits representation (6)–(7).

Existing studies on MVPPs look at models that have mostly a balanced design, i.e., $R_x(\mathbb{X}) = r$, $x \in \mathbb{X}$, and assume irreducibility-like conditions for $(R_x)_{x \in \mathbb{X}}$, see [8,9,14,15] and Remark 4 in [16]. In contrast, RRPPs require that $R_x(\{x\}^c) = 0$, and so are excluded from the analysis in those papers. In fact, this difference in reinforcement mechanisms mirrors the dichotomy within k -color urn models, where the replacement R is best described in terms of a matrix with random elements. There, the class of randomly reinforced urns [17] assumes an R with zero off-diagonal elements (i.e., we reinforce only the color of the observed ball), whereas the generalized Pólya urn models require the mean replacement matrix to be irreducible. Similarly to the k -color case, RRPPs need the use of different techniques, which yield completely different results than those in [8,9,14–16]. As an example, Theorem 1 in [16] and our Theorem 2 prove convergence of the kind (2), yet the limit probability measure in [16] is non-random.

The RRPP has been implicitly studied by [17–23], among others, with the focus being on the process $(X_n)_{n \geq 1}$. Those papers deal primarily with the k -color case (with the exception of [18,19,23]) and can be categorized on the basis of their assumptions on $(W_n)_{n \geq 1}$. For example, [18,19,21,22] assume that W_n and $(X_1, W_1, \dots, X_{n-1}, W_{n-1}, X_n)$ are independent, in which case the process $(X_n)_{n \geq 1}$ is conditionally identically distributed (c.i.d.) [21], that is, conditionally on current information, all future observations are identically distributed. It follows from [21] that c.i.d. processes preserve many of the properties of exchangeable sequences and, in particular, satisfy (2)–(3). In contrast, [17,20,23] assume that the reinforcement W_n depends on the particular color X_n , and prove a version of (2) where \tilde{P} is concentrated on the set of dominant colors for which the expected reinforcement is maximum. In this work, we reconsider the above models in the framework of RRPPs. For the c.i.d. case, we prove results whose analogues have already been established by [23] for the model with dominant colors. In particular, we extend the convergence in (2) to be in total variation and give a unified central limit theorem. We also examine the number of distinct values that are generated by the sequence $(X_n)_{n \geq 1}$.

In some applications, the definition of an MVPP can be too restrictive as it assumes that the probability law of the reinforcement R is known. However, we can envisage situations where the law is itself random, so we extend the definition of an MVPP by introducing

a random parameter V . The resulting generalized measure-valued Pólya urn process (GMVPP) turns out to be a mixture of Markov processes and admits representation (4)–(5), conditional on the parameter V . When the reinforcement measure R_x is concentrated on x , we call $(\mu_n)_{n \geq 0}$ a generalized randomly reinforced Pólya process (GRRPP). We give a characterization of GRRPPs with exchangeable weights $(W_n)_{n \geq 1}$ and show that the process $((X_n, W_n))_{n \geq 1}$ is partially conditionally identically distributed (partially c.i.d) [24], that is, conditionally on the past observations and the concurrent observation from the other sequence, the future observations are marginally identically distributed. We also extend some of the results for RRPPs to the generalized setting.

The paper is structured as follows. In Section 2.1, we recall the definition of a measure-valued Pólya urn process and prove representation (4)–(5) for a suitably selected sequence $(X_n)_{n \geq 1}$. Section 2.2 defines a particular subclass of MVPPs, called randomly reinforced Pólya processes (RRPP), which share with exchangeable Pólya sequences the property of reinforcing only the observed color. Section 3 is devoted to the study of the asymptotic properties of RRPPs. In Section 4, we give the definition of GMVPPs and GRRPPs, and obtain basic results.

2. Definitions and a Representation Result

Let (\mathbb{X}, d) be a complete separable metric space, endowed with its Borel σ -field \mathcal{X} . Denote by

$$\mathbb{M}_F(\mathbb{X}), \quad \mathbb{M}_F^*(\mathbb{X}), \quad \mathbb{M}_P(\mathbb{X}),$$

the collections of measures μ on \mathbb{X} that are finite, finite and non-null, and probability measures, respectively. We regard $\mathbb{M}_F(\mathbb{X})$, $\mathbb{M}_F^*(\mathbb{X})$ and $\mathbb{M}_P(\mathbb{X})$ as measurable spaces equipped with the σ -fields generated by $\mu \mapsto \mu(B)$, $B \in \mathcal{X}$. We further let

$$\mathbb{K}_F(\mathbb{X}, \mathbb{Y}), \quad \mathbb{K}_P(\mathbb{X}, \mathbb{Y}),$$

be the collections of transition kernels K from \mathbb{X} to \mathbb{Y} that are finite and probability kernels, respectively. Any non-null measure $\mu \in \mathbb{M}_F^*(\mathbb{X})$ has a normalized version $\mu' = \mu/\mu(\mathbb{X})$. If $f : \mathbb{X} \rightarrow \mathbb{Y}$ is measurable, then $f^\# : \mathbb{M}_F(\mathbb{X}) \rightarrow \mathbb{M}_F(\mathbb{Y})$ denotes the induced mapping of measures, $f^\#(\mu)(\cdot) = \mu(f^{-1}(\cdot))$, $\mu \in \mathbb{M}_F(\mathbb{X})$.

All random quantities are defined on a common probability space $(\Omega, \mathcal{H}, \mathbb{P})$, which is assumed to be rich enough to support any required randomization. The symbol ‘ \perp ’ will be used to denote independence between random objects, and ‘ $\stackrel{d}{=}$ ’ equality in distribution.

2.1. Measure-Valued Pólya urn Processes

Let $\mu \in \mathbb{M}_F^*(\mathbb{X})$ describe the contents of an urn, as in Section 1. Once a ball is picked at random from μ , the urn is reinforced according to a replacement rule, which is formally a kernel $R \in \mathbb{K}_F(\mathbb{X}, \mathbb{X})$ that maps colors $x \mapsto R_x(\cdot)$ to finite measures; thus,

$$\mu + R_x, \tag{9}$$

represents the updated urn composition if a ball of color x has been observed. In general, R is random and there exists a probability kernel $\mathcal{R} \in \mathbb{K}_P(\mathbb{X}, \mathbb{M}_F(\mathbb{X}))$ such that $R_x \sim \mathcal{R}_x$, $x \in \mathbb{X}$. Then, the distribution of (9) prior to the sampling of the urn is given by

$$\hat{R}_\mu(\cdot) = \int_{\mathbb{X}} \psi_\mu^\#(\mathcal{R}_x)(\cdot) \mu'(dx), \tag{10}$$

where ψ_μ is the measurable map $\nu \mapsto \nu + \mu$ from $\mathbb{M}_F(\mathbb{X})$ to $\mathbb{M}_F^*(\mathbb{X})$. By Lemma 3.3 in [9], $\mu \mapsto \hat{R}_\mu$ is a measurable map from $\mathbb{M}_F^*(\mathbb{X})$ to $\mathbb{M}_P(\mathbb{M}_F^*(\mathbb{X}))$.

Definition 1 (Measure-Valued Pólya Urn Process [9]). *A sequence $(\mu_n)_{n \geq 0}$ of random finite measures on \mathbb{X} is called a measure-valued Pólya urn process (MVPP) with parameters $\mu_0 \in \mathbb{M}_F^*(\mathbb{X})$*

and $\mathcal{R} \in \mathbb{K}_P(\mathbb{X}, \mathbb{M}_F(\mathbb{X}))$ if it is a Markov process with transition kernel \hat{R} given by (10). If, in particular, $\mathcal{R}_x = \delta_{R_x}$ for some $R \in \mathbb{K}_F(\mathbb{X}, \mathbb{X})$, then $(\mu_n)_{n \geq 0}$ is said to be a deterministic MVPP.

The representation theorem below formalizes the idea of MVPP as an urn scheme.

Theorem 1. A sequence $(\mu_n)_{n \geq 0}$ of random finite measures is an MVPP with parameters (μ_0, \mathcal{R}) if and only if, for every $n \geq 1$,

$$\mu_n = \mu_{n-1} + R_{X_n} \quad \text{a.s.}, \tag{11}$$

where $(X_n)_{n \geq 1}$ is a sequence of \mathbb{X} -valued random variables such that $X_1 \sim \mu'_0$ and, for $n \geq 2$,

$$\mathbb{P}(X_n \in \cdot \mid X_1, \mu_1, \dots, X_{n-1}, \mu_{n-1}) = \mu'_{n-1}(\cdot), \tag{12}$$

and R is a random finite transition kernel on \mathbb{X} such that

$$\mathbb{P}(R_{X_n} \in \cdot \mid X_1, \mu_1, \dots, X_{n-1}, \mu_{n-1}, X_n) = \mathcal{R}_{X_n}(\cdot). \tag{13}$$

Proof. If $(\mu_n)_{n \geq 0}$ satisfies (11)–(13) for every $n \geq 1$, then it holds a.s. that

$$\mathbb{P}(\mu_n \in \cdot \mid \mu_1, \dots, \mu_{n-1}) = \mathbb{E}[\psi_{\mu_{n-1}}^\#(\mathcal{R}_{X_n})(\cdot) \mid \mu_1, \dots, \mu_{n-1}] = \hat{R}_{\mu_{n-1}}(\cdot).$$

Conversely, suppose $(\mu_n)_{n \geq 0}$ is a MVPP with parameters (μ_0, \mathcal{R}) . As \mathcal{R} is a probability kernel from \mathbb{X} to $\mathbb{M}_F(\mathbb{X})$ and $\mathbb{M}_F(\mathbb{X})$ is Polish, then there exists by Lemma 4.22 in [25] a measurable function $f(x, u)$ such that, for every $x \in \mathbb{X}$,

$$f(x, U) \sim \mathcal{R}_x,$$

whenever U is a uniform random variable on $[0, 1]$, denoted $U \sim \text{Unif}[0, 1]$.

Let us prove by induction that there exists a sequence $((X_n, U_n))_{n \geq 1}$ such that $X_1 \sim \mu'_0$, $U_1 \perp X_1$, $U_1 \sim \text{Unif}[0, 1]$, $\mu_1 = \mu_0 + f(X_1, U_1)$ a.s., $(\mu_2, \mu_3, \dots) \perp (X_1, U_1) \mid \mu_1$, and, for every $n \geq 2$,

- (i) $\mathbb{P}(X_n \in \cdot \mid X_1, U_1, \mu_1, \dots, X_{n-1}, U_{n-1}, \mu_{n-1}) = \mu'_{n-1}(\cdot)$;
- (ii) $U_n \sim \text{Unif}[0, 1]$ and $U_n \perp (X_1, U_1, \mu_1, \dots, X_{n-1}, U_{n-1}, \mu_{n-1}, X_n)$;
- (iii) $\mu_n = \mu_{n-1} + f(X_n, U_n)$ a.s.;
- (iv) $(\mu_{n+1}, \mu_{n+2}, \dots) \perp (X_n, U_n) \mid (X_1, U_1, \mu_1, \dots, X_{n-1}, U_{n-1}, \mu_{n-1}, \mu_n)$;
- (v) $\mu_{n+1} \perp (X_1, U_1, \dots, X_n, U_n) \mid (\mu_1, \dots, \mu_n)$.

Then, Equations (11)–(13) follow from (i)–(iii) with $R_{X_n} = f(X_n, U_n)$.

Regarding the base case, let \tilde{X}_1 and \tilde{U}_1 be independent random variables such that $\tilde{U}_1 \sim \text{Unif}[0, 1]$ and $\tilde{X}_1 \sim \mu'_0$. It follows that, for any measurable set $B \subseteq \mathbb{M}_F(\mathbb{X})$,

$$\mathbb{P}(\mu_1 \in B) = \hat{R}_{\mu_0}(B) = \mathbb{E}[\psi_{\mu_0}^\#(\mathcal{R}_{\tilde{X}_1})(B)] = \mathbb{P}((\mu_0 + f(\tilde{X}_1, \tilde{U}_1)) \in B);$$

thus, $\mu_1 \stackrel{d}{=} \mu_0 + f(\tilde{X}_1, \tilde{U}_1)$. By Theorem 8.17 in [25], there exist random variables X_1 and U_1 such that

$$(\mu_1, X_1, U_1) \stackrel{d}{=} (\mu_0 + f(\tilde{X}_1, \tilde{U}_1), \tilde{X}_1, \tilde{U}_1),$$

and $(\mu_2, \mu_3, \dots) \perp (X_1, U_1) \mid \mu_1$. Then, in particular, $(X_1, U_1) \stackrel{d}{=} (\tilde{X}_1, \tilde{U}_1)$ and $(\mu_1, \mu_0 + f(X_1, U_1)) \stackrel{d}{=} (\mu_0 + f(\tilde{X}_1, \tilde{U}_1), \mu_0 + f(\tilde{X}_1, \tilde{U}_1))$, so

$$\mu_1 = \mu_0 + f(X_1, U_1) \quad \text{a.s.}$$

Regarding the induction step, assume that (i)–(v) hold true until some $n > 1$. Let \tilde{X}_{n+1} and \tilde{U}_{n+1} be such that $\tilde{U}_{n+1} \sim \text{Unif}[0, 1]$, $\tilde{U}_{n+1} \perp (X_1, U_1, \mu_1, \dots, X_n, U_n, \mu_n, \tilde{X}_{n+1})$, and

$$\mathbb{P}(\tilde{X}_{n+1} \in \cdot \mid X_1, U_1, \mu_1, \dots, X_n, U_n, \mu_n) = \mu'_n(\cdot).$$

It follows from (v) that, for any measurable set $B \subseteq \mathbb{M}_F(\mathbb{X})$,

$$\begin{aligned} \mathbb{P}(\mu_{n+1} \in B | X_1, U_1, \mu_1, \dots, X_n, U_n, \mu_n) &= \mathbb{E}[\psi_{\mu_n}^\#(\mathcal{R}_{\tilde{X}_{n+1}})(B) | X_1, U_1, \mu_1, \dots, X_n, U_n, \mu_n] \\ &= \mathbb{P}((\mu_n + f(\tilde{X}_{n+1}, \tilde{U}_{n+1})) \in B | X_1, U_1, \mu_1, \dots, X_n, U_n, \mu_n); \end{aligned}$$

thus, $\mu_{n+1} \stackrel{d}{=} \mu_n + f(\tilde{X}_{n+1}, \tilde{U}_{n+1}) \mid X_1, U_1, \mu_1, \dots, X_n, U_n, \mu_n$. By Theorem 8.17 in [25], there exist random variables X_{n+1} and U_{n+1} such that

$$\begin{aligned} (\mu_{n+1}, X_1, U_1, \mu_1, \dots, X_n, U_n, \mu_n, X_{n+1}, U_{n+1}) \\ \stackrel{d}{=} (\mu_n + f(\tilde{X}_{n+1}, \tilde{U}_{n+1}), X_1, U_1, \mu_1, \dots, X_n, U_n, \mu_n, \tilde{X}_{n+1}, \tilde{U}_{n+1}), \end{aligned}$$

and $(\mu_{n+2}, \mu_{n+3}, \dots) \perp (X_{n+1}, U_{n+1}) \mid (X_1, U_1, \mu_1, \dots, X_n, U_n, \mu_n, \mu_{n+1})$. Then, in particular, $U_{n+1} \sim \text{Unif}[0, 1]$, $U_{n+1} \perp (X_1, U_1, \mu_1, \dots, X_n, U_n, \mu_n, X_{n+1})$, and

$$\mathbb{P}(X_{n+1} \in \cdot \mid X_1, U_1, \mu_1, \dots, X_n, U_n, \mu_n) = \mu'_n(\cdot).$$

Moreover,

$$(\mu_{n+1}, \mu_n + f(X_{n+1}, U_{n+1})) \stackrel{d}{=} (\mu_n + f(\tilde{X}_{n+1}, \tilde{U}_{n+1}), \mu_n + f(\tilde{X}_{n+1}, \tilde{U}_{n+1}));$$

therefore,

$$\mathbb{P}(\mu_{n+1} = \mu_n + f(X_{n+1}, U_{n+1})) = \mathbb{P}(\mu_n + f(\tilde{X}_{n+1}, \tilde{U}_{n+1}) = \mu_n + f(\tilde{X}_{n+1}, \tilde{U}_{n+1})) = 1.$$

By Theorem 8.12 in [25], statement (v) with $n + 1$ is equivalent to $\mu_{n+2} \perp (X_1, U_1) \mid (\mu_1, \dots, \mu_{n+1})$ and $\mu_{n+2} \perp (X_{k+1}, U_{k+1}) \mid (X_1, U_1, \dots, X_k, U_k, \mu_1, \dots, \mu_{n+1})$, $k = 1, \dots, n$. The latter follows from the induction hypothesis since, by (iv), we have $(\mu_{k+2}, \dots, \mu_{n+2}) \perp (X_{k+1}, U_{k+1}) \mid (X_1, U_1, \dots, X_k, U_k, \mu_1, \dots, \mu_{k+1})$ for every $k = 1, \dots, n$. \square

The process $(X_n)_{n \geq 1}$ in Theorem 1 corresponds to the sequence of observed colors from the implied urn sampling scheme. Furthermore, the replacement rule takes the form $R_{X_n} = f(X_n, U_n)$, where f is some measurable function, $U_n \sim \text{Unif}[0, 1]$, and $U_n \perp (X_1, U_1, \dots, X_{n-1}, U_{n-1}, X_n)$, from which it follows that

$$\mu_n = \mu_{n-1} + f(X_n, U_n), \tag{14}$$

and

$$\mathbb{P}(X_{n+1} \in \cdot \mid X_1, \dots, X_n, (U_m)_{m \geq 1}) = \frac{\mu_0(\cdot) + \sum_{i=1}^n f(X_i, U_i)(\cdot)}{\mu_0(\mathbb{X}) + \sum_{i=1}^n f(X_i, U_i)(\mathbb{X})}. \tag{15}$$

Thus, the sequence $(U_n)_{n \geq 1}$ models the additional randomness in the reinforcement measure R . Janson [9] obtains a rather similar result; Theorem 1.3 in [9] states that any MVPP $(\mu_n)_{n \geq 0}$ can be coupled with a deterministic MVPP $(\bar{\mu}_n)_{n \geq 0}$ on $\mathbb{X} \times [0, 1]$ in the sense that

$$\bar{\mu}_n = \mu_n \times \lambda, \tag{16}$$

where λ is the Lebesgue measure on $[0, 1]$, and $\mu_n \times \lambda$ is the product measure on $\mathbb{X} \times [0, 1]$. In our case, the MVPP defined by $\bar{\mu}_0 = \mu_0 \times \lambda$ and, for $n \geq 1$,

$$\bar{\mu}_n = \bar{\mu}_{n-1} + f(X_n, U_n) \times \lambda,$$

has a non-random replacement rule $R_{x,\mu} = f(x, u) \times \lambda$ and satisfies (16) on a set of probability one.

2.2. Randomly Reinforced Pólya Processes

It follows from (8) that any Pólya sequence generates a deterministic MVPP through

$$\mu_n = \mu_{n-1} + \delta_{X_n}.$$

Here, we consider a randomly reinforced extension of Pólya sequences in the form of an MVPP with replacement rule $R_x = W(x) \cdot \delta_x$, $x \in \mathbb{X}$, where $W(x)$ is a non-negative random variable.

Definition 2 (Randomly Reinforced Pólya Process). *We call an MVPP with parameters (μ_0, \mathcal{R}) a randomly reinforced Pólya process (RRPP) if there exists $\eta \in \mathbb{K}_P(\mathbb{X}, \mathbb{R}_+)$ such that $\mathcal{R}_x = \xi_x^\#(\eta_x)$, $x \in \mathbb{X}$, where $\xi_x : \mathbb{R}_+ \rightarrow \mathbb{M}_F(\mathbb{X})$ is the map $w \mapsto w\delta_x$.*

Observe that, for RRPPs, the reinforcement measure $f(x, u)$ in (14)–(15) concentrates its mass on x ; thus, we obtain the following variant of the representation result in Theorem 1.

Proposition 1. *Let $(\mu_n)_{n \geq 0}$ be an RRPP with parameters (μ_0, η) . Then, there exist a measurable function $h : \mathbb{X} \times [0, 1] \rightarrow \mathbb{R}_+$ and a sequence $((X_n, U_n))_{n \geq 1}$ such that, using $W_n = h(X_n, U_n)$, we have for every $n \geq 1$ that*

$$\mu_n = \mu_{n-1} + W_n \delta_{X_n} \quad \text{a.s.}, \tag{17}$$

where $X_1 \sim \mu'_0$ and, for $n \geq 1$, $U_n \sim \text{Unif}[0, 1]$, $U_n \perp (X_1, U_1, \dots, X_{n-1}, U_{n-1}, X_n)$, and

$$\mathbb{P}(X_{n+1} \in \cdot \mid X_1, W_1, \dots, X_n, W_n) = \sum_{i=1}^n \frac{W_i}{\mu_0(\mathbb{X}) + \sum_{j=1}^n W_j} \delta_{X_i}(\cdot) + \frac{\mu_0(\mathbb{X})}{\mu_0(\mathbb{X}) + \sum_{j=1}^n W_j} \mu'_0(\cdot). \tag{18}$$

Moreover,

$$\mathbb{P}(W_n \in \cdot \mid X_1, W_1, \dots, X_{n-1}, W_{n-1}, X_n) = \eta_{X_n}(\cdot). \tag{19}$$

It follows from (19) that $W(x) \equiv h(x, U) \sim \eta_x$, $x \in \mathbb{X}$, whenever $U \sim \text{Unif}[0, 1]$. Then, the random measure

$$R_x = W(x) \cdot \delta_x \tag{20}$$

is such that $R_x \sim \mathcal{R}_x$, where \mathcal{R}_x appears in Definition 2.

3. Asymptotic Properties of RRPP

In this section, we study the asymptotic properties of RRPPs through the sequence $(X_n)_{n \geq 1}$ in the representation (17). We show that the limit behavior of $(\mu_n)_{n \geq 0}$ depends on the relationship between weights and observations. In particular, when $W(x) \equiv W$ in (20) is constant with respect to the color x , the process $(X_n)_{n \geq 1}$ is conditionally identically distributed (c.i.d.) and, for every $A \in \mathcal{X}$, the normalized sequence $(\mu'_n(A))_{n \geq 0}$ is a bounded martingale. We consider the c.i.d. case in Section 3.3. In contrast, if some colors x have a higher expected reinforcement, then they tend to dominate the observation process and, as n grows to infinity, the probability measure μ'_n concentrates its mass on the subset of dominant colors, see Theorem 2.

3.1. Preliminaries

Our focus is on the convergence of the normalized sequence $(\mu'_n)_{n \geq 0}$, which by Theorem 1 is a.s. equal to the predictive distributions (18). We also consider the sequence of empirical frequencies of $(X_n)_{n \geq 1}$, defined for $n \geq 1$ by

$$\hat{\mu}'_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}.$$

We obtain conditions under which the convergence in (2) extends to convergence in total variation, where the total variation distance between any two probability measures $\alpha, \beta \in \mathbb{M}_P(\mathbb{X})$ is given by

$$d_{TV}(\alpha, \beta) = \sup_{B \in \mathcal{X}} |\alpha(B) - \beta(B)|.$$

To state some of the results, we recall the definition of support of a probability measure $\gamma \in \mathbb{M}_P(\mathbb{R}_+)$,

$$\text{supp}(\gamma) = \{u \geq 0 : \gamma((u - \epsilon, u + \epsilon)) > 0, \forall \epsilon > 0\}.$$

Of particular interest is the conditional probability of observing a new color, given by

$$\theta_n \equiv \mathbb{P}(X_{n+1} \notin \{X_1, \dots, X_n\} \mid X_1, W_1, \dots, X_n, W_n) = \frac{\theta}{\theta + \sum_{j=1}^n W_j} \mu'_0(\{X_1, \dots, X_n\}^c),$$

for $n \geq 1$, where $\theta = \mu_0(\mathbb{X})$. This would inform us on the number of distinct values in a sample (X_1, \dots, X_n) of size n ,

$$L_n = \max\{k \in \{1, \dots, n\} : X_k \notin \{X_1, \dots, X_{k-1}\}\},$$

since $\theta_n = \mathbb{P}(L_{n+1} = L_n + 1 \mid X_1, W_1, \dots, X_n, W_n)$.

The following modes of convergence are used when we investigate the rate of convergence of the distance between μ'_n and $\hat{\mu}_n$.

Almost sure (a.s.) conditional convergence. Let $\mathcal{G} = (\mathcal{G}_n)_{n \geq 0}$ be a filtration and $\tilde{Q} \in \mathbb{K}_P(\Omega, \mathbb{X})$. A sequence $(Y_n)_{n \geq 1}$ is said to converge to \tilde{Q} in the sense of a.s. conditional convergence w.r.t. \mathcal{G} if the conditional distribution of Y_n , given \mathcal{G}_n , converges weakly on a set of probability one to \tilde{Q} , that is, as $n \rightarrow \infty$,

$$\mathbb{P}(Y_n \in \cdot \mid \mathcal{G}_n) \xrightarrow{w} \tilde{Q}(\cdot) \quad \text{a.s.}$$

We refer to [22] for more details.

Stable convergence. Stable convergence is a strong form of convergence in distribution, albeit weaker than a.s. conditional convergence. A sequence $(Y_n)_{n \geq 1}$ is said to converge stably to \tilde{Q} if

$$\mathbb{E}[V f(Y_n)] \rightarrow \mathbb{E}\left[V \int_{\mathbb{X}} f(x) \tilde{Q}(dx)\right],$$

for all continuous bounded functions f and any integrable random variable V . The main application of stable convergence is in central limit theorems that allow for mixing variables in the limit. See [26] for a complete reference on stable convergence.

In the sequel, the stable and a.s. conditional limits will be some Gaussian law, which we denote by $\mathcal{N}(\mu, \sigma^2)$ for parameters (μ, σ^2) , where $\mathcal{N}(\mu, 0) = \delta_\mu$.

3.2. RRPP with Dominant Colors

Using (20), let us define, for $x \in \mathbb{X}$,

$$w(x) = \mathbb{E}[W(x)] \quad \text{and} \quad \bar{w} = \sup_{x \in \mathbb{X}} w(x).$$

We further let

$$\mathcal{D} = \{x \in \mathbb{X} : w(x) = \bar{w}\},$$

be the set of dominant colors. The model (18) with $\mathcal{D} \subset \mathbb{X}$ has been studied by [23] under the assumption that \bar{w} is strictly greater than the next largest value of $w(\cdot)$ in the support of $w^\sharp(\mu'_0)$. Then, the probability of observing a non-dominant color, $x \in \mathcal{D}^c$, vanishes, and the predictive and the empirical distributions converge in total variation to a common random

probability measure, which is concentrated on \mathcal{D} . For completeness reasons, we report here the main results from [23].

Theorem 2 ([23], Theorem 3.3). *For any RRPP $(\mu_n)_{n \geq 0}$ that satisfies*

$$\begin{aligned} W(x) &\leq \beta < \infty; \\ \bar{w} &\in \text{supp}(w^\#(\mu'_0)); \\ \bar{w} &> \bar{w}^c \equiv \sup\{u \geq 0 : u \in \text{supp}(w^\#(\mu'_0(\cdot|\mathcal{D}^c)))\}, \end{aligned} \tag{21}$$

there exists a random probability measure \tilde{P} on \mathbb{X} with $\tilde{P}(\mathcal{D}) = 1$ a.s. such that

$$d_{TV}(\mu'_n, \tilde{P}) \xrightarrow{a.s.} 0 \quad \text{and} \quad d_{TV}(\hat{\mu}'_n, \tilde{P}) \xrightarrow{a.s.} 0.$$

Under conditions (21), Theorem 3.3 in [23] implies $\sum_{i=1}^n W_i/n \xrightarrow{a.s.} \bar{w}$. If μ_0 is further diffuse, then $\sum_{n=1}^\infty \theta_n = \infty$ a.s., and so $L_n \xrightarrow{a.s.} \infty$ by Theorem 1 in [27]; thus, by Theorem 1 in [27], Proposition 3.4 in [23] shows that the actual growth rate is that of a Pólya sequence,

$$\frac{L_n}{\log n} \xrightarrow{a.s.} \frac{\theta}{\bar{w}}. \tag{22}$$

In addition to the uniform convergence in Theorem 2, the authors in [23] obtain set-wise rates of convergence. To state their result, we introduce, for any $A \in \mathcal{X}$,

$$q_A = \lim_{n \rightarrow \infty} \mathbb{E}[W_{n+1}^2 \delta_{X_{n+1}}(A) | X_1, W_1, \dots, X_n, W_n],$$

which exists a.s. under the assumptions of Theorem 2.

Theorem 3 ([23], Theorem 4.2). *Let $(\mu_n)_{n \geq 0}$ be an RRPP satisfying (21). Suppose $\bar{w} > 2\bar{w}^c$. Define*

$$V(A) = \frac{1}{\bar{w}^2} \{(\tilde{P}(A^c))^2 q_A + (\tilde{P}(A))^2 q_{A^c}\} \quad \text{and} \quad U(A) = V(A) - \tilde{P}(A)\tilde{P}(A^c).$$

Then,

$$\sqrt{n}(\mu'_n(A) - \hat{\mu}'_n(A)) \xrightarrow{\text{stably}} \mathcal{N}(0, U(A)),$$

and

$$\sqrt{n}(\mu'_n(A) - \tilde{P}(A)) \xrightarrow{a.s.\text{cond.}} \mathcal{N}(0, V(A)) \quad \text{w.r.t. } (\mathcal{F}_n^{X,W})_{n \geq 1},$$

where $\mathcal{F}_n^{X,W} = \sigma(X_1, W_1, \dots, X_n, W_n)$, $n \geq 1$ is the filtration generated by $((X_n, W_n))_{n \geq 1}$.

3.3. RRPP with Independent Weights

Let $(\mu_n)_{n \geq 0}$ be an RRPP with reinforcement distribution $\eta_x \equiv \eta$ that does not depend on x . Using the notation of Section 3.2, we have

$$w(x) \equiv \bar{w}, \tag{23}$$

and, thus, $\mathcal{D} = \mathbb{X}$. An equivalent formulation can be given in terms of the sequence of weights $(W_n)_{n \geq 1}$ in Proposition 1, whereby

$$W_n = h(U_n), \tag{24}$$

for some measurable function h , with $U_n \perp (X_1, U_1, \dots, X_{n-1}, U_{n-1}, X_n)$ and $U_n \sim \text{Unif}[0, 1]$. Then, $W_n \stackrel{i.i.d.}{\sim} \eta$ and $W_n \perp (X_1, \dots, X_n)$, which implies that $\mathbb{E}[W_1] = \bar{w}$.

The model (18) with weights (24) has been studied by [18,19,22], among others, where the authors obtain central limit theorems and study the growth rate of L_n when $\bar{w} < \infty$.

Their results rely on the fact that $(X_n)_{n \geq 1}$ is conditionally identically distributed (c.i.d.) with respect to the filtration generated by $((X_n, W_n))_{n \geq 1}$. By [21], an \mathbb{X} -valued random sequence $(Y_n)_{n \geq 1}$ that is adapted to a filtration $(\mathcal{F}_n)_{n \geq 1}$ is said to be c.i.d. with respect to $(\mathcal{F}_n)_{n \geq 1}$ if and only if $(Y_n)_{n \geq 1}$ is identically distributed and, for every $n, k \geq 1$,

$$\mathbb{P}(Y_{n+k} \in \cdot \mid \mathcal{F}_n) = \mathbb{P}(Y_{n+1} \in \cdot \mid \mathcal{F}_n). \tag{25}$$

Proposition 2 ([19], Lemma 6). *For any RRPP $(\mu_n)_{n \geq 0}$ with $\eta_x \equiv \eta$, the observation process $(X_n)_{n \geq 1}$ is c.i.d. with respect to the filtration generated by $((X_n, W_n))_{n \geq 1}$.*

C.i.d. processes preserve many of the properties of exchangeable sequences, see [21]. For example, if $(Y_n)_{n \geq 1}$ is c.i.d., then there exists a random probability measure such that (2)–(3) hold true with respect to the filtration used in the definition (25). It follows for the model in Proposition 2 that there exists $\tilde{P} \in \mathbb{K}_p(\Omega, \mathbb{X})$ such that, for every $A \in \mathcal{X}$,

$$\mu'_n(A) \xrightarrow{a.s.} \tilde{P}(A).$$

In fact, by (25), the sequence $(\mu'_n(A))_{n \geq 0}$ is a bounded martingale. On the other hand, (23) implies that $\mathcal{D} = \mathbb{X}$; therefore, any RRPP with $\eta_x \equiv \eta$ whose weights are bounded, $W_1 \leq \beta < \infty$, satisfies the assumptions of Theorem 2. In that case,

$$d_{TV}(\mu'_n, \tilde{P}) \xrightarrow{a.s.} \tilde{P}.$$

It follows from Theorem 4.2 in [23] that the boundedness condition in (21) is needed to show that (i) $\sum_{i=1}^n W_i/n \xrightarrow{a.s.} \bar{w}$; and (ii) μ'_n converge set-wise to \tilde{P} , which is non-trivial in that setting. Here, (i) is granted as $(W_n)_{n \geq 1}$ is i.i.d., and (ii) has already been established; thus, we obtain the following result for RRPPs with independent weights.

Theorem 4. *For any RRPP $(\mu_n)_{n \geq 0}$ with $\eta_x \equiv \eta$, there exists a random probability measure \tilde{P} on \mathbb{X} such that*

$$d_{TV}(\mu'_n, \tilde{P}) \xrightarrow{a.s.} 0 \quad \text{and} \quad d_{TV}(\hat{\mu}'_n, \tilde{P}) \xrightarrow{a.s.} 0.$$

Proof. Let $((X_n, W_n))_{n \geq 1}$ be the joint observation process associated to $(\mu_n)_{n \geq 0}$ by Proposition 1. As $\eta_x \equiv \eta$, Equation (19) implies that $W_n \stackrel{i.i.d.}{\sim} \eta$; thus, by the strong law of large numbers,

$$\frac{1}{n} \sum_{i=1}^n W_i \xrightarrow{a.s.} \bar{w} \leq \infty. \tag{26}$$

Let us define, for $n \geq 1$,

$$P_n(\cdot) = \mathbb{P}(X_{n+1} \in \cdot \mid \mathcal{F}_n^{X,W}), \quad \text{where } \mathcal{F}_n^{X,W} = \sigma(X_1, W_1, \dots, X_n, W_n).$$

By Proposition 2, $(X_n)_{n \geq 1}$ is c.i.d. with respect to $(\mathcal{F}_n^{X,W})_{n \geq 1}$, so there exists by Lemmas 2.1 and 2.4 in [21] a random probability measure \tilde{P} on \mathbb{X} such that, for every $A \in \mathcal{X}$,

$$P_n(A) \xrightarrow{a.s.} \tilde{P}(A). \tag{27}$$

Moreover, $\int_{\mathbb{X}} f(x) P_n(dx) = \mathbb{E}[\int_{\mathbb{X}} f(x) \tilde{P}(dx) \mid \mathcal{F}_n^{X,W}]$ a.s. for every bounded measurable $f : \mathbb{X} \rightarrow \mathbb{R}$. Fix $m \geq 1$. By a monotone class argument, we can show that, for every bounded measurable $f : \mathbb{X}^2 \rightarrow \mathbb{R}$,

$$\int_{\mathbb{X}} f(X_m, x) P_n(dx) = \mathbb{E}[\int_{\mathbb{X}} f(X_m, x) \tilde{P}(dx) \mid \mathcal{F}_n^{X,W}] \quad \text{a.s., for all } n > m;$$

thus, $P_n(\{X_m\}) = \mathbb{E}[\tilde{P}(\{X_m\}) | \mathcal{F}_n^{X,W}]$ a.s., and so $(P_n(\{X_m\}))_{n>m}$ is a uniformly integrable martingale. It follows from martingale convergence that, as $n \rightarrow \infty$,

$$P_n(\{X_m\}) \xrightarrow{a.s.} \tilde{P}(\{X_m\}). \tag{28}$$

Using (26)–(28), we can repeat the argument in the proof of Proposition 3.1 in [23] to show that (i) $d_{TV}(P_n, \tilde{P}) \xrightarrow{a.s.} 0$, and so $d_{TV}(\mu'_n, \tilde{P}) \xrightarrow{a.s.} 0$ by Proposition 1; and (ii) $d_{TV}(\hat{\mu}'_n, \tilde{P}) \xrightarrow{a.s.} 0$. \square

Equation (26) implies that $\theta_n \xrightarrow{a.s.} 0$. If, in addition, $\bar{w} < \infty$, then $\sum_{n=1}^\infty \theta_n = \infty$ a.s. and $L_n \xrightarrow{a.s.} \infty$. In fact, as long as $\bar{w} < \infty$, the sequence $(L_n)_{n \geq 1}$ grows at the same rate as (22).

Proposition 3 ([18], Lemma 6). *Let $\eta \in \mathbb{M}_P(\mathbb{X})$ and μ_0 be diffuse. If $\bar{w} < \infty$, then*

$$\frac{L_n}{\log n} \xrightarrow{a.s.} \frac{\theta}{\bar{w}}.$$

If $\bar{w} = \infty$, then θ_n may approach zero fast enough that we stop seeing new observations as $n \rightarrow \infty$. For example, let us consider random reinforcement with a totally skewed stable distribution $S_\alpha(1, \sigma, 0)$ for $\alpha \in (0, 2]$ and $\sigma > 0$. If $\alpha < 1$, then $\bar{w} = \infty$, and we show that $n^{1/\alpha}\theta_n$ is stochastically bounded, which implies that L_n converges to a finite limit.

Proposition 4. *Let η be a $S_\alpha(1, \sigma, 0)$ distribution with stability parameter $\alpha < 1$, and μ_0 be diffuse. Then, $\theta_n = O_p(n^{-1/\alpha})$ and*

$$\lim_{n \rightarrow \infty} L_n < \infty \quad a.s.$$

Proof. From the properties of stable distributions, we obtain $n^{-1/\alpha} \sum_{i=1}^n W_i \stackrel{d}{=} W_1$ for every $n \geq 1$ and, as a consequence,

$$\theta_n = n^{-1/\alpha} \frac{\theta}{n^{-1/\alpha}\theta + n^{-1/\alpha} \sum_{i=1}^n W_i} \stackrel{d}{=} n^{-1/\alpha} \frac{\theta}{n^{-1/\alpha}\theta + W_1} \leq n^{-1/\alpha} \frac{\theta}{W_1}.$$

By Theorem 5.4.1 in [28], $\mathbb{E}[1/W_1] < \infty$, and so $1/W_1 < \infty$ a.s. It follows for every $M > 0$ that $\mathbb{P}(n^{1/\alpha}\theta_n > M) \leq \mathbb{P}(\theta/W_1 > M)$, which can be made arbitrarily small by taking M large enough. Regarding the second assertion, as $1/\alpha > 1$, we have

$$\mathbb{E}[\lim_{n \rightarrow \infty} L_n] = \lim_{n \rightarrow \infty} \sum_{i=1}^n \mathbb{E}[\mathbb{1}_{\{L_i=L_{i-1}+1\}}] = \sum_{n=1}^\infty \mathbb{E}[\theta_n] \leq \sum_{n=1}^\infty \frac{\theta}{n^{1/\alpha}} \mathbb{E}[1/W_1] < \infty.$$

\square

Proposition 4 can be extended for any fat tailed reinforcement distribution η by means of a generalized central limit theorem (see, e.g., [28] (p. 62)).

The rate of convergence of (18) and $\hat{\mu}'_n$ has already been studied for the model with independent weights under different assumptions, see, e.g., [19] (p. 1363), Examples 4.2 and 4.5 in the technical report to [18], Corollary 4.1 in [22] for $\mathbb{X} = \{0, 1\}$. In the next theorem, we combine ideas from [18,20] to give a fairly general result.

Theorem 5. *Let $\eta \in \mathbb{M}_P(\mathbb{R}_+)$. If $\mathbb{E}[W_1^2] < \infty$, then*

$$\sqrt{n}(\mu'_n(A) - \hat{\mu}_n(A)) \xrightarrow{stably} \mathcal{N}(0, U(A)), \quad \text{where } U(A) = \frac{\text{Var}(W_1)}{\mathbb{E}[W_1^2]} \tilde{P}(A)\tilde{P}(A^c). \tag{29}$$

If, in addition, $\mathbb{E}[W_1^4] < \infty$, then, with respect to the filtration generated by $((X_n, W_n))_{n \geq 1}$,

$$\sqrt{n}(\mu'_n(A) - \tilde{P}(A)) \xrightarrow{a.s. cond.} \mathcal{N}(0, V(A)), \quad \text{where } V(A) = \frac{\mathbb{E}[W_1^2]}{\bar{w}^2} \tilde{P}(A) \tilde{P}(A^c). \quad (30)$$

Proof. Let us define, for $n \geq 1$,

$$P_n(\cdot) = \mathbb{P}(X_{n+1} \in \cdot \mid X_1, W_1, \dots, X_n, W_n).$$

The assertions in Theorem 5 have already been established by [18] when $W_1 \geq \gamma > 0$. In that case, Examples 4.2 and 4.5 in the technical report to [18] show that (29) is a consequence of the fact that

$$\mathbb{E} \left[\max_{1 \leq k \leq n} |Y_{n,k}| \right] \longrightarrow 0 \quad \text{and} \quad \sum_{k=1}^n Y_{n,k}^2 \xrightarrow{P} U(A), \quad (31)$$

where $Y_{n,k} = \frac{1}{\sqrt{n}} \{ \delta_{X_k}(A) - kP_k(A) + (k-1)P_{k-1}(A) \}$, and (30) follows from

$$\mathbb{E} \left[\sup_{n \geq 1} \sqrt{n} |P_{n-1}(A) - P_n(A)| \right] < \infty \quad \text{and} \quad n \sum_{k \geq n} (P_{k-1}(A) - P_k(A))^2 \xrightarrow{a.s.} V(A).$$

Replicating the approach of Proposition 9 in [20], we avoid using the assumption $W_1 \geq \gamma > 0$ by conditioning on the sets $H_n = \{ 2 \sum_{i=1}^n W_i \geq n\bar{w} \}$, $n \geq 1$. By (26), $\mathbb{1}_{H_n} \xrightarrow{a.s.} 1$, so (29) follows from (31) with

$$Y_{n,k} = \frac{1}{\sqrt{n}} \mathbb{1}_{H_{k-1}} \{ \delta_{X_k}(A) - kP_k(A) + (k-1)P_{k-1}(A) \},$$

whereas (30) is, ultimately, a result of

$$\mathbb{E} \left[\sup_{n \geq 1} \sqrt{n} \cdot \mathbb{1}_{H_n} |P_{n-1}(A) - P_n(A)| \right] < \infty \quad \text{and} \quad n \sum_{k \geq n} (P_{k-1}(A) - P_k(A))^2 \xrightarrow{a.s.} V(A).$$

□

4. Generalized Measure-Valued Pólya Urn Processes

The definition of an MVPP assumes that the law of the reinforcement \mathcal{R} is fixed, yet, in some situations, \mathcal{R} can itself be random (e.g., RRPP with exchangeable weights, see Section 4.1). To avoid measurability issues, we assume a parametric model for \mathcal{R} , with the parameter taking values in a Polish space \mathbb{V} .

Definition 3 (Generalized Measure-Valued Pólya Urn Process). *Let V be a \mathbb{V} -valued random variable. A sequence $(\mu_n)_{n \geq 0}$ of random finite measure on \mathbb{X} is called a generalized measure-valued Pólya urn process (GMVPP) with uncertainty parameter V , initial state $\mu_0 \in \mathbb{M}_F^*(\mathbb{X})$ and replacement rule $\mathcal{R} \in \mathbb{K}_P(\mathbb{V} \times \mathbb{X}, \mathbb{M}_F(\mathbb{X}))$ if $\mu_1 \mid V \sim \hat{R}_{\mu_0}^V$, and, for every $n \geq 2$,*

$$\mathbb{P}(\mu_n \in \cdot \mid V, \mu_1, \dots, \mu_{n-1}) = \hat{R}_{\mu_{n-1}}^V(\cdot),$$

where \hat{R} is the transition probability kernel from $\mathbb{V} \times \mathbb{M}_F^*(\mathbb{X})$ to $\mathbb{M}_F^*(\mathbb{X})$ given by

$$(v, \mu) \mapsto \hat{R}_\mu^v(\cdot) = \int_{\mathbb{X}} \psi_\mu^\#(\mathcal{R}(v, x))(\cdot) \mu'(dx),$$

and ψ_μ is the map $v \mapsto v + \mu$.

It follows from Definition 3 that any GMVPP is a mixture of Markov chains with initial state μ_0 and transition kernel \hat{R}^V . A separate modeling approach, which we do not examine here, defines a measure-valued Markov chain with transition kernel

$$\mu \mapsto \int_{\mathbb{X}} \psi_{\mu}^{\sharp}(\mathcal{R}(\mu, x))(\cdot) \mu'(dx).$$

In fact, some of the predictive constructions in [11,29] can be framed in such a way.

Theorem 1 extends to GMVPPs, provided that we condition all quantities on the parameter V . As a consequence, there exists a measurable function f from $\mathbb{V} \times \mathbb{X} \times [0, 1]$ to $\mathbb{M}_F(\mathbb{X})$ and a random sequence $((X_n, U_n))_{n \geq 1}$ such that

$$\mu_n = \mu_{n-1} + f(V, X_n, U_n) \quad \text{a.s.}, \tag{32}$$

where $U_n \sim \text{Unif}[0, 1]$, $U_n \perp (V, X_1, U_1, \dots, X_{n-1}, U_{n-1}, X_n)$, $X_1 \mid (V, (U_m)_{m \geq 1}) \sim \mu'_0$, and, for $n \geq 1$,

$$\mathbb{P}(X_{n+1} \in \cdot \mid V, X_1, \dots, X_n, (U_m)_{m \geq 1}) = \frac{\mu_0(\cdot) + \sum_{i=1}^n f(V, X_i, U_i)(\cdot)}{\mu_0(\mathbb{X}) + \sum_{i=1}^n f(V, X_i, U_i)(\mathbb{X})}, \tag{33}$$

and

$$\mathbb{P}(f(V, X_n, U_n) \in \cdot \mid V, X_1, U_1, \dots, X_{n-1}, U_{n-1}, X_n) = \mathcal{R}(V, X_n)(\cdot). \tag{34}$$

The definition of a randomly reinforced Pólya process is similarly generalized to cover the case of a random reinforcement distribution η .

Definition 4 (Generalized Randomly Reinforced Pólya Process). *We call a GMVPP with parameters (V, μ_0, \mathcal{R}) a generalized randomly reinforced Pólya process (GRRPP) if there exists $\eta \in \mathbb{K}_P(\mathbb{V} \times \mathbb{X}, \mathbb{R}_+)$ such that $\mathcal{R}(v, x) = \xi_x^{\sharp}(\eta(v, x))$, where $\xi_x : \mathbb{R}_+ \rightarrow \mathbb{M}_F(\mathbb{X})$ is the map $w \mapsto w\delta_x$.*

For GRRPPs, the function f in the representation (32)–(34) can be written as

$$f(v, x, u) = h(v, x, u) \cdot \delta_x,$$

where h is a measurable function from $\mathbb{V} \times \mathbb{X} \times [0, 1]$ to \mathbb{R}_+ such that $h(v, x, U) \sim \eta(v, x)$ for all $v \in \mathbb{V}$ and $x \in \mathbb{X}$, whenever $U \sim \text{Unif}[0, 1]$. Letting $W_n = h(V, X_n, U_n)$, we obtain

$$\mu_n = \mu_{n-1} + W_n \delta_{X_n} \quad \text{a.s.}, \tag{35}$$

where

$$\mathbb{P}(X_{n+1} \in \cdot \mid V, X_1, \dots, X_n, (U_m)_{m \geq 1}) = \frac{\mu_0(\cdot) + \sum_{i=1}^n W_i \delta_{X_i}(\cdot)}{\mu_0(\mathbb{X}) + \sum_{i=1}^n W_i}, \tag{36}$$

and

$$\mathbb{P}(W_n \in \cdot \mid V, X_1, U_1, \dots, X_{n-1}, U_{n-1}, X_n) = \eta(V, X_n)(\cdot). \tag{37}$$

The weights W_n in (36) allow us to incorporate additional information about the observations $(X_n)_{n \geq 1}$. As an example, consider the problem of computer-based classification, where the output usually includes confidence scores, which reflect the software’s confidence that the classifications are correct. In analyzing the number and dimension of the types already discovered, or the probability of detecting a new type, a typical procedure would take into account only those classifications whose confidence scores are above a certain threshold. Alternatively, we could adopt a Bayesian perspective and weigh each classification according to its confidence score. Denoting by $((X_n, W_n))_{n \geq 1}$ the sequence of classifications and confidence scores, we would model the distribution of the next classification by (36).

4.1. GRRPP with Exchangeable Weights

Let $(\mu_n)_{n \geq 0}$ be a GRRPP with reinforcement distribution $\eta(v)$ that does not depend on x . Then,

$$W_n = h(V, U_n),$$

for some measurable function $h(v, u)$. The next result shows that the sequence $(W_n)_{n \geq 1}$ is exchangeable with directing random measure $\tilde{\eta} \equiv \eta(V)$. Moreover, $(\mu_n)_{n \geq 0}$ is completely parameterized by $(\mu_0, \tilde{\eta})$.

Theorem 6. *A sequence $(\mu_n)_{n \geq 0}$ of random finite measures is a GRRPP with parameters $(\mu_0, \tilde{\eta})$ for $\tilde{\eta} \in \mathbb{K}_P(\Omega, \mathbb{R}_+)$ if and only if $\mu_0 = \theta v$ and, for every $n \geq 1$,*

$$\mu_n = \mu_{n-1} + W_n \delta_{X_n} \quad \text{a.s.},$$

where $\theta \in (0, \infty)$, $v \in \mathbb{M}_P(\mathbb{X})$, $(W_n)_{n \geq 1}$ is an exchangeable process with directing random measure $\tilde{\eta}$, and $(X_n)_{n \geq 1}$ is a sequence of \mathbb{X} -valued random variables such that $X_1 \mid (W_k)_{k \geq 1} \sim v$ and, for $n \geq 1$,

$$\mathbb{P}(X_{n+1} \in \cdot \mid X_1, \dots, X_n, (W_k)_{k \geq 1}) = \sum_{i=1}^n \frac{W_i}{\theta + \sum_{j=1}^n W_j} \delta_{X_i}(\cdot) + \frac{\theta}{\theta + \sum_{j=1}^n W_j} v(\cdot). \quad (38)$$

Proof. Let $(\mu_n)_{n \geq 0}$ be a GRRPP with parameters $(\mu_0, \tilde{\eta})$, and consider the representation (35)–(37). Put $\theta = \mu_0(\mathbb{X})$ and $v = \mu'_0$. It follows from (37) that

$$W_n \mid \tilde{\eta} \stackrel{i.i.d.}{\sim} \tilde{\eta};$$

thus, $(W_n)_{n \geq 1}$ is exchangeable. Moreover, $W_n = h(V, U_n)$, $n \geq 1$, so (38) follows from (36).

Conversely, suppose $\mu_n = \mu_{n-1} + W_n \delta_{X_n}$, where the process $((X_n, W_n))_{n \geq 1}$ is as described. It follows from (38) and Theorem 8.12 in [25] that

$$(W_k)_{k \geq 1} \perp X_1 \quad \text{and} \quad (W_{n+k})_{k \geq 1} \perp (X_1, \dots, X_{n+1}) \mid (W_1, \dots, W_n), \quad n \geq 1. \quad (39)$$

Since $(W_n)_{n \geq 1}$ is exchangeable with directing random measure $\tilde{\eta}$, we have

$$W_n \mid \tilde{\eta} \stackrel{i.i.d.}{\sim} \tilde{\eta}. \quad (40)$$

Furthermore, $\tilde{\eta}$ is measurable with respect to the tail σ -field of $(W_n)_{n \geq 1}$, so, by (39),

$$\tilde{\eta} \perp X_1 \quad \text{and} \quad \tilde{\eta} \perp (X_1, \dots, X_{n+1}) \mid (W_1, \dots, W_n), \quad n \geq 1. \quad (41)$$

Using (39)–(41), we can show that

$$W_1 \perp X_1 \mid \tilde{\eta} \quad \text{and} \quad W_{n+1} \perp (X_1, W_1, \dots, X_n, W_n, X_{n+1}) \mid \tilde{\eta}, \quad n \geq 1.$$

Then, $\mathbb{P}(\mu_1 \in \cdot \mid \tilde{\eta}) = \mathbb{P}(\mu_0 + W_1 \delta_{X_1} \in \cdot \mid \tilde{\eta}) = \int_{\mathbb{X}} \psi_{\mu_0}^\#(\xi_x^\#(\tilde{\eta}))(\cdot) \mu'_0(dx)$ and, for $n \geq 2$,

$$\begin{aligned} \mathbb{P}(\mu_n \in \cdot \mid \tilde{\eta}, \mu_1, \dots, \mu_{n-1}) &= \mathbb{E}[\mathbb{P}(\mu_{n-1} + W_n \delta_{X_n} \in \cdot \mid \tilde{\eta}, X_1, \dots, W_{n-1}, X_n) \mid \tilde{\eta}, \mu_1, \dots, \mu_{n-1}] \\ &= \mathbb{E}[\mathbb{E}[\psi_{\mu_{n-1}}^\#(\xi_{X_n}^\#(\tilde{\eta}))(\cdot) \mid X_1, \dots, X_{n-1}, (W_m)_{m \geq 1}] \mid \tilde{\eta}, \mu_1, \dots, \mu_{n-1}] \\ &= \int_{\mathbb{X}} \psi_{\mu_{n-1}}^\#(\xi_x^\#(\tilde{\eta}))(\cdot) \mu'_{n-1}(dx). \end{aligned}$$

□

It follows from the proof of Theorem 6 that $(X_1, W_1) \sim \mu'_0 \times \mathbb{E}[\tilde{\eta}]$ and, for $n \geq 1$,

$$\mathbb{P}((X_{n+1}, W_{n+1}) \in \cdot \mid X_1, W_1, \dots, X_n, W_n) = (\mu'_n \times \mathbb{E}[\tilde{\eta} \mid W_1, \dots, W_n])(\cdot). \tag{42}$$

As μ'_n and $\mathbb{E}[\tilde{\eta} \mid W_1, \dots, W_n]$ are both symmetric with respect to $((X_1, W_1), \dots, (X_n, W_n))$, then (42) is a symmetric function of $((X_1, W_1), \dots, (X_n, W_n))$. This is a necessary but not sufficient condition for $((X_n, W_n))_{n \geq 1}$ to be exchangeable, see Proposition 3.2 and Example 3.1 in [2]. In Proposition 5, we show that $((X_n, W_n))_{n \geq 1}$ is exchangeable if and only if either μ'_0 is degenerate or the weights are a.s. identical. On the other hand, for every $n, k \geq 1$, the sequence $((X_n, W_n))_{n \geq 1}$ satisfies

$$\mathbb{P}(W_k \in \cdot \mid X_1) = \mathbb{P}(W_1 \in \cdot \mid X_1), \quad \mathbb{P}(X_k \in \cdot \mid W_1) = \mathbb{P}(X_1 \in \cdot \mid W_1), \tag{43}$$

and

$$\begin{aligned} \mathbb{P}(W_{n+k} \in \cdot \mid X_1, W_1, \dots, X_n, W_n, X_{n+1}) &= \mathbb{P}(W_{n+1} \in \cdot \mid X_1, W_1, \dots, X_n, W_n, X_{n+1}), \\ \mathbb{P}(X_{n+k} \in \cdot \mid X_1, W_1, \dots, X_n, W_n, W_{n+1}) &= \mathbb{P}(X_{n+1} \in \cdot \mid X_1, W_1, \dots, X_n, W_n, W_{n+1}). \end{aligned} \tag{44}$$

By [24], Equations (43) and (44) are defining a process that is partially conditionally identity distributed (partially c.i.d.). Analogously to the c.i.d. case, partially c.i.d. processes preserve many of the properties of partially exchangeable sequences, see [24].

Proposition 5. *Under the conditions of Theorem 6, $((X_n, W_n))_{n \geq 1}$ is partially c.i.d. Moreover, $((X_n, W_n))_{n \geq 1}$ is exchangeable if and only if either μ'_0 is degenerate or $W_n = W_1$ a.s., $n \geq 1$. In that case, $((X_n, W_n))_{n \geq 1}$ is partially exchangeable.*

Proof. It follows that $((X_n, W_n))_{n \geq 1}$ is partially c.i.d. if and only if $X_2 \stackrel{d}{=} X_1 \mid W_1, W_2 \stackrel{d}{=} W_1 \mid X_1$, and (44) is true for every $n \geq 1$ with $k = 2$. By hypothesis, $(W_n)_{n \geq 1}$ is exchangeable and $(W_n)_{n \geq 1} \perp X_1$, so $W_2 \stackrel{d}{=} W_1 \mid X_1$. Moreover, applying (39) repeatedly, we obtain

$$\begin{aligned} \mathbb{P}(W_{n+2} \in \cdot \mid X_1, W_1, \dots, X_n, W_n, X_{n+1}) &= \mathbb{E}[\mathbb{P}(W_{n+2} \in \cdot \mid X_1, \dots, W_{n+1}, X_{n+2}) \mid X_1, W_1, \dots, X_n, W_n, X_{n+1}] \\ &= \mathbb{E}[\mathbb{P}(W_{n+2} \in \cdot \mid W_1, \dots, W_{n+1}) \mid W_1, \dots, W_n] \\ &= \mathbb{P}(W_{n+1} \in \cdot \mid W_1, \dots, W_n) = \mathbb{P}(W_{n+1} \in \cdot \mid X_1, W_1, \dots, X_n, W_n, X_{n+1}). \end{aligned}$$

On the other hand, by (38),

$$\begin{aligned} \mathbb{P}(X_{n+2} \in \cdot \mid X_1, W_1, \dots, X_n, W_n, W_{n+1}) &= \mathbb{E}[\mu'_{n+1}(\cdot) \mid X_1, W_1, \dots, X_n, W_n, W_{n+1}] \\ &= \frac{\mu_n(\cdot) + W_{n+1} \cdot \mu'_n(\cdot)}{\mu_{n+1}(\mathbb{X})} = \mu'_n(\cdot) \\ &= \mathbb{P}(X_{n+1} \in \cdot \mid X_1, W_1, \dots, X_n, W_n, W_{n+1}). \end{aligned}$$

Analogously, $\mathbb{P}(X_2 \in \cdot \mid W_1) = \mu_1(\cdot) = \mathbb{P}(X_1 \in \cdot \mid W_1)$, which completes the proof of the first part.

If μ'_0 is degenerate, then $((X_n, W_n))_{n \geq 1}$ is trivially exchangeable. If $W_n = W_1$ a.s. instead, then one can show that $((X_n, W_n))_{n \geq 1}$ satisfies condition (b) of Proposition 3.2 in [2], which, together with the symmetry of (42), implies by Theorem 3.1 in [2] that $((X_n, W_n))_{n \geq 1}$ is exchangeable.

Conversely, suppose that $((X_n, W_n))_{n \geq 1}$ is exchangeable. As $((X_n, W_n))_{n \geq 1}$ is partially c.i.d., the predictive distributions (42) converge to a product random measure [24]. It follows from de Finetti’s theorem that $((X_n, W_n))_{n \geq 1}$ is partially exchangeable, so, in particular,

$$(X_1, W_1, X_2, W_2) \stackrel{d}{=} (X_1, W_2, X_2, W_1).$$

However, $W_2 \perp X_2 \mid (X_1, W_1)$ from (36), so $W_1 \perp X_2 \mid (X_1, W_2)$. Thus, for every bounded measurable function \tilde{f} , there exists a measurable function $g_{\tilde{f}}$ such that

$$\mathbb{E}[\tilde{f}(X_2) \mid X_1, W_1, W_2] = g_{\tilde{f}}(X_1, W_2) \quad \text{a.s.}$$

Integrating $\tilde{f}(X_2)$ with respect to (38) and rearranging the terms, we obtain

$$W_1(\tilde{f}(X_1) - g_{\tilde{f}}(X_1, W_2)) = \theta(g_{\tilde{f}}(X_1, W_2) - \mathbb{E}[\tilde{f}(X_1)]) \quad \text{a.s.}$$

Assume that μ'_0 is non-degenerate. Then, there is an \tilde{f} such that $\mathbb{P}(\tilde{f}(X_1) = \mathbb{E}[\tilde{f}(X_1)]) = 0$; e.g., take $\tilde{f} = \mathbb{1}_B$ for some $B \in \mathcal{X}$ such that $0 < \mathbb{P}(X_1 \in B) < 1$. It follows that

$$\begin{aligned} \mathbb{P}(\tilde{f}(X_1) = g_{\tilde{f}}(X_1, W_2) = 0) &= \mathbb{P}(\tilde{f}(X_1) = \mathbb{E}[\tilde{f}(X_2) \mid X_1, W_1, W_2]) \\ &= \mathbb{P}(\tilde{f}(X_1) = \mathbb{E}[\tilde{f}(X_1)]) = 0; \end{aligned}$$

therefore,

$$W_1 = \frac{\theta(g_{\tilde{f}}(X_1, W_2) - \mathbb{E}[\tilde{f}(X_1)])}{\tilde{f}(X_1) - g_{\tilde{f}}(X_1, W_2)} \quad \text{a.s.}$$

In other words, there exists a measurable function \tilde{h} such that $W_1 = \tilde{h}(X_1, W_2)$ a.s., and so $W_2 = \tilde{h}(X_1, W_1)$ a.s. by partial exchangeability. It follows from $X_1 \perp (W_1, W_2)$ that, for every $A \in \mathcal{B}(\mathbb{R}_+)$,

$$\mathbb{P}(W_2 \in A \mid W_1) = \mathbb{P}(W_2 \in A \mid X_1, W_1) = \mathbb{1}_A(W_2) \quad \text{a.s.};$$

thus, $W_2 = W_1$ a.s. and, from exchangeability, $W_n = W_1$ a.s., $n \geq 1$. \square

4.2. Asymptotic Properties of GRRPP with Exchangeable Weights

It follows from (38) that the GRRPP with exchangeable weights is a mixture of RRPPs with independent weights, with the mixing distribution affecting only the sequence $(W_n)_{n \geq 1}$. Thus, we expect that the results in Section 3.3 carry over to this more general setting. In this section, we concentrate on the behavior of θ_n and the sequence $(L_n)_{n \geq 1}$.

Assume that $\mathbb{P}(W_1 > 0 \mid \tilde{\eta}) > 0$. If $\mathbb{E}[W_1] < \infty$, then $0 < \mathbb{E}[W_1 \mid \tilde{\eta}] < \infty$ a.s., and, by the law of large numbers for exchangeable random variables (see [1], Section 2),

$$\frac{1}{n} \sum_{i=1}^n W_i \xrightarrow{a.s.} \mathbb{E}[W_1 \mid \tilde{\eta}] \in (0, +\infty).$$

Then, if μ_0 is diffuse, $n \cdot \theta_n \xrightarrow{a.s.} \theta / \mathbb{E}[W_1 \mid \tilde{\eta}]$ and $\sum_{i=1}^n \theta_n = \infty$ a.s., so Theorem 1 in [27] implies

$$\frac{L_n}{\log n} = \frac{L_n}{\sum_{k=1}^n \theta_k} \left(\frac{1}{\log n} \sum_{k=1}^n \frac{1}{k} (k \cdot \theta_k) \right) \xrightarrow{a.s.} \frac{\theta}{\mathbb{E}[W_1 \mid \tilde{\eta}]}.$$

If $\mathbb{E}[W_1] = \infty$, then L_n may converge to a finite limit, as $n \rightarrow \infty$. For example, let us consider a strictly stable reinforcement distribution as in Proposition 4.

Proposition 6. *Let $(\mu_n)_{n \geq 0}$ be a GRRPP with parameters (V, μ_0, η) such that V is a strictly positive random variable with $\mathbb{E}[V^{-1}] < \infty$, μ_0 is diffuse, and $\eta(v)$, $v > 0$ is a $S_\alpha(1, v, 0)$ distribution with stability parameter $\alpha < 1$. Then, $\theta_n = O_P(n^{-1/\alpha})$ and*

$$\lim_{n \rightarrow \infty} L_n < \infty \quad \text{a.s.}$$

Proof. It follows from how the weights in the representation (35) are chosen that we can take

$$W_n = VF^{-1}(U_n),$$

where $U_n \sim \text{Unif}[0, 1]$, $U_n \perp (V, X_1, U_1, \dots, X_{n-1}, U_{n-1}, X_n)$, and F^{-1} is the inverse of the $S_\alpha(1, 1, 0)$ distribution function. Then,

$$\theta_n = \frac{\theta}{\theta + \sum_{i=1}^n W_i} \leq n^{-1/\alpha} \frac{\theta}{Vn^{-1/\alpha} \sum_{i=1}^n F^{-1}(U_i)} \stackrel{d}{=} n^{-1/\alpha} \frac{\theta}{VY},$$

for some $Y \sim S_\alpha(1, 1, 0)$ such that $Y \perp V$. It follows for every $M > 0$ that $\mathbb{P}(n^{1/\alpha}\theta_n > M) \leq \mathbb{P}(\theta/VY > M)$, which can be made arbitrarily small by taking M large enough. Regarding the second assertion, as $1/\alpha > 1$ and $\mathbb{E}[\theta/(VY)] < \infty$ by Theorem 5.4.1 in [28], we have

$$\mathbb{E}[\lim_{n \rightarrow \infty} L_n] = \lim_{n \rightarrow \infty} \sum_{i=1}^n \mathbb{E}[\mathbb{1}_{\{L_i=L_{i-1}+1\}}] = \sum_{n=1}^{\infty} \mathbb{E}[\theta_n] \leq \sum_{n=1}^{\infty} \frac{\theta}{n^{1/\alpha}} \mathbb{E}[1/VY] < \infty.$$

□

Extensions of Proposition 6 can be obtained by exploiting the central limit theorems for exchangeable random variables, which are found in [30,31].

5. Discussion

In this paper, we study the extension of randomly reinforced urns [17] to an unbounded set of possible colors. The resulting measure-valued urn process provides a predictive characterization of the law of an asymptotically exchangeable sequence of random variables, which corresponds to the observation process of an implied urn sampling scheme. In fact, the model (6)–(7) fits into a line of recent research, which explores efficient predictive constructions for fast online prediction or approximately-Bayesian solutions, see [11,29,32] and references therein. To that end, one direction for future work is to generalize the functional relationship in (7) and/or, as one referee suggested, to consider finitely-additive measures, along the lines discussed in [33].

We investigate the asymptotic properties of the sequences of predictive distributions and empirical frequencies of the observation process, and prove their convergence in total variation distance to a common random limit. The rate of convergence of their difference is given set-wise; so, another possible direction for future research is to consider a stronger distance. As far as we know, the topic of merging of the predictive and empirical distributions is largely unexplored. Within the relevant literature, we mention the works of [4,34], where the authors study the rate of convergence of the Wasserstein or Prokhorov distances under exchangeability, and the papers by Berti et al. [21,35], who consider the c.i.d. case and regard the difference between the predictive and empirical measures as a map in the space of real bounded functions.

Author Contributions: Formal analysis, S.F., S.P., H.S.; writing—original draft preparation, S.F., S.P., H.S.; writing—review and editing, S.F., S.P., H.S. All authors have read and agreed to the published version of the manuscript.

Funding: This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 817257). H.S. was partially supported by the Bulgarian Ministry of Education and Science under the National Research Programme “Young scientists and postdoctoral students” approved by DCM No. 577/17.08.2018.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Acknowledgments: We wish to express our sincere gratitude to Regazzini for his deeply inspiring ideas and for instilling in us his passion for research. We thank the four anonymous referees for the valuable comments.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Aldous, D.J. Exchangeability and related topics. *École D’Été De Probab. De St.-Flour XIII 1983* **1985**, *1117*, 1–198.
2. Fortini, S.; Ladelli, L.; Regazzini, E. Exchangeability, predictive distributions and parametric models. *Sankhya Ser. A* **2000**, *62*, 86–109.
3. Cifarelli, D.M.; Regazzini, E. De Finetti’s contribution to probability and statistics. *Statist. Sci.* **1996**, *11*, 253–282. [[CrossRef](#)]
4. Cifarelli, D.M.; Dolera, E.; Regazzini, E. Frequentistic approximations to Bayesian prevision of exchangeable random elements. *Int. J. Approx. Reason.* **2016**, *78*, 138–152. [[CrossRef](#)]
5. Fortini, S.; Petrone, S. Predictive distribution (de Finetti’s view). In *Wiley StatsRef: Statistics Reference Online*; Wiley Online Library: 2014; pp. 1–9. Available online: <https://onlinelibrary.wiley.com/doi/full/10.1002/9781118445112.stat07831> (accessed on 4 October 2021).
6. Regazzini, E. Old and recent results on the relationship between predictive inference and statistical modeling either in nonparametric or parametric form. In *Bayesian Statistics 6*; Oxford University Press: Oxford, UK, 1999; pp. 571–588.
7. Fortini, S.; Petrone, S. Predictive construction of priors in Bayesian nonparametrics. *Braz. J. Probab. Stat.* **2012**, *26*, 423–449. [[CrossRef](#)]
8. Mailler, C.; Marckert, J.F. Measure-valued Pólya urn processes. *Electron. Commun. Probab.* **2017**, *22*, 33. [[CrossRef](#)]
9. Janson, S. Random replacements in Pólya urns with infinitely many colours. *Electron. Commun. Probab.* **2019**, *24*, 11. [[CrossRef](#)]
10. Aletti, G.; Ghiglietti, A.; Rosenberger, W.F. Nonparametric covariate-adjusted response-adaptive design based on a functional urn model. *Ann. Stat.* **2018**, *46*, 3838–3866. [[CrossRef](#)]
11. Fortini, S.; Petrone, S. Quasi-Bayes properties of a procedure for sequential learning in mixture models. *J. R. Stat. Soc. Ser. B* **2020**, *82*, 1087–1114. [[CrossRef](#)]
12. Zhang, L.X.; Hu, F.; Cheung, S.H.; Chan, W.S. Immigrated urn models—Theoretical properties and applications. *Ann. Stat.* **2011**, *39*, 643–671. [[CrossRef](#)]
13. Blackwell, D.; MacQueen, J.B. Ferguson distributions via Pólya urn schemes. *Ann. Stat.* **1973**, *1*, 353–355. [[CrossRef](#)]
14. Bandyopadhyay, A.; Thacker, D. Pólya urn schemes with infinitely many colors. *Bernoulli* **2017**, *23*, 3243–3267. [[CrossRef](#)]
15. Janson, S. A.s. convergence for infinite colour Pólya urns associated with random walks. *Ark. Mat.* **2021**, *59*, 87–123. [[CrossRef](#)]
16. Mailler, C.; Villemonais, D. Stochastic approximation on non-compact measure spaces and application to measure-valued Pólya processes. *Ann. Appl. Probab.* **2020**, *30*, 2393–2438. [[CrossRef](#)]
17. Muliere, P.; Paganoni, A.M.; Secchi, P. A randomly reinforced urn. *J. Stat. Plan. Inference* **2006**, *136*, 1853–1874. [[CrossRef](#)]
18. Bassetti, F.; Crimaldi, I.; Leisen, F. Conditionally identically distributed species sampling sequences. *Adv. Appl. Probab.* **2010**, *42*, 433–459. [[CrossRef](#)]
19. Berti, P.; Crimaldi, I.; Pratelli, L.; Rigo, P. Rate of convergence of predictive distributions for dependent data. *Bernoulli* **2009**, *15*, 1351–1367. [[CrossRef](#)]
20. Berti, P.; Crimaldi, I.; Pratelli, L.; Rigo, P. Central limit theorems for multicolor urns with dominated colors. *Stoch. Process. Appl.* **2010**, *120*, 1473–1491. [[CrossRef](#)]
21. Berti, P.; Pratelli, L.; Rigo, P. Limit theorems for a class of identically distributed random variables. *Ann. Probab.* **2004**, *32*, 2029–2052. [[CrossRef](#)]
22. Crimaldi, I. An almost sure conditional convergence result and an application to a generalized Pólya urn. *Int. Math. Forum* **2009**, *4*, 1139–1156.
23. Sariev, H.; Fortini, S.; Petrone, S. Infinite-Color Randomly Reinforced Urns with Dominant Colors. 2021. Preprint. Available online: <https://arxiv.org/abs/2106.04307> (accessed on 4 October 2021).
24. Fortini, S.; Petrone, S.; Sporysheva, P. On a notion of partially conditionally identically distributed sequences. *Stoch. Process. Appl.* **2018**, *128*, 819–846. [[CrossRef](#)]
25. Kallenberg, O. *Foundations of Modern Probability*, 3rd ed.; Springer: New York, NY, USA, 2021.

26. Häusler, E.; Luschgy, H. *Stable Convergence and Stable Limit Theorems*; Springer: Cham, Switzerland, 2015.
27. Dubins, L.; Freedman, D. A sharper form of the Borel-Cantelli lemma and the strong law. *Ann. Math. Stat.* **1965**, *36*, 800–807. [[CrossRef](#)]
28. Uchaikin, V.V.; Zolotarev, V.M. *Chance and Stability: Stable Distributions and Their Applications*; Walter de Gruyter: Berlin, Germany, 2011.
29. Fong, E.; Holmes, C.; Walker, S. Martingale Posterior Distributions 2021. Preprint. Available online: <https://arxiv.org/abs/2103.15671> (accessed on 4 October 2021).
30. Fortini, S.; Ladelli, L.; Regazzini, E. A central limit problem for partially exchangeable random variables. *Theory Probab. Appl.* **1997**, *41*, 224–246. [[CrossRef](#)]
31. Fortini, S.; Ladelli, L.; Regazzini, E. Central limit theorem with exchangeable summands and mixtures of stable laws as limits. *Boll. Unione Mat. Ital.* **2012**, *5*, 515–542.
32. Berti, P.; Dreassi, E.; Pratelli, L.; Rigo, P. A class of models for Bayesian predictive inference. *Bernoulli* **2021**, *27*, 702–726. [[CrossRef](#)]
33. de Cooman, G.; Bock, J.D.; Diniz, M.A. Coherent predictive inference under exchangeability with imprecise probabilities. *J. Artif. Intell. Res.* **2015**, *52*, 1–95. [[CrossRef](#)]
34. Dolera, E.; Regazzini, E. Uniform rates of the Glivenko-Cantelli convergence and their use in approximating Bayesian inferences. *Bernoulli* **2019**, *25*, 2982–3015. [[CrossRef](#)]
35. Berti, P.; Pratelli, L.; Rigo, P. Limit theorems for empirical processes based on dependent data. *Electron. J. Probab.* **2012**, *17*, 1–18. [[CrossRef](#)]