

UNIVERSITÁ COMMERCIALE "LUIGI BOCCONI"

PhD SCHOOL

PhD in Statistics

31st Cycle

Disciplinary field: SECS-S/01

Thesis Title: *Stochastic processes of the urn type with
convergent predictive distributions*

Advisor: Prof. Sandra FORTINI

Co-Advisor: Prof. Sonia PETRONE

PhD Thesis by

Hristo SARIEV

ID number: 1326180

Academic Year: 2019/2020

THESIS DECLARATION

I, the undersigned

Sariev, Hristo

Student ID number: *1326180*

Thesis Title: *Stochastic processes of the urn type with convergent predictive distributions*

PhD in Statistics

31st Cycle

SECS-S/01

DECLARE

Under *my* responsibility:

- 1) that, according to Italian Republic Presidential Decree no. 445, 28th December 2000, mendacious declarations, falsifying records and the use of false records are punishable under the Italian penal code and related special laws. Should any of the above prove true, all benefits included in this declaration and those of the temporary "embargo" are automatically forfeited from the beginning;
- 2) that the University has the obligation, according to art. 6, par. 11, Ministerial Decree no. 224, 30th April 1999, to keep a copy of the thesis on deposit at the "Biblioteche Nazionali Centrali" (Italian National Libraries) in Rome and Florence, where consultation will be permitted, unless there is a temporary "embargo" protecting the rights of external bodies and the industrial/commercial exploitation of the thesis;
- 3) that I will submit the thesis online in unalterable format to Bocconi University, that, by means of "Institutional Research Information System (IRIS)", will permit online consultation of the complete text (except in cases of temporary "embargo");
- 4) the thesis is protected by the regulations governing copyright (Italian law no. 633, 22nd April 1941 and subsequent modifications). The exception is the right of Università Commerciale "Luigi Bocconi" to reproduce the same, quoting the source, for research and teaching purposes;
- 5) that the copy of the thesis submitted online is identical to the copies handed in/sent to the members of the Thesis Board and to any other paper or digital copy deposited at the University

offices, and, as a consequence, the University is absolved from any responsibility regarding errors, inaccuracy or omissions in the contents of the thesis;

- 6) that the contents and organization of the thesis is an original work carried out by the undersigned and does not in any way compromise the rights of third parties (Italian law, no. 633, 22nd April 1941 and subsequent integrations and modifications), including those regarding security of personal details; therefore the University is in any case absolved from any responsibility whatsoever, civil, administrative or penal, and shall be exempt from any requests or claims from third parties;
- 7) that the thesis is not subject to "embargo", i.e. that it is not the result of work included in the regulations governing industrial property; it was not written as part of a project financed by public or private bodies with restrictions on the diffusion of the results; is not subject to patent or protection registrations.

25th November 2019

Sariev, Hristo

Contents

I	Introduction	1
1.1	Motivation	1
1.2	Literature review	3
1.3	Preliminaries	4
II	Randomly reinforced Pólya sequence	8
2.1	Introduction	8
2.1.1	Model	9
2.2	Generalized Ottawa sequences	10
2.3	Measure-valued Pólya urn processes	13
III	Conditionally identically distributed randomly reinforced Pólya sequence	18
3.1	Introduction	18
3.1.1	Model	19
3.2	Characterization of directing measure	19
3.2.1	Almost sure discreteness	19
3.2.2	Random partition	27
3.2.3	C.i.d. RRPS with 0-1 weights	27
3.3	Central limit theorem	30
3.3.1	Central limit theorem	30
3.3.2	Almost sure discreteness - revisited	35
3.4	Inference	36
3.4.1	Case 1: W_1, \dots, W_n are observable	36

3.4.2	Case 2: W_1, \dots, W_n are non-observable	37
IV Dominant Pólya sequence		39
4.1	Introduction	39
4.1.1	Model	40
4.2	First-order convergence results	41
4.2.1	DPS with strictly increasing w	41
4.2.2	DPS with general w	49
4.2.3	Convergence on the dominant subsequence	52
4.2.4	DPS with a discontinuity at \mathcal{D}	55
4.2.5	Random partition	57
4.3	Central limit theorems	59
4.3.1	Central limit theorem for dominant subsequence	59
4.3.2	Central limit theorem for original sequence	64
4.4	DPS with finite number of colors	66
4.5	Inference	70
4.5.1	Confidence intervals for the random limit	70
4.5.2	Credible intervals for the random limit	71
V Further discussion		73
5.1	Weighted Chinese Restaurant Process	73
5.2	Clinical trials	74
5.3	Randomly reinforced Poisson-Dirichlet sequence	76
5.4	Multi-experiment data	77
5.4.1	Partially c.i.d. interacting RRPS	78
5.4.2	Related models	79
5.5	Conclusion	80
A Miscellaneous results		82

Acknowledgements

First and foremost, this thesis would have been impossible without the guidance and teachings of my advisors, Sandra Fortini and Sonia Petrone. In fact, none of it would have been possible had Sonia not given me the chance of a lifetime almost five years ago. I can say only now that I had been at a constant pressure this whole time trying not to disappoint you. This work should, at least in part, reflect that. As for Sandra, I am grateful to have not missed the opportunity of working with you for the second time around. You are the academic that I want to see me become.

My sincere gratitudes go to the faculty of the Department of Decision Sciences and the administrative staff at Bocconi University for the indispensable support they had provided me with throughout my studies. I am also grateful to the group at IMT Lucca, headed by Irene Crimaldi, for the exchange of ideas and experiences. As a matter of fact, large parts of the thesis would have been hard for me to come up with without the work done by Irene and her coauthors. I owe a personal thank you to Stefano Favaro for the time and attention paid, and all of his input on the thesis.

I would like to thank my fellow doctoral students, not the least of which the ones who left pre-maturely, including my dear friend Ko Sugiura, for providing a positively competitive environment that allowed me to develop myself. I am happy to acknowledge that the people I graduate with – Andrea Arfé, Giovanni Rabitti, Tommaso Rigon, are all excellent researchers and amazing human beings. I wish them nothing but the best ahead.

Last but not least, I would like to pay my deepest respects to my family and close friends for the understanding, love and unconditional support.

I dedicate this work to my mother and father. I only hope that I have made you proud. To my brother, I hope I have given you something to look up to.

Rada and Ivaila, I love you more than anything.

Abstract

In this work we propose a general class of stochastic processes with random reinforcement that are extensions of the celebrated Pólya sequence by Blackwell and MacQueen [*Ann. Stat.* **1** (1973) 353–355]. The resulting randomly reinforced Pólya sequence (RRPS) can be described as an urn scheme with countable number of colors and a general replacement rule. Under assumptions of conditional independence between reinforcement and observation, a RRPS becomes conditionally identically distributed (in the sense of [*Ann. Probab.* **32** (2004) 2029–2052]), and thus predictively convergent, in which case we show that it is asymptotically equivalent in law to an exchangeable species sampling sequence. This result has important implications on the generated random partition, which can be visualized as a weighted version of the Chinese Restaurant Process. We then provide complete distributional characterization of the predictive limit for the model with dichotomous reinforcements. Throughout the second part of the thesis, we consider an alternative specification of the replacement mechanism of a RRPS, whereby we deem some colors to be probabilistically dominant. In this situation the predictive and empirical distributions evaluated near the set of dominant colors both tend to 1. In fact, under some further restrictions on the reinforcement, the predictive and empirical distributions converge in the sense of almost sure weak convergence to one and the same random probability measure, whose mass is concentrated on the dominant set. As a consequence, the process becomes asymptotically exchangeable and its law – directed by the above random measure, so that the data structure gets relatively sparse with time. The predictive limit for both models is generally unknown, however, so we derive central limit results, with which to approximate its distribution. The last chapter of the thesis is addressed towards applications of the RRPS, with the dominant-color model being considered in the context of clinical trials with response-adaptive design. Sections discussing uni- and multivariate extensions of the RRPS complete our study.

Chapter I

Introduction

1.1 Motivation

There are several, more or less overlapping ideas that underlie the study of the class of stochastic processes that we propose. The most apparent reason for their development is that they act as non-trivial extensions of already existing models, which are particularly well-established within Bayesian statistics. On the other hand, different specifications of our basic model can be considered as relatively simple examples of some broader classes of processes, whose probabilistic behavior goes beyond exchangeability. In fact, as the title of the thesis suggests, we study stochastic processes with reinforcement (see Pemantle, 2007, for an insightful review) and the latter is a powerful precondition for the system of predictive distributions to be convergent. In turn, predictive convergence implies that the sequence of observations becomes asymptotically exchangeable (Aldous, 1985); thus, we envision the use of our model for the purposes of Bayesian analysis, which has actually inspired us to take an overall predictive approach to modeling.

The central tenet of Bayesian philosophy as set forth by Bruno de Finetti is the assessment of probability as being subjective in nature, stemming from the individual's incomplete information about observable events. The fundamental probabilistic framework to work within in Bayesian statistics is that of exchangeability, which amounts to an invariance judgement on the group of permutations. A famous result due to de Finetti (1931) and Hewitt and Savage (1955) asserts that any exchangeable random sequence may be represented by a unique *prior* probability distribution on some induced parameter space. Complete characterization of the process requires the adoption of additional probabilistic assumptions, which would ideally involve only observable quantities, instead of a prior guess, and there is already an extensive literature on various characterizations of exchangeable random sequences through the system of predictive distributions (refer to Fortini and Petrone, 2012b, for a review). In fact, predictive constructions of exchangeable processes have played a crucial part in the development of nonparametric Bayesian theory as they characterize some fundamental prior distributions (see Blackwell and MacQueen, 1973; Pitman and Yor, 1997, etc.) and are closely related to the theory of exchangeable random partitions (Pitman, 1995, 1996). Moreover, de Finetti (1937) himself emphasized the role of the prediction in the Bayesian statistical approach as its ultimate goal.

Existing exchangeable predictive constructions have one core feature in common in that they give rise to stochastic processes with reinforcement (and convergent predictive distributions). As a consequence, the sampling mechanism that they describe can often be interpreted in terms of sequential draws with replacement from an urn containing balls of different colors. The model, which we propose, exhibits a similar behavior, yet we require the reinforcement to be further randomized, which in general fails to guarantee the exchangeability of the process or even that the limit of the predictive distributions exists. At the same time, however, the urn scheme implied by our process acts as an extension of some well-known urn models to the case of infinite colors, for which there is already an established theory that leads to the existence of a predictive limit.

In general, we envision the following two roles that non-exchangeable stochastic processes with convergent predictive distributions would have in Bayesian statistics. On the one hand, these models could be adopted in situations, where exchangeability is broken by asymmetries, forms of selection, competition, temporary disequilibrium, or interactions with other concomitant processes (references include Bassetti et al., 2010; Fortini et al., 2018), but the process at hand is assumed to have some stabilizing behavior with time. Indeed, predictively convergent sequences are asymptotically exchangeable, which means, roughly speaking, that one becomes progressively more probabilistically indifferent to the order in which observations arrive. Probability laws with convergent predictive distributions could then be used as approximations of the exchangeable laws that are directed by the former’s predictive limits since they become asymptotically indistinguishable. This strategy can be of particular interest if the predictive distributions of the non-exchangeable process are actually easier to handle because of the less structure imposed on them, thereby speeding up the estimation process at the cost of some imprecision (see Fortini and Petrone, 2019, for an example).

A particular probabilistic framework that is consistent with predictive convergence is that of conditional identity in distribution. A sequence of random variables is said to be conditionally identically distributed (c.i.d.) if at any time of the observation process future observations have the same marginal distribution, given all past information. Conditional identity in distribution is studied extensively by Berti et al. (2004), although Kallenberg (1988) has shown earlier that it is equivalent to exchangeability under stationarity. Besides being sufficient for predictive convergence, conditional identity in distribution implies that the empirical process itself converges to the same random limit. The asymptotic agreement of frequency and prediction constitutes one of the main pillars of Bayesian analysis as it provides a “frequentist basis of the subject’s probabilistic learning” (Fortini et al., 2018), and hence c.i.d. sequences form an important subclass of predictively convergent processes, which has only recently been used for the purposes of Bayesian inference (see, e.g. Airoldi et al., 2014; Cassese et al., 2019).

In Chapter II of the thesis we propose a general family of stochastic processes with random reinforcement that can simultaneously be regarded as a generalization of the predictive construction of the Dirichlet process prior in the direction beyond exchangeability, and as an extension of some finite-color urn models to broader sampling schemes that allow for a continuum of colors. We then investigate two specifications of our basic model, the first of which is studied in Chapter III and is associated with a c.i.d. sequence of random variables. In contrast, the model of Chapter IV is motivated by a concrete application that requires reinforcement and observation to be contemporaneously dependent, which violates the conditional identity in distribution assumption. Nonetheless, the stochastic processes from both chapters have convergent pre-

dictive distributions, whose random limits are the main object of study in this thesis. Results will be mostly probabilistic in nature, so in Chapter V we relate them to a wide-range of applied problems such as cluster analysis, clinical trials with response-adaptive design, and Bayesian inference in multi-experiment setting. Throughout the rest of this chapter we try to make some ideas more concrete by going through the relevant literature and by providing some technical definitions and results that will be used later on.

1.2 Literature review

Consider the classical *two-color Pólya urn* scheme (Pólya, 1930; Freedman, 1965), which describes the composition of an urn containing balls of two different colors that are then sequentially sampled and reinforced with an additional ball of the observed color. This represents arguably the simplest example of a (exchangeable) stochastic process with reinforcement. Janson (2019) outlines the different ways, in which one might generalize the two-color Pólya urn; namely, consider additional colors, reinforce with balls of any color, do so with a non-negative real quantity (in that case *number* of balls becomes a misnomer), randomize the reinforcement, remove balls, or conceptualize to the infinite-color case. The book of Mahmoud (2008) covers many of the suggested modifications, with some notable extensions of the Pólya urn scheme to k colors being the generalized Pólya urn (Athreya and Ney, 1972), the randomly reinforced urn (Muliere et al., 2006), and the immigrated urn model (Zhang et al., 2011). Recent developments to k -color urn models include the introduction of covariates (Aletti et al., 2018a) and thresholds (Aletti et al., 2018b) into the reinforcement.

The *Pólya sequence* of Blackwell and MacQueen (1973), which characterizes the predictive construction of the Dirichlet process prior, is the conceptual extension of the two-color Pólya urn to the case of countable colors. A proper urn model to explain the sampling procedure of a Pólya sequence is *Hoppe's urn* (Hoppe, 1984; Fortini and Petrone, 2012a), which is another way of stating the famous Chinese Restaurant Process (CRP) metaphor (Aldous, 1985) and can be described as follows. Imagine an urn containing initially only black balls. The reinforcement rule then reads: each time a black ball is picked, it is returned together with a ball, whose color is generated from some given distribution on colors; should a non-black ball be sampled, it is returned with another ball of the same color. Strictly speaking, the CRP records just the ordering of the drawn balls, regardless of the actual labels attached to them, whereas what we have described above is a "colored" version of the Hoppe's urn/CRP.

The property of Pólya sequences to generate new colors only when the need occurs is a fundamental feature of exchangeable species sampling sequences (Pitman, 1996) or more generally of *generalized species sampling sequences* (Bassetti et al., 2010), which both treat the Pólya sequence as a basic example. On the other hand, the urn composition of a Pólya sequence can be regarded as a measure-valued process, in which case reinforcement amounts to a "summation" of measures. The latter condition is defining for the *measure-valued Pólya urn processes* of Janson (2019), which also emphasize the Markovianity of the process. Both of these frames of reference will prove useful in the subsequent chapters.

Our main model, which is defined in Chapter II, is an extension of the Pólya sequence along the lines of random reinforcement, hence the moniker, *randomly reinforced Pólya sequence* (RRPS). Under some conditional independence assumptions (Chapter II, Section 2.2), a RRPS becomes a member of the class of

c.i.d. generalized species sampling sequences. As such its system of predictive distributions constitutes a measure-valued martingale and converges in the sense of almost sure weak convergence to a random probability measure. Chapter III is devoted to the study of that random limit, with our results showing, in particular, that c.i.d. RRPSs become asymptotically equivalent to exchangeable species sampling sequences that are directed by the former's predictive limits. In Chapter V, we look at the random partition associated to a c.i.d. RRPS in terms of a weighted version of the CRP.

Under quite different conditions (Chapter II, Section 2.3), the predictive distributions of a RRPS can be regarded as a measure-valued Pólya urn process. Our study of these constructions is further motivated by the introduction of a mechanism, through which some colors will be able to dominate the rest. The resulting RRPS, which we call a *dominant Pólya sequence* (DPS), is an extension of the k -color randomly reinforced urn model (see Muliere et al., 2006; Berti et al., 2010) towards a continuum of colors. It follows from results in Chapter IV that in some situations the predictive distributions of a DPS converge to a random limit measure with a sparse support, in the sense that the predictive distributions will tend to concentrate their total mass on a small subset of colors, namely the dominant ones. This feature of the DPS will be discussed in Chapter V in relation to response-adaptive designs of clinical trials.

The RRPS, overall, and the c.i.d. RRPS and the DPS, in particular, help fill the gap on Pólya urn models with countable number of colors and random reinforcement. As such, they establish a general modelling framework that, in addition to accomodating many of the existing finite-color urn schemes, provides important examples of some extensive classes of models.

1.3 Preliminaries

Definitions and notation

Throughout the thesis, (\mathbb{X}, d) will be a complete and separable metric space (c.s.m.s), $\tau_{\mathbb{X}}$ the generated topology, and \mathcal{X} the associated Borel σ -algebra. All random quantities will be defined on the same probability space $(\Omega, \mathcal{H}, \mathbb{P})$. The following symbols

$$M(\mathcal{X}), \quad M_b(\mathcal{X}), \quad M_+(\mathcal{X}), \quad M_0(\mathcal{X}), \quad C(\mathbb{X}), \quad L^p(\mathcal{X}),$$

are used to denote the collections of all numerical functions, $f : \mathbb{X} \rightarrow \bar{\mathbb{R}}$, that are \mathcal{X} -measurable, bounded and \mathcal{X} -measurable, non-negative and \mathcal{X} -measurable, simple and \mathcal{X} -measurable, continuous, and p th power integrable with $p \geq 1$, respectively. In general, $\mathcal{F} = (\mathcal{F}_n)_{n \geq 0}$ will be a filtration on (Ω, \mathcal{H}) with $\mathcal{F}_0 = \{\emptyset, \Omega\}$. Whenever the filtration is associated to a stochastic process, say, $X = (X_n)_{n \geq 1}$, then it will have a corresponding superscript attached, e.g. $\mathcal{F}^X = (\mathcal{F}_n^X)_{n \geq 0}$ given by $\mathcal{F}_n^X = \sigma(X_1, \dots, X_n)$, for $n \geq 1$. The σ -algebra $\mathcal{F}_\infty := \bigvee_{n=1}^{\infty} \mathcal{F}_n$ captures, loosely speaking, the total information in the random experiment. We will also denote by

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

the collections of all measures on \mathbb{X} , finite measures on \mathbb{X} , finite non-null measures on \mathbb{X} , and probability measures on \mathbb{X} , respectively. By Theorem 1.5 of Kallenberg (2017), $\mathbb{M}_F(\mathbb{X})$ (and $\mathbb{M}_F^*(\mathbb{X})$, $\mathbb{M}_P(\mathbb{X})$) is c.s.m.s., in which case the σ -algebra $\mathcal{M}_F(\mathbb{X})$ (and $\mathcal{M}_F^*(\mathbb{X})$, $\mathcal{M}_P(\mathbb{X})$) generated by the topology of weak convergence

is the smallest one that makes the map $\mu \mapsto \mu(B)$ measurable, for each $B \in \mathcal{X}$. Let $\nu \in \mathbb{M}_P(\mathbb{X})$. The support of ν is defined as the set of points $x \in \mathbb{X}$, whose every open neighborhood has a positive measure, namely,

$$\text{supp}(\nu) := \{x \in \mathbb{X} : \nu(U) > 0, \text{ for each } U \in \tau_{\mathbb{X}} : x \in U\}.$$

Let X be an \mathbb{X} -valued random variable on $(\Omega, \mathcal{H}, \mathbb{P})$, $\nu \in \mathbb{M}(\mathbb{X})$, and $f \in M(\mathcal{X})$. By $X \sim \nu$ we would mean that $\mathbb{P}(X \in \cdot) = \nu(\cdot)$, whereas $\nu \circ f^{-1}(\cdot) = \nu(f^{-1}(\cdot)) = \mathbb{P}(f(X) \in \cdot)$ is used to denote the measure induced by f .

Let $(\mathbb{Y}, \mathcal{Y})$ be a measurable space. A *transition kernel* from \mathbb{Y} to \mathbb{X} is a function $K : \mathbb{Y} \times \mathcal{X} \rightarrow \bar{\mathbb{R}}_+$ such that the map $y \mapsto K(y, B)$ is \mathcal{Y} -measurable, for every $B \in \mathcal{X}$, and the map $B \mapsto K(y, B)$ is a measure on \mathbb{X} , for every $y \in \mathbb{Y}$. The transition kernel K said to be finite if $K(y, \mathbb{X}) < \infty$, for $y \in \mathbb{Y}$. A *transition probability kernel* from \mathbb{Y} to \mathbb{X} is a transition kernel K from \mathbb{Y} to \mathbb{X} such that $K(y, \mathbb{X}) = 1$. The collections of all transition kernels, finite transition kernels and transition probability kernels from \mathbb{Y} to \mathbb{X} are indicated by

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

respectively. We would denote by $\mathcal{N}(\mu, \sigma^2)$ the Gaussian transition probability kernel with (random) parameters $(\mu, \sigma^2) \in \mathbb{R} \times \mathbb{R}_+$.

Let $\nu \in \mathbb{M}_P(\mathbb{Y})$ and $K \in \mathbb{K}_P(\mathbb{Y}, \mathbb{X})$. The set function $K \times \nu$, given by

$$K \times \nu(B) := \int_{\mathbb{Y}} \int_{\mathbb{X}} \mathbb{1}_B(x, y) K(y, dx) \nu(dy), \quad \text{for } B \in \mathcal{X} \otimes \mathcal{Y},$$

defines a probability measure on the product space $\mathbb{X} \times \mathbb{Y}$, endowed with the σ -algebra $\mathcal{X} \otimes \mathcal{Y}$ that is generated by the measurable rectangles. If $K(y, \cdot) = \mu(\cdot)$, for all $y \in \mathbb{Y}$ and some $\mu \in \mathbb{M}_P(\mathbb{X})$, then $K \times \nu$ coincides with the product probability measure on $\mathbb{X} \times \mathbb{Y}$, which we denote by $\mu \otimes \nu$.

Let $K, L \in \mathbb{K}_P(\mathbb{Y}, \mathbb{X})$. The product transition probability kernel between K and L is the transitional probability kernel $K \otimes L$ from \mathbb{Y} to \mathbb{X} , defined by

$$(K \otimes L)(y, \cdot) := K(y, \cdot) \otimes L(y, \cdot), \quad \text{for } y \in \mathbb{Y}.$$

A *random probability measure* on \mathbb{X} is a transition probability kernel $\tilde{P} \in \mathbb{K}_P(\Omega, \mathbb{X})$. It follows that \tilde{P} can be regarded on c.s.m.s. spaces as an $\mathcal{H} \setminus \mathcal{M}_P(\mathbb{X})$ -measurable map $\tilde{P} : \Omega \rightarrow \mathbb{M}_P(\mathbb{X})$. An example of a random probability measure \tilde{P} is the Dirichlet process with parameters $\theta \in \mathbb{R}_+$ and $\nu \in \mathbb{M}_P(\mathbb{X})$, denoted $\tilde{P} \sim \text{DP}(\theta, \nu)$, where $(\tilde{P}(B_1), \dots, \tilde{P}(B_n)) \sim \text{Dir}(\theta\nu(B_1), \dots, \theta\nu(B_n))$ is Dirichlet distributed, for any finite partition $\{B_i\}_{i=1}^n$ of \mathbb{X} in \mathcal{X} .

Let X be an \mathbb{X} -valued random variable on $(\Omega, \mathcal{H}, \mathbb{P})$, and $\mathcal{F} \subseteq \mathcal{H}$ be a sub- σ -algebra. The *conditional distribution* of X given \mathcal{F} is any random probability measure \tilde{P} on \mathbb{X} such that

$$\tilde{P}(B) = \mathbb{P}(X \in B | \mathcal{F}) \quad \text{a.s.}[\mathbb{P}], \text{ for } B \in \mathcal{X}.$$

From the fundamental property of conditional distributions, $\mathbb{E}[f(X) | \mathcal{F}] = \int_{\mathbb{X}} f(x) \tilde{P}(dx)$ a.s. $[\mathbb{P}]$, for every $f \in M_b(\mathcal{X})$, where it is implied that $\int_{\mathbb{X}} f(x) \tilde{P}(dx) \in M_b(\mathcal{F})$. All conditional distributions both here and throughout the thesis are meant as regular versions.

Conditional identity in distribution and asymptotic exchangeability

A stochastic process $X = (X_n)_{n \geq 1}$ that is adapted to a filtration $\mathcal{F} = (\mathcal{F}_n)_{n \geq 0}$ is said to be *conditionally identically distributed* w.r.t. \mathcal{F} (or \mathcal{F} -c.i.d.) if it holds, for each $k, n \geq 1$ and any $f \in M_b(\mathcal{X})$, that

$$\mathbb{E}[f(X_{n+k})|\mathcal{F}_{n-1}] = \mathbb{E}[f(X_n)|\mathcal{F}_{n-1}].$$

If $\mathcal{F} \equiv \mathcal{F}^X$, then X is said to be simply c.i.d. As the above condition is equivalent to requiring the sequence of predictive distributions be a measure-valued martingale, then any \mathcal{F} -c.i.d. process is also c.i.d.; in fact, it is c.i.d. with respect to any coarser filtration, to which X is adapted. This kind of stochastic dependence is hinted in Kallenberg (1988) and introduced and studied in a systematic way by Berti et al. (2004). It turns out that any exchangeable sequence is conditionally identically distributed, whereas any stationary c.i.d. sequence is exchangeable. An important property of c.i.d. processes is that their predictive distributions converge in the sense of almost sure weak convergence to a random probability measure \tilde{P} on \mathbb{X} ,

$$\mathbb{P}(X_{n+1} \in \cdot | \mathcal{F}_n) \xrightarrow{w} \tilde{P}(\cdot) \quad \text{a.s.}[\mathbb{P}].$$

On c.s.m.s. spaces Berti et al. (2006, Theorem 2.2) show that this is equivalent to

$$\mathbb{E}[f(X_{n+1})|\mathcal{F}_n] \longrightarrow \int_{\mathbb{X}} f(x) \tilde{P}(dx) \quad \text{a.s.}[\mathbb{P}], \text{ for each } f \in C_b(\mathbb{X}),$$

which is an instance of almost sure conditional convergence (see below). Moreover, $\mathbb{P}(X_{n+1} \in B | \mathcal{F}_n) \xrightarrow{\text{a.s.}} \tilde{P}(B)$, for each $B \in \mathcal{X}$. On the other hand, c.i.d. processes satisfy a law of large numbers and have the additional property that

$$\frac{1}{n} \sum_{i=1}^n \delta_{X_i} \xrightarrow{w} \tilde{P} \quad \text{a.s.}[\mathbb{P}],$$

which implies that \tilde{P} is essentially $\sigma(X_1, X_2, \dots)$ -measurable. Now \tilde{P} , being a weak limit, is uniquely determined by the empirical distributions, and hence we will refer to it as the *directing measure* of X , similarly to the theory on exchangeable processes.

Note that $\mathbb{P}(X_{n+1} \in \cdot | \mathcal{F}_n)$ are not the actual predictive distributions, $\mathbb{P}(X_{n+1} \in \cdot | X_1, \dots, X_n)$, unless $\mathcal{F} \equiv \mathcal{F}^X$. However, a slightly generalized martingale convergence theorem (see Blackwell and Dubins, 1962, Theorem 2) implies for any \mathcal{F} -adapted process X , whose $\mathbb{P}(X_{n+1} \in \cdot | \mathcal{F}_n)$ and $\frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ converge to one and the same $\tilde{P} \in \mathbb{K}_P(\Omega, \mathbb{X})$, that

$$\mathbb{E}[f(X_{n+1})|X_1, \dots, X_n] = \mathbb{E}[\mathbb{E}[f(X_{n+1})|\mathcal{F}_n]|X_1, \dots, X_n] \longrightarrow \mathbb{E}[\tilde{P}_f|X_1, X_2, \dots] = \tilde{P}_f \quad \text{a.s.}[\mathbb{P}],$$

where $\tilde{P}_f = \int_{\mathbb{X}} f(x) \tilde{P}(dx)$ a.s. $[\mathbb{P}]$, for $f \in C_b(\mathbb{X})$. As a consequence,

$$\mathbb{P}(X_{n+1} \in \cdot | X_1, \dots, X_n) \xrightarrow{w} \tilde{P}(\cdot), \quad \text{a.s.}[\mathbb{P}].$$

Moreover, predictive convergence implies that X is *asymptotically exchangeable*, in the sense that

$$(X_{n+1}, X_{n+2}, \dots) \xrightarrow{d} (Z_1, Z_2, \dots),$$

for some exchangeable infinite sequence $(Z_n)_{n \geq 1}$ with directing measure \tilde{P} (see Aldous, 1985, Lemma 8.2).

Modes of convergence

Let $\mathcal{F} = (\mathcal{F}_n)_{n \geq 0}$ be a filtration and $Y = (Y_n)_{n \geq 1}$ be an \mathbb{X} -valued sequence that is *non-necessarily* adapted to \mathcal{F} . Then Y is said to converge in the sense of *almost sure conditional convergence* to some $K \in \mathbb{K}_P(\Omega, \mathbb{X})$ w.r.t. \mathcal{F} , denoted $Y_n \xrightarrow{a.s.cond.} K$, if

$$\mathbb{E}[f(Y_n)|\mathcal{F}_n] \longrightarrow \int_{\mathbb{X}} f(x)K(dx) \quad \text{a.s.}[\mathbb{P}], \text{ for each } f \in C_b(\mathbb{X}).$$

The first systematic study of a.s conditional convergence dates to Crimaldi (2009). If the above convergence is in probability, then Y is said to converge *stably in the strong sense* to K w.r.t. \mathcal{F} , denoted $Y_n \xrightarrow{s.stably} K$ (see Crimaldi et al., 2007, for a reference). *Stable convergence* of Y w.r.t. some sub- σ -field $\mathcal{G} \subseteq \mathcal{H}$, denoted $Y_n \xrightarrow{\mathcal{G}\text{-stably}} K$, is a weaker form of convergence, which requires

$$\mathbb{E}[V \cdot \mathbb{E}[f(Y_n)|\mathcal{G}]] \longrightarrow \int_{\Omega} V(\omega) \int_{\mathbb{X}} f(x)K(\omega, dx)\mathbb{P}(d\omega), \quad \text{for each } f \in C_b(\mathbb{X}) \text{ and } V \in L^1(\mathcal{H}).$$

It can be shown that $Y_n \xrightarrow{s.stably} K$ w.r.t. \mathcal{F} implies $Y_n \xrightarrow{\mathcal{F}_\infty\text{-stably}} K$. Moreover, if it holds $Y_n \xrightarrow{\mathcal{G}_1\text{-stably}} K$, for any sub- σ -fields $\mathcal{G}_1, \mathcal{G}_2 \subseteq \mathcal{H}$ such that $\mathcal{G}_2 \subseteq \mathcal{G}_1$, then $Y_n \xrightarrow{\mathcal{G}_2\text{-stably}} K$. In the case Y_n converges stably to K w.r.t. \mathcal{H} , we say simply that $Y_n \xrightarrow{stably} K$. The main application of stable convergence is in central limit theorems that allow for mixing variables in the limit. In fact, stable convergence can be seen as an intermediate form of convergence between convergence in distribution and convergence in probability as $Y_n \xrightarrow{stably} K$ implies $Y_n \xrightarrow{d} Z$, for some random variable $Z \sim \mathbb{E}[K(\cdot)]$, whereas it holds $Y_n \xrightarrow{p} Y$ if and only if $Y_n \xrightarrow{stably} \delta_Y$, for any random variable Y . In addition, stable convergence is transferrable in the sense that $Y_n \xrightarrow{stably} K$ and $d(X_n, Y_n) \xrightarrow{p} 0$ both imply $X_n \xrightarrow{stably} K$. For more information on stable convergence refer to Häusler and Luschgy (2015).

Note on measurability

A final point to make is that, whenever we have a sequence $(X_n)_{n \geq 1}$ of \mathbb{X} -valued random variables on $(\Omega, \mathcal{H}, \mathbb{P})$ that is adapted to a filtration $\mathcal{F} = (\mathcal{F}_n)_{n \geq 0}$, and a \mathbb{P} -a.s.-finite \mathcal{F} -stopping time T , then by X_T we would mean the random variable, defined by $X_T(\omega) := X_{T(\omega)}(\omega)$, for $\omega \in \{T < \infty\}$, and $X_T(\omega) := \vartheta$ otherwise, for some element $\vartheta \notin \mathbb{X}$. Then X_T is a map from Ω into $\mathbb{X}^* = \mathbb{X} \cup \{\vartheta\}$ that is $\mathcal{H} \setminus \mathcal{X}^*$ -measurable, where $\mathcal{X}^* = \sigma(\mathcal{X})$ on \mathbb{X}^* . Ideally, we would then extend each $f \in M(\mathcal{X})$ and $\nu \in \mathbb{M}_P(\mathbb{X})$ to \mathbb{X}^* by taking $f(\vartheta) = 0$, in which case the extended function would be guaranteed to be \mathcal{X}^* -measurable, and $\nu(B) = \nu(B \setminus \{\vartheta\})$ with $\nu(\{\vartheta\}) = 0$. As $T < \infty$ a.s. $[\mathbb{P}]$, as long as we work with identities and inequalities that hold \mathbb{P} -a.s., all of the above subtleties would be inconsequential, so we need not worry about ϑ , and we will consider X_T to be for all intents and purposes \mathbb{X} -valued.

Chapter II

Randomly reinforced Pólya sequence

2.1 Introduction

Random processes with reinforcement have long been a subject of interest to probabilists as surveyed by Pemantle (2007). The quintessential example of a stochastic process with reinforcement is the two-color Pólya urn (Pólya, 1930), which describes the composition of an urn containing balls of two different colors that are then sequentially sampled and reinforced with an additional ball of the observed color. Denote by $X_n \in \{0, 1\}$ the color of the ball drawn at step $n \geq 1$. The sampling procedure of a Pólya urn starts off with a Bernoulli pick, $X_1 \sim \text{Ber}(\frac{n_1}{n_0+n_1})$, and then proceeds according to

$$P(X_{n+1} = 1 | X_1, \dots, X_n) = \frac{n_1 + \sum_{i=1}^n X_i}{n_0 + n_1 + n},$$

where $n_0, n_1 \in \mathbb{N}$ represent the initial number of balls of each color. The predictive rule above postulates that the probability of picking a ball of color 1 is equal to the proportion of balls in the urn that have color 1. It follows that this system of predictive distributions characterizes an exchangeable sequence of random variables with a $\text{Beta}(n_0, n_1)$ prior distributions (see Freedman, 1965).

There has been an increasing interest from fields as diverse as design of clinical trials (Rosenberger, 2002), economics (Beggs, 2005), information science (Martin and Ho, 2002), etc. in generalizing the classical Pólya urn scheme along the lines of time-varying or random reinforcement. Examples include the time-dependent Pólya urn of Pemantle (1990), Durham et al. (1998)'s randomized Pólya urn, and the general class of immigrated urn models by Zhang et al. (2011). Most of these constructions suppose the existence of more than two, but a finite number of colors. However, in many applications either the number of colors is unknown, possibly infinite, or it may be that the set of colors varies from sample to sample.

The Pólya sequence of Blackwell and MacQueen (1973) is the conceptual extension of the above sampling scheme to the case of countable colors. Fix $\theta > 0$ and $\nu \in \mathbb{M}_P(\mathbb{X})$. A sequence of \mathbb{X} -valued random variables $(X_n)_{n \geq 1}$ is called a Pólya sequence with parameters (θ, ν) if $X_1 \sim \nu$ and the conditional distribution of

X_{n+1} given (X_1, \dots, X_n) is the transition probability kernel

$$\mathbb{P}(X_{n+1} \in \cdot | X_1, \dots, X_n) = \sum_{i=1}^n \frac{1}{\theta + n} \delta_{X_i}(\cdot) + \frac{\theta}{\theta + n} \nu(\cdot), \quad \text{for } n \geq 1.$$

Blackwell and MacQueen (1973) show that the Pólya sequence is exchangeable and has a Dirichlet process prior with parameters θ and ν . In addition, $P(X_i = X_j) > 0$, for $i \neq j$, so different observations may coincide. The probability measure ν is called the *base measure* of the sequence and, according to the colored Hoppe’s urn interpretation given in Chapter I, is used to “generate” new colors when the need occurs. It follows from the form of the predictive distributions that a researcher would find at stage $n + 1$ of the experiment a previously seen color X_k^* with probability $\sum_{i=1}^n \delta_{X_i}(\{X_k^*\})/(\theta + n)$, where $X_1^*, \dots, X_{L_n}^*$ denote the distinct colors among the already collected sample (X_1, \dots, X_n) , or a new color (i.e. pick a black ball from the corresponding Hoppe’s urn) with probability $\theta/(\theta + n)$. The above scheme represents the simplest case of a reinforcement mechanism as the probability of picking a previously seen (non-black) color is just the proportion of times that the same color has been observed in the past. Under this framework past observations are weighted equally, and thus our intent is to study the properties of the model, for which this condition is relaxed.

In this chapter of the thesis we define a large class of Pólya sequences with random reinforcement, called *randomly reinforced Pólya sequences* (RRPSs), through the introduction of a weighting scheme to the observation process. We then infer general conditions on the sequence of weights, under which a RRPS can be regarded as a *generalized species sampling sequence* (Bassetti et al., 2010) or as a *measure-valued Pólya urn process* (Janson, 2019). A RRPS that is a member of the first class of models has predictive distributions that form a measure-valued martingale, whereas as part of the latter they constitute a measure-valued Markov process. Each of the two model specifications has a later chapter devoted to it.

2.1.1 Model

Definition 2.1.1. *A sequence of \mathbb{X} -valued random variables $X = (X_n)_{n \geq 1}$ is called a randomly reinforced Pólya sequence (RRPS) if there exist a probability measure ν on \mathbb{X} such that $X_1 \sim \nu$, a constant $\theta > 0$, and a sequence of non-negative random variables $W = (W_n)_{n \geq 1}$ such that a version of the conditional distribution of X_{n+1} given $\mathcal{F}_n := \mathcal{F}_n^X \vee \mathcal{F}_n^W$ is the transition probability kernel*

$$\mathbb{P}(X_{n+1} \in \cdot | \mathcal{F}_n) = \sum_{i=1}^n \frac{W_i}{\theta + \sum_{j=1}^n W_j} \delta_{X_i}(\cdot) + \frac{\theta}{\theta + \sum_{j=1}^n W_j} \nu(\cdot), \quad \text{for } n \geq 1. \quad (\text{II.1})$$

Remark. It should be noted that the triplet $(\nu, \theta, (W_n))$ characterizing a RRPS is not identifiable as any transformation $(\nu, \gamma\theta, (\gamma W_n))$, for some $\gamma > 0$, would lead to the same predictive rules. Therefore, one has to fix either θ or W_1 in a statistical application.

The dynamics of a RRPS can be interpreted through a weighted version of the colored Hoppe’s urn, with the only difference being that non-black balls carry a weight (mass) of W_n . For simplicity, assume that the base measure ν is diffuse. Then the probability of picking any of the existing non-black balls from the urn,

labeled $k = 1, \dots, L_n$, is $\sum_{i: X_i = X_k^*} W_i / (\theta + \sum_{j=1}^n W_j)$, whereas the probability of picking a black ball is $\theta / (\theta + \sum_{j=1}^n W_j)$, in which case we would generate a new color $X_{L_{n+1}}^*$ from ν .

As anticipated, if $W_n = 1$ one recovers the Pólya sequence. In case the state space \mathbb{X} is finite, RRPSs constitute a particular subclass of generalized Pólya urn models (refer to Pemantle (2007) for definition) and as such have already garnered a lot of research interest (see Aletti et al., 2009; Berti et al., 2011; Crimaldi, 2009; May et al., 2005, among others). On the other hand, the current literature on RRPSs with infinite colors is, to the best of our knowledge, contained in Berti et al. (2009), Bassetti et al. (2010) and Fortini et al. (2018) and concerns only situations, in which the probabilistic model generates a conditionally identically distributed (c.i.d.) sequence. In fact, a RRPS is c.i.d., whenever the future weighting of the observations is done independently of what happens at the moment, conditionally given the past, that is

$$\mathbf{Assumption\ A.1.} \quad X_{n+1} \text{ is independent of } (W_{n+j})_{j \geq 1} \text{ given } \mathcal{F}_n, \text{ for each } n \geq 0. \quad (\text{A.1})$$

In Chapter III we formalize this statement and study the properties of the resulting process. In Chapter IV, in comparison, we examine a non-c.i.d. RRPS with convergent predictive distributions that can be regarded as an extension of the particular k -color urn scheme in Berti et al. (2010) to the case of countable colors. Before going into that direction, though, we investigate the relationship between RRPSs and two other classes of processes.

2.2 Generalized Ottawa sequences

Most of the existing extensions of the Pólya sequence are constructed in such a way as to preserve exchangeability and these include the two-parameter Poisson-Dirichlet process (see Pitman and Yor, 1997) or, more generally, the class of *species sampling sequences* of Pitman (1996) that are characterized by

$$\mathbb{P}(X_{n+1} \in \cdot | X_1, \dots, X_n) = \sum_{i=1}^n p_{n,i}(\Pi_n) \delta_{X_i}(\cdot) + r_n(\Pi_n) \nu(\cdot),$$

for some measurable functions $p_{n,i}$ and r_n from \mathcal{P}_n in $[0, 1]$, where \mathcal{P}_n denotes the set of all partitions on $\{1, \dots, n\}$ and Π_n is the random partition on $\{1, \dots, n\}$ that is generated by (X_1, \dots, X_n) . The fact that $p_{n,i}$ and r_n depend only on the partition is a necessary and sufficient condition for the sequence $(X_n)_{n \geq 1}$ to be exchangeable (see Hansen and Pitman, 2000, Theorem 1). However, this assumption may not be reasonable in situations, where the actual labels of the X_n 's matter, in which case knowledge of the partition is not enough. In addition, the clustering behavior of the generated sequence of observations cannot in general be described in terms of an urn sampling scheme.¹ Nonetheless, $p_{n,i}$ and r_n can be interpreted within a species sampling framework, with $p_{n,i}$ being the probability of observing at stage $n+1$ of the experiment the species of the i th subject that had previously appeared, given all that has happened up to time n , and r_n – the probability of discovering a new species, all provided ν is diffuse.

¹Regarding the Poisson-Dirichlet process with parameters (α, θ) , for $0 \leq \alpha < 1$ and $\theta > 0$, we can imagine a modification of Hoppe's urn, where a black pick would result in the addition of a black ball of mass α and a non-black ball of mass $1 - \alpha$, all else being equal.

In order to address some of the shortcomings of the basic model, Bassetti et al. (2010) define a notion of a generalized species sampling sequence by allowing the predictive distributions to further depend on the realizations of a latent process $(Y_n)_{n \geq 1}$, so that

$$\mathbb{P}(X_{n+1} \in \cdot | X_1, Y_1, \dots, X_n, Y_n) = \sum_{i=1}^n p_{n,i}(\Pi_n, Y_1, \dots, Y_n) \delta_{X_i}(\cdot) + r_n(\Pi_n, Y_1, \dots, Y_n) \nu(\cdot), \quad (\text{II.2})$$

for some measurable functions $p_{n,i}$ and r_n from $\mathcal{P}_n \times \mathbb{R}_+^n$ in $[0, 1]$. Then, unless $(X_n)_{n \geq 1}$ and $(Y_n)_{n \geq 1}$ are independent, the process is generally no longer exchangeable. In particular, Bassetti et al. (2010) call a sequence of \mathbb{X} -valued random variables $(X_n)_{n \geq 1}$ with predictive distributions as in (II.2) a *generalized Ottawa sequence* (GOS) if $p_{n,i}$ and r_n depend only on (Y_1, \dots, Y_n) and are such that $1 = r_0 \geq r_1 \geq r_2 \geq \dots > 0$,

$$p_{n,i}(y(n)) = \frac{r_n(y(n))}{r_{n-1}(y(n-1))} p_{n-1,i}(y(n-1)), \quad \text{for } i = 1, \dots, n-1,$$

$$p_{n,n}(y(n)) = 1 - \frac{r_n(y(n))}{r_{n-1}(y(n-1))},$$

for each $n \geq 1$ and every $y(n) = (y_1, \dots, y_n) \in \mathbb{R}_+^n$, and provided $(X_n, Y_n)_{n \geq 1}$ satisfies the condition that

$$X_{n+1} \text{ is independent of } (Y_{n+j})_{j \geq 1} \text{ given } \mathcal{G}_n, \text{ for each } n \geq 0,$$

where $\mathcal{G}_n := \mathcal{F}_n^X \vee \mathcal{F}_n^Y$. It turns out that any GOS is determined by the sequence $(r_n)_{n \geq 0}$ as we have from the above recursions that

$$\mathbb{P}(X_{n+1} \in \cdot | \mathcal{G}_n) = \sum_{i=1}^n r_n(Y_1, \dots, Y_n) \left(\frac{1}{r_i(Y_1, \dots, Y_i)} - \frac{1}{r_{i-1}(Y_1, \dots, Y_{i-1})} \right) \delta_{X_i}(\cdot) + r_n(Y_1, \dots, Y_n) \nu(\cdot).$$

The next result investigates the connection between the generalized Ottawa and the randomly reinforced Pólya sequences. It is immediate to see that the predictive distributions (II.1) of a RRPS with parameters $(\theta, \nu, (W_n))$ have the same structure as that of a GOS with parameters $(\nu, (r_n), (W_n))$, where

$$r_n(W_1, \dots, W_n) = \frac{\theta}{\theta + \sum_{j=1}^n W_j}.$$

More interesting is the fact that any GOS can be reformulated into a RRPS, provided one fixes θ beforehand. Such an equivalence relationship gives flexibility in modeling the probabilistic behavior of $(X_n)_{n \geq 1}$; for example, it may be more natural to make assumptions on the sequence of weights $(W_n)_{n \geq 1}$ rather than on $(Y_n)_{n \geq 1}$, which does not necessarily have any physical interpretation (Theorem 3.3.1 in Chapter III is one such example). On the other hand, some models such as the one in Airoidi et al. (2014) are best defined through the sequence $(r_n)_{n \geq 0}$, even though the GOS reparameterization loses the interpretation of a weighted Hoppe's urn.

Proposition 2.2.1. *Let $X = (X_n)_{n \geq 1}$ be a generalized Ottawa sequence with parameters ν and $(r_n)_{n \geq 0}$, and latent process $(Y_n)_{n \geq 1}$. Fix $\theta > 0$. Then there exists a sequence of non-negative random variables $(W_n)_{n \geq 1}$ such that X is a randomly reinforced Pólya sequence with parameters θ, ν and $(W_n)_{n \geq 1}$. Moreover, X satisfies assumption (A.1).*

Proof. Define

$$W_n := \theta \left(\frac{1}{r_n(Y_1, \dots, Y_n)} - \frac{1}{r_{n-1}(Y_1, \dots, Y_{n-1})} \right), \quad \text{and} \quad \mathcal{F}_n := \mathcal{F}_n^X \vee \mathcal{F}_n^W, \quad \text{for } n \geq 1.$$

It is immediate to see that W_n is \mathcal{F}_n^Y -measurable and non-negative as $r_{n-1} \geq r_n$ from the definition of a GOS. In addition, $r_n(Y_1, \dots, Y_n) = \theta / (\theta + \sum_{j=1}^n W_j)$, and thus

$$\mathbb{P}(X_{n+1} \in \cdot | \mathcal{G}_n) = \sum_{i=1}^n \frac{W_i}{\theta + \sum_{j=1}^n W_j} \delta_{X_i}(\cdot) + \frac{\theta}{\theta + \sum_{j=1}^n W_j} \nu(\cdot).$$

Let $f \in M_b(\mathcal{X})$. As $\mathcal{F}_n \subseteq \mathcal{G}_n$, it follows \mathbb{P} -a.s. that

$$\begin{aligned} \mathbb{E}[f(X_{n+1}) | \mathcal{F}_n] &= \mathbb{E}[\mathbb{E}[f(X_{n+1}) | \mathcal{G}_n] | \mathcal{F}_n] = \\ &= \mathbb{E} \left[\sum_{i=1}^n \frac{W_i}{\theta + \sum_{j=1}^n W_j} f(X_i) + \frac{\theta}{\theta + \sum_{j=1}^n W_j} \mathbb{E}[f(X_1)] \middle| \mathcal{F}_n \right] = \\ &= \sum_{i=1}^n \frac{W_i}{\theta + \sum_{j=1}^n W_j} f(X_i) + \frac{\theta}{\theta + \sum_{j=1}^n W_j} \mathbb{E}[f(X_1)] = \mathbb{E}[f(X_{n+1}) | \mathcal{G}_n]. \end{aligned}$$

Therefore,

$$\mathbb{P}(X_{n+1} \in \cdot | \mathcal{F}_n) = \sum_{i=1}^n \frac{W_i}{\theta + \sum_{j=1}^n W_j} \delta_{X_i}(\cdot) + \frac{\theta}{\theta + \sum_{j=1}^n W_j} \nu(\cdot).$$

Regarding the last part of the proposition, denote $\mathcal{F}_n^{W,c} := \sigma(W_{n+1}, W_{n+2}, \dots)$ and $\mathcal{G}_n^{Y,c} := \sigma(Y_{n+1}, Y_{n+2}, \dots)$, for $n \geq 1$. Let $m \geq 1$, $V \in M_b(\mathcal{F}_n)$ and $h_1, \dots, h_m \in M_b(\mathcal{B}(\mathbb{R}_+))$. Then

$$\begin{aligned} \mathbb{E}[V \cdot \mathbb{E}[h_1(W_{n+1}) \cdots h_m(W_{n+m}) f(X_{n+1}) | \mathcal{F}_n]] &= \mathbb{E}[V h_1(W_{n+1}) \cdots h_m(W_{n+m}) f(X_{n+1})] = \\ &= \mathbb{E}[V h_1(W_{n+1}) \cdots h_m(W_{n+m}) \mathbb{E}[f(X_{n+1}) | \mathcal{F}_n \vee \mathcal{F}_n^{W,c}]] = \\ &= \mathbb{E} \left[V h_1(W_{n+1}) \cdots h_m(W_{n+m}) \mathbb{E}[\mathbb{E}[f(X_{n+1}) | \mathcal{G}_n \vee \mathcal{G}_n^{Y,c}] | \mathcal{F}_n \vee \mathcal{F}_n^{W,c}] \right] = \\ &= \mathbb{E} \left[V h_1(W_{n+1}) \cdots h_m(W_{n+m}) \mathbb{E}[\mathbb{E}[f(X_{n+1}) | \mathcal{G}_n] | \mathcal{F}_n \vee \mathcal{F}_n^{W,c}] \right] = \\ &= \mathbb{E} \left[V h_1(W_{n+1}) \cdots h_m(W_{n+m}) \mathbb{E}[\mathbb{E}[f(X_{n+1}) | \mathcal{F}_n] | \mathcal{F}_n \vee \mathcal{F}_n^{W,c}] \right] = \\ &= \mathbb{E}[V h_1(W_{n+1}) \cdots h_m(W_{n+m}) \mathbb{E}[f(X_{n+1}) | \mathcal{F}_n]] = \\ &= \mathbb{E}[V \cdot \mathbb{E}[h_1(W_{n+1}) \cdots h_m(W_{n+m}) | \mathcal{F}_n] \mathbb{E}[f(X_{n+1}) | \mathcal{F}_n]], \end{aligned}$$

where we have used conditional determinism and $\mathbb{E}[f(X_{n+1}) | \mathcal{G}_n] = \mathbb{E}[f(X_{n+1}) | \mathcal{F}_n]$ from before. But m , V , f and h_1, \dots, h_m are arbitrary, so X satisfies (A.1). \square

Bassetti et al. (2010, Section 5) show that any GOS is c.i.d. with respect to $(\mathcal{G}_n^*)_{n \geq 0}$, where $\mathcal{G}_n^* := \mathcal{F}_n^X \vee \mathcal{F}_\infty^Y$, so it follows from Proposition 2.2.1 that a RRPS satisfying assumption (A.1) is c.i.d. with respect to $(\mathcal{F}_n^*)_{n \geq 0}$, for $\mathcal{F}_n^* := \mathcal{F}_n^X \vee \mathcal{F}_\infty^W$. In Chapter III we study the large n limiting behavior of the latter process and provide results that either generalize the existing ones from Bassetti et al. (2010) or are completely new as regards to the class of RRPSs/GOSs.

2.3 Measure-valued Pólya urn processes

In this section we consider the predictive distributions (II.1) of a RRPS as measure-valued random variables and establish a connection between the probabilistic behavior of the resulting measure-valued process and the sequence $(W_n)_{n \geq 1}$. To that end, define $\mu_0 := \theta\nu$ and

$$\mu_n := \theta\nu + \sum_{i=1}^n W_i \delta_{X_i}, \quad \text{for } n \geq 1, \quad (\text{II.3})$$

which forms a sequence of finite random measures on \mathbb{X} . In case ν is diffuse, $\mu_n(B)$ denotes the total mass of the balls, whose colors lie in B , for $B \in \mathcal{X}$, with respect to the weighted Hoppe's urn interpretation. Note that

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

so μ_n is the sum of the last term in the sequence with another finite random measure that captures the reinforcement mechanism. The above structure is inherent in all urn schemes and species sampling sequences, and yet Bandyopadhyay and Thacker (2014, 2016, 2017) and Mailler and Marckert (2017) are arguably the first to concentrate their study on the measure-valued random sequence $(\mu_n)_{n \geq 0}$. In fact, their models are instances of the general class of Pólya urn processes proposed by Janson (2018, 2019), for which we lay the groundwork next.

Let $\mu \in \mathbb{M}_F^*(\mathbb{X})$. Denote by $\mu' = \mu/\mu(\mathbb{X})$. As the composition of an urn is described by such a finite measure μ , we consider a kernel $R \in \mathbb{K}_F(\mathbb{X}, \mathbb{X})$, which maps colors $x \mapsto R_x$ to finite measures, to model the *replacement rule* of the corresponding urn scheme. Reinforcements will be in general random, so we assume the existence of a probability kernel $\mathcal{R} \in \mathbb{K}_P(\mathbb{X}, \mathbb{M}_F(\mathbb{X}))$ such that $R_x \sim \mathcal{R}_x$, for $x \in \mathbb{X}$. The function $\phi_\mu : \mathbb{X} \rightarrow \mathbb{M}_F^*(\mathbb{X})$, defined by

$$\phi_\mu(x) := \mu + R_x, \quad \text{for } x \in \mathbb{X},$$

records the updated urn composition after μ has been reinforced with R_x , given that a ball of color x has been drawn with distribution μ' . Moreover, $\phi_\mu(x) \sim \Phi_\mu(x) := \mathcal{R}_x \circ \psi_\mu^{-1}$, where $\psi_\mu : \mathbb{M}_F(\mathbb{X}) \mapsto \mathbb{M}_F^*(\mathbb{X})$ is the map $\nu \mapsto \nu + \mu$, and thus $\phi_\mu(x) = \psi_\mu(R_x)$. In fact, $\Phi_\mu \in \mathbb{K}_P(\mathbb{X}, \mathbb{M}_F^*(\mathbb{X}))$ (see Janson, 2019). The unconditional distribution of the updated composition is the probability measure that results from averaging Φ_μ out across possible observations,

$$\hat{R}_\mu(\cdot) := \int_{\mathbb{X}} \Phi_\mu(x)(\cdot) \mu'(dx).$$

It follows that the function \hat{R} that maps $\mu \mapsto \hat{R}(\mu) = \hat{R}_\mu$ is a probability kernel from $\mathbb{M}_F^*(\mathbb{X})$ to $\mathbb{M}_F^*(\mathbb{X})$ (Janson, 2019, Lemma 3.3). In case R is deterministic, i.e. $\mathcal{R}_x = \delta_{R_x}$, one has $\Phi_\mu(x) = \delta_{\phi_\mu(x)}$, and thus $\hat{R}_\mu(\cdot) = \mu' \circ \phi_\mu^{-1}(\cdot)$.

Formally, Janson (2019) calls a sequence of finite random measures $(\mu_n)_{n \geq 0}$ a *measure-valued Pólya urn process* with a random replacement rule $\mathcal{R} \in \mathbb{K}_P(\mathbb{X}, \mathbb{M}_F(\mathbb{X}))$, provided

$$\mu_n = \mu_{n-1} + R_{X_n}, \quad \text{for } n \geq 1,$$

where

$$\mathbb{P}(X_n \in \cdot | \mu_0, \dots, \mu_{n-1}) = \mathbb{P}(X_n \in \cdot | \mu_{n-1}) = \mu'_{n-1}(\cdot),$$

denotes the color of the ball that has been drawn at stage n , given the urn composition μ_{n-1} , and

$$\mathbb{P}(R_{X_n} \in \cdot | \mu_0, \dots, \mu_{n-1}, X_1, \dots, X_n) = \mathcal{R}_{X_n}.$$

The above equation implies, in particular, that the reinforcement R_{X_n} is independent of the accumulated information $(\mu_0, \mu_1, X_1, \dots, \mu_{n-1}, X_{n-1})$ given X_n . As a consequence,

$$\begin{aligned} \mathbb{P}(\mu_n \in \cdot | \mu_0, \dots, \mu_{n-1}) &= \mathbb{E}[\mathbb{P}(\psi_{\mu_{n-1}}(R_{X_n}) \in \cdot | \mu_0, \dots, \mu_{n-1}, X_n) | \mu_0, \dots, \mu_{n-1}] = \\ &= \mathbb{E}[\mathbb{P}(\psi_{\mu_{n-1}}(R_{X_n}) \in \cdot | \mu_{n-1}, X_n) | \mu_0, \dots, \mu_{n-1}] = \int_{\mathbb{X}} \mathcal{R}_x \circ \psi_{\mu_{n-1}}^{-1}(\cdot) \mu'_{n-1}(dx) = \hat{R}_{\mu_{n-1}}(\cdot); \end{aligned}$$

thus, $(\mu_n)_{n \geq 0}$ is a Markov process with initial state μ_0 and transition kernel \hat{R} . The next result states that each measure-valued Pólya urn process with random R can be represented as a measure-valued Pólya urn process with a deterministic replacement rule that is defined on a larger color space.

Theorem 2.3.1 (Janson 2019, Theorem 1.3). *Let \mathbb{X} be c.s.m.s., and $(\mu_n)_{n \geq 0}$ be a measure-valued Pólya urn process with initial state $\mu_0 \in \mathbb{M}_F^*(\mathbb{X})$ and random replacement rule $\mathcal{R} \in \mathbb{K}_P(\mathbb{X}, \mathbb{M}_F(\mathbb{X}))$. Then there exists a measure-valued Pólya urn process $(\bar{\mu}_n)_{n \geq 0}$ on $\mathbb{X} \times [0, 1]$ with deterministic replacement rule such that*

$$\bar{\mu}_n = \mu_n \otimes \lambda, \quad \text{for } n \geq 0,$$

where λ denotes the Lebesgue measure on $[0, 1]$.

Let $(\mu_n)_{n \geq 0}$ be a measure-valued Pólya process. We will derive the predictive distributions of $(X_n)_{n \geq 1}$ next. Suppose first that $(\mu_n)_{n \geq 0}$ has a deterministic replacement rule $R \in \mathbb{K}_P(\mathbb{X}, \mathbb{X})$. Define $\tilde{\mu}_0 := \mu_0$ and

$$\tilde{\mu}_n(x_1, \dots, x_n; B) := \frac{\mu_0(B) + \sum_{i=1}^n R_{x_i}(B)}{\mu_0(\mathbb{X}) + \sum_{i=1}^n R_{x_i}(\mathbb{X})}, \quad \text{for } x_1, \dots, x_n \in \mathbb{X}, B \in \mathcal{X}, n \geq 1.$$

Then $\tilde{\mu}_n \in \mathbb{K}_P(\mathbb{X}^n, \mathbb{X})$, for $n \geq 1$. As $\mu'_n(\cdot) = \tilde{\mu}(X_1, \dots, X_n; \cdot)$, it follows that

$$\mathbb{P}(X_{n+1} \in \cdot | X_1, \dots, X_n) = \mu'_n(\cdot).$$

Suppose $(\mu_n)_{n \geq 0}$ is a measure-valued Pólya urn process with random replacement rule $\mathcal{R} \in \mathbb{K}_P(\mathbb{X}, \mathbb{M}_F(\mathbb{X}))$. From Lemma 3.22 in Kallenberg (2002), there exists a measurable function $f : \mathbb{X} \times [0, 1] \rightarrow \mathbb{M}_F(\mathbb{X})$ such that $f(x, U) \sim \mathcal{R}_x$, for $x \in \mathbb{X}$. As a consequence, the function $\bar{R} : \mathbb{X} \times [0, 1] \rightarrow \mathbb{M}_F(\mathbb{X} \times [0, 1])$, defined by

$$\bar{R}_{x,u} := f(x, u) \otimes \lambda, \quad \text{for } x \in \mathbb{X}, u \in [0, 1],$$

is a transition kernel from $\mathbb{X} \times [0, 1]$ to $\mathbb{X} \times [0, 1]$. Let $(\bar{\mu}_n)_{n \geq 0}$ be a measure-valued Pólya process with initial state $\bar{\mu}_0 = \mu_0 \otimes \lambda$ and (deterministic) replacement rule \bar{R} . Then

$$\bar{\mu}_n = \bar{\mu}_{n-1} + \bar{R}_{X_n, U_n}, \quad \text{for } n \geq 1,$$

where $(X_1, U_1) \sim \bar{\mu}_0$ and

$$\mathbb{P}((X_{n+1}, U_{n+1}) \in \cdot | \bar{\mu}_0, \dots, \bar{\mu}_n) = \bar{\mu}'_n(\cdot).$$

Suppose $\bar{\mu}_k = \mu_k \otimes \lambda$, for some $k \geq 1$. Then $\sigma(\bar{\mu}_0, \dots, \bar{\mu}_k) \subseteq \sigma(\mu_0, \dots, \mu_k)$ and

$$\mathbb{P}((X_{k+1}, U_{k+1}) \in \cdot | \mu_0, \dots, \mu_k) = \mu'_k \otimes \lambda(\cdot).$$

It follows that $U_{k+1} \sim \text{Unif}[0, 1]$ is independent of $(\mu_0, X_1, \mu_1, \dots, X_k, \mu_k, X_{k+1})$. By Lemma A.6 in the Appendix,

$$\begin{aligned} \mathbb{P}(f(X_{k+1}, U_{k+1}) \in \cdot | \mu_0, \dots, \mu_k, X_1, \dots, X_{k+1}) &= \mathbb{P}(f(X_{k+1}, U_{k+1}) \in \cdot | X_{k+1}) = \\ &= \int_{[0,1]} f(X_{k+1}, u) \lambda(du) = \mathbb{P} \circ (f(X_{k+1}, U))^{-1}(\cdot) = \mathcal{R}_{X_{k+1}}(\cdot). \end{aligned}$$

As a result, we can couple the two measure-valued Pólya processes and take $\mu_{k+1} = \mu_k + f(X_{k+1}, U_{k+1})$, which implies

$$\bar{\mu}_{k+1} = \bar{\mu}_k + \bar{R}_{X_k, U_k} = (\mu_k + f(X_k, U_k)) \otimes \lambda = \mu_{k+1} \otimes \lambda.$$

By induction, $\bar{\mu}_n = \mu_n \otimes \lambda$, for $n \geq 1$, so $\mu_n = \mu_{n-1} + f(X_n, U_n)$ and

$$\mathbb{P}(X_{n+1} \in \cdot | X_1, U_1, \dots, X_n, U_n) = \frac{\mu_0(\cdot) + \sum_{i=1}^n f(X_i, U_i)(\cdot)}{\mu_0(\mathbb{X}) + \sum_{i=1}^n f(X_i, U_i)(\mathbb{X})},$$

where $(U_n)_{n \geq 1}$ is a sequence of i.i.d. $\text{Unif}[0, 1]$ random variables such that U_n is independent of (X_1, \dots, X_n) .

The generalized Pólya urn model of Athreya and Ney (1972) studies the situation, where the state space \mathbb{X} is finite, in which case the urn composition μ_n becomes an atomic measure and R is best described in terms of a matrix with random elements. When the matrix has zero off-diagonal elements, i.e. we add only balls of the observed color, R collapses to a random atomic measure with just one atom and the resulting urn scheme is better known as a randomly reinforced urn (see Muliere et al., 2006). On the other hand, the Pólya sequence of Blackwell and MacQueen (1973) on general \mathbb{X} corresponds to a measure-valued Pólya urn process with the deterministic replacement rule $R_x = \delta_x$. As a consequence, RRPSs can simultaneously be regarded as extensions of both randomly reinforced urns and Pólya sequences through a weighted replacement rule of the form $R_x = W(x) \cdot \delta_x$, where $W(x)$ is a non-negative random variable, for $x \in \mathbb{X}$, which ensures that the observed color is reinforced with an additional random number of balls. We provide below necessary and sufficient conditions for the composition measures (II.3) of a RRPS to be a measure-valued Pólya process. In particular, this result concerns the specification of the weights $(W_n)_{n \geq 1}$ and motivates our modeling choices in Chapter IV. As a prerequisite, let $\xi_x : \mathbb{R}_+ \rightarrow \mathbb{M}_F(\mathbb{X})$ be the map $w \mapsto w\delta_x$. Then ξ_x is $\mathcal{B}(\mathbb{R}_+) \setminus \mathcal{M}_F(\mathbb{X})$ -measurable and such that $\xi_x(\mathbb{R}_+) = \{\mu \in \mathbb{M}_F(\mathbb{X}) : \mu = w\delta_x, w \in \mathbb{R}_+\}$.

Proposition 2.3.2. *Let \mathbb{X} be c.s.m.s., and $(\mu_n)_{n \geq 0}$ be a measure-valued Pólya urn process with initial state $\mu_0 \in \mathbb{M}_F^*(\mathbb{X})$ and random replacement rule $\mathcal{R} \in \mathbb{K}_P(\mathbb{X}, \mathbb{M}_F(\mathbb{X}))$ such that $\mathcal{R}_x = \eta_x \circ \xi_x^{-1}$, for all $x \in \mathbb{X}$ and some $(\eta_x)_{x \in \mathbb{X}} \subseteq \mathbb{M}_P(\mathcal{B}(\mathbb{R}_+))$. Then $(\mu_n)_{n \geq 0}$ generates the probability law of a randomly reinforced Pólya sequence $(X_n)_{n \geq 1}$, satisfying*

$$\mathbb{P}(X_{n+1} \in \cdot | X_1, W_1, \dots, X_n, W_n) = \mu'_n(\cdot),$$

where $W_n = h(X_n, U_n)$, for some $h \in M(\mathcal{X} \otimes \mathcal{B}[0, 1])$, and $(U_n)_{n \geq 1}$ is a sequence of i.i.d. $\text{Unif}[0, 1]$ random variables such that U_n is independent of (X_1, \dots, X_n) .

Conversely, if $(X_n)_{n \geq 1}$ is a randomly reinforced Pólya sequence with parameters θ, ν and $(W_n)_{n \geq 1}$ such that $W_n = h(X_n, U_n)$, for some $h \in M_+(\mathcal{X} \otimes \mathcal{B}[0, 1])$, where $(U_n)_{n \geq 1}$ are i.i.d. $\text{Unif}[0, 1]$ and U_n is independent of (X_1, \dots, X_n) , then the sequence of random measures $(\mu_n)_{n \geq 0}$, defined by $\mu_0 := \theta\nu$ and

$$\mu_n := \theta\nu + \sum_{i=1}^n W_i \delta_{X_i}, \quad \text{for } n \geq 1,$$

is a measure-valued Pólya urn process with random replacement rule $\mathcal{R}_x = \eta_x \circ \xi_x^{-1}$, where $\eta_x = \mathbb{P} \circ h(x, U)^{-1}$, for $x \in \mathbb{X}$ and $U \sim \text{Unif}[0, 1]$.

Proof. Part I. Suppose $(\mu_n)_{n \geq 0}$ is the prescribed measure-valued Pólya urn process. From Lemma 3.22 in Kallenberg (2002), there exists a measurable function $h : \mathbb{X} \times [0, 1] \rightarrow \mathbb{R}_+$ such that $h(x, U) \sim \eta_x$, for $x \in \mathbb{X}$, where $U \sim \text{Unif}[0, 1]$. Then $h(x, U)\delta_x \sim \mathcal{R}_x$. Following the discussion prior to this proposition, $\mu_1 = \mu_0 + h(X_1, U_1)$ with $(X_1, U_1) \sim \mu_0 \otimes \lambda$ and

$$\mu_{n+1} = \mu_n + h(X_{n+1}, U_{n+1})\delta_{X_{n+1}}, \quad \text{for } n \geq 1,$$

where

$$\mathbb{P}((X_{n+1}, U_{n+1}) \in \cdot | X_1, U_1, \dots, X_n, U_n) = \mu_n \otimes \lambda,$$

with λ the Lebesgue measure on $[0, 1]$. Moreover, $(U_n)_{n \geq 1}$ is a sequence of i.i.d. $\text{Unif}[0, 1]$ random variables such that U_n is independent of X_1, \dots, X_n . Define $W_n := h(X_n, U_n)$, for $n \geq 1$. Then $W_n \in M_+(\mathcal{H})$. From $\mu_n = \mu_0 + \sum_{i=1}^n W_i \delta_{X_i}$, it follows that μ_n is measurable with respect to $\sigma(X_1, W_1, \dots, X_n, W_n)$, and hence

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

As a result, $(X_n)_{n \geq 1}$ is a RRPS with parameters $\theta = \mu_0(\mathbb{X})$, $\nu = \mu'_0$ and $(W_n)_{n \geq 1}$.

Part II. Let $(X_n)_{n \geq 1}$ be a RRPS such that $W_n = h(X_n, U_n)$, for $h \in M_+(\mathcal{X} \otimes \mathcal{B}[0, 1])$. Define $\mu_0 := \theta\nu$,

$$\mu_n := \theta\nu + \sum_{i=1}^n W_i \delta_{X_i}, \quad \text{and} \quad \mathcal{R}_x := (\mathbb{P} \circ h(x, U)^{-1}) \circ \xi_x^{-1}, \quad \text{for } x \in \mathbb{X},$$

where $U \sim \text{Unif}[0, 1]$. From Lemma 2.6 in Janson (2019), it holds $\mathcal{R} \in \mathbb{K}_P(\mathbb{X}, \mathbb{M}_F(\mathbb{X}))$ if and only if, for every $h \in M_{b,+}(\mathcal{X})$, the map

$$x \mapsto \int_{\mathbb{M}_F(\mathbb{X})} e^{-\int_{\mathbb{X}} g(x)\nu(dx)} \mathcal{R}_x(d\nu) = \int_{\mathbb{R}_+} e^{-cg(x)} \mathbb{P} \circ h(x, U)^{-1}(dc) = \int_{[0,1]} e^{-h(x,u)g(x)} \lambda(du),$$

from \mathbb{X} to \mathbb{R}_+ is measurable. In fact, the latter is true by Proposition 6.9 in (Cinlar, 2011, Chapter I). On the other hand, $(\mu_n)_{n \geq 0}$ is a sequence of finite random measures such that $\mu_n = \mu_{n-1} + h(X_n, U_n)\delta_{X_n}$. It follows that μ_n is $\sigma(X_1, W_1, \dots, X_n, W_n)$ -measurable, so

$$\mathbb{P}(X_n \in \cdot | \mu_0, \dots, \mu_{n-1}) = \mathbb{E}[\mathbb{P}(X_n \in \cdot | X_1, W_1, \dots, X_{n-1}, W_{n-1}) | \mu_0, \dots, \mu_{n-1}] = \mu'_{n-1}(\cdot).$$

Moreover, by Lemma A.6 in the Appendix,

$$\begin{aligned} \mathbb{P}(h(X_n, U_n)\delta_{X_n} \in \cdot | \mu_0, \dots, \mu_{n-1}, X_1, \dots, X_n) &= \mathbb{P}(h(X_n, U_n)\delta_{X_n} \in \cdot | X_n) = \\ &= \int_{[0,1]} h(X_n, u)\delta_{X_n} \lambda(du) = \mathcal{R}_{X_n}(\cdot); \end{aligned}$$

therefore, $(\mu_n)_{n \geq 0}$ is a measure-valued Pólya urn process with initial state $\theta\nu$ and replacement rule \mathcal{R} .

□

Note that the RRPSs from the above proposition do not satisfy assumption (A.1) in general. However, the particular model with $W_n = \bar{h}(U_n)$, for some $\bar{h} \in M(\mathcal{B}[0, 1])$, does, in which case the RRPS generates a c.i.d. sequence of random variables and the urn composition forms a measure-valued Pólya urn process. A necessary and sufficient condition is given by $\eta_x = \eta$, for some $\eta \in \mathbb{M}_P(\mathbb{X})$. The RRPS studied by Berti et al. (2009) and the one given by Bassetti et al. (2010) as an example of a GOS are such that $W_n = \bar{h}(U_n)$, though, in the latter paper the authors require $(U_n)_{n \geq 1}$ and $(X_n)_{n \geq 1}$ to be completely independent.

A remark on measure-valued Pólya urn processes that we do not explore further in this thesis concerns the fact that the representation

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

could potentially allow one to develop novel prior distributions for Bayesian nonparametric inference such that the random limit of the predictive distributions is \mathbb{P} -a.s. absolutely continuous. Such constructions would require that R_x is itself absolutely continuous, which is a step beyond what existing urn models imply.

Chapter III

Conditionally identically distributed randomly reinforced Pólya sequence

3.1 Introduction

In this chapter of the thesis we study a particular subclass of randomly reinforced Pólya sequences (RRPS), whose predictive distributions form a measure-valued martingale, in which case the sequence of observations becomes conditionally identically distributed (c.i.d.). The earliest examples of non-stationary c.i.d. RRPSs date to Pemantle (1990) and Li et al. (1996), although they deal with the finite state space. Berti et al. (2004, Example 1.3) are arguably the first to suggest a construction of this type and to state explicitly its c.i.d. property. Further developments of their model can be found in May et al. (2005) and Crimaldi (2009). On the other hand, Berti et al. (2009, Section 4) and Fortini et al. (2018, Example 3.8) cover some aspects of the infinite-state case, whereas Bassetti et al. (2010, Example 5.3) consider a c.i.d. RRPS with independent weights as part of a more general family of models, called generalized species sampling sequences, and in particular as an example of the subfamily of *generalized Ottawa sequences* (GOSs). In Chapter II we showed that c.i.d. RRPSs and GOSs are, in fact, different reparameterizations of the same model.

Our contributions to the study of these processes concern the properties of the random limit of the predictive distributions (II.1). The main result of this chapter, which is proven in Theorem 3.2.1, states that any c.i.d. RRPS behaves asymptotically as a species sampling sequence (Pitman, 1996) such that the number of individuals from each species in the population is non-negligible (Proposition 3.3.3). In Theorem 3.3.1 we provide central limit results for c.i.d. RRPSs in terms of stable and a.s. conditional convergence. As these are centered on the random limit measure, we show in Section 3.4 how to use them for approximate posterior inference, both when the reinforcement weights are observable and in the case when they are latent to the model. In Theorem 3.2.3 we recover the exact distribution of the predictive limit for the particular RRPS with binary weights.

3.1.1 Model

Let $(X_n)_{n \geq 1}$ be a RRPS with parameters θ , ν and $(W_n)_{n \geq 1}$. In Chapter II we introduced the following conditional independence assumption on the pair $(X_n, W_n)_{n \geq 1}$, which was part of the definition of a GOS,

$$\textbf{Assumption A.1.} \quad X_{n+1} \text{ is independent of } (W_{n+j})_{j \geq 1} \text{ given } \mathcal{F}_n, \text{ for each } n \geq 0; \quad (\text{A.1})$$

where the filtration $\mathcal{F} = (\mathcal{F}_n)_{n \geq 0}$ is defined by $\mathcal{F}_n := \mathcal{F}_n^X \vee \mathcal{F}_n^W$, for $n \geq 1$, and \mathcal{F}_0 is the trivial σ -algebra. It follows that any RRPS satisfying assumption (A.1) is c.i.d. with respect to the filtration $\mathcal{F}^* = (\mathcal{F}_n^*)_{n \geq 0}$, given by $\mathcal{F}_n^* := \mathcal{F}_n^X \vee \mathcal{F}_\infty^W$, where $\mathcal{F}_\infty^W := \bigvee_{n \in \mathbb{N}} \mathcal{F}_n^W$. In some of the results, however, we would need the weaker condition that

$$\textbf{Assumption A.2.} \quad X_{n+1} \text{ is independent of } W_{n+1} \text{ given } \mathcal{F}_n, \text{ for each } n \geq 0; \quad (\text{A.2})$$

thus, according to (A.2), weights and observations are contemporaneously independent conditionally on all past information, in which case the RRPS becomes c.i.d. with respect to \mathcal{F} . The two results are contained in the proposition below, whose proof can be found by adapting Example 3.8 in Fortini et al. (2018) to the univariate case.

Proposition 3.1.1. *Let $X = (X_n)_{n \geq 1}$ be a randomly reinforced Pólya sequence with parameters θ , ν and $(W_n)_{n \geq 1}$. If X satisfies assumption (A.1), then X is \mathcal{F}^* -c.i.d. If X satisfies assumption (A.2), then X is \mathcal{F} -c.i.d.*

3.2 Characterization of directing measure

Recall from Chapter I that the predictive distributions of any c.i.d. process $X = (X_n)_{n \geq 1}$ converge in the sense of a.s. weak convergence to some random probability measure \tilde{P} on \mathbb{X} ,

$$\mathbb{P}(X_{n+1} \in \cdot | X_1, \dots, X_n) \xrightarrow{w} \tilde{P}(\cdot) \quad \text{a.s.}[\mathbb{P}],$$

implying that X is asymptotically exchangeable with directing measure \tilde{P} . As a consequence, knowledge of the distribution of \tilde{P} would inform us on the eventual behavior of the sequence of observations and provide a basis for Bayesian inference. In this section we give a general description of the random limit measure of a c.i.d. RRPS and complete probabilistic characterization in the case $W_n \in \{0, 1\}$. We also state some known facts about the random partition.

3.2.1 Almost sure discreteness

The following theorem describes the directing measure of any RRPS that satisfies assumption (A.1) as a mixture of the base measure ν and a discrete measure with atoms at the distinct values of X , listed in order

of appearance. Similar distributional results can be found in Pitman (1996) regarding species sampling sequences and in Berti et al. (2019) for a class of c.i.d. models with a recursive structure. The proof of our result uses arguments from Berti et al. (2013), Berti et al. (2019) and Krasker and Pratt (1986).

We will adopt the language of species sampling from Pitman (1996) (see also Chapter II, Section 2.2) in order to specify the sequence of distinct observations as they appear. For that purpose, define the sequence of random variables $(T_n)_{n \geq 1}$ by $T_1 := 1$ and

$$T_n := \inf\{m \in \mathbb{N} : m > T_{n-1}, X_m \notin \{X_1, \dots, X_{m-1}\}\}, \quad \text{for } n \geq 2,$$

that mark the time, at which a new species has been discovered. The actual observations at the discovery times are captured by the process $(X_k^*)_{k \geq 1}$, defined by

$$X_k^*(\omega) := \begin{cases} X_{T_k(\omega)}(\omega), & \omega \in \{T_k < \infty\}; \\ \vartheta, & \text{otherwise} \end{cases}, \quad \text{for } \omega \in \Omega \text{ and } k \geq 1.$$

for some $\vartheta \notin \mathbb{X}$. The sequence $(L_n)_{n \geq 1}$, given by $L_1 := 1$ and

$$L_n := \max\{k \in \{1, \dots, n\} : X_k \notin \{X_1, \dots, X_{k-1}\}\}, \quad \text{for } n \geq 2,$$

is used to count the number of distinct species at each stage n of the experiment. Lastly, we denote by

$$P_n(\cdot) := \mathbb{P}(X_{n+1} \in \cdot | \mathcal{F}_n), \quad \text{for } n \geq 1.$$

Theorem 3.2.1. *Let $X = (X_n)_{n \geq 1}$ be a randomly reinforced Pólya sequence with parameters θ , ν and $(W_n)_{n \geq 1}$. Suppose X satisfies assumption (A.1). If ν is non-atomic, then*

$$\|P_n - \tilde{P}\| := \sup_{B \in \mathcal{X}} |P_n(B) - \tilde{P}(B)| \longrightarrow 0 \quad \text{a.s.}[\mathbb{P}], \quad (\text{III.1})$$

for some random probability measure \tilde{P} on \mathbb{X} that is of the form

$$\tilde{P} = \sum_k p_k^* \delta_{X_k^*} + \left(1 - \sum_k p_k^*\right) \nu,$$

where

$$p_k^* = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\{X_k^*\}) \quad \text{a.s.}[\mathbb{P}],$$

and $(X_k^*)_{k \geq 1}$ are i.i.d. (ν) conditionally given $(p_k^*)_{k \geq 1}$.

Remark.

- (a) Non-atomicity of ν ensures that already observed species are not reinforced further through ν . Otherwise, the structure of the random partition becomes more involved and the form of \tilde{P} is not necessarily the one from above (see Bassetti et al., 2018).

(b) Predictive and empirical convergence both imply

$$\frac{1}{\sum_{j=1}^n W_j} \sum_{i=1}^n W_i \delta_{X_i}(\{X_k^*\}) \approx \frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\{X_k^*\}),$$

so that the relative frequency and the relative weighted frequency approach each other as $n \rightarrow \infty$.

(c) If $(W_n)_{n \geq 1}$ are i.i.d. with $W_n \leq \beta < \infty$, then one has from Proposition 3.3.3 that

$$\mathbb{P}(p_k^* = p) = 0, \quad \text{for all } p \in [0, 1].$$

As a consequence, $p_k^* \neq p_l^*$ a.s.[\mathbb{P}], for $k \neq l$, and $p_k^* > 0$ a.s.[\mathbb{P}], i.e. all discovered species survive in the end. The latter fact is not true in general and a counterexample is given implicitly in Pitman (1996, Section 3).

Proof of Theorem 3.2.1.

Part I: Discreteness. The sequence $(L_n)_{n \geq 1}$ is increasing, so the limit $L_\infty := \lim_{n \rightarrow \infty} L_n$ exists in $\bar{\mathbb{R}}$. Then $T_k < \infty$ if and only if $k \leq L_\infty$. Indeed, if $T_k < \infty$, then $k \leq L_{T_k} \leq L_\infty$. Conversely, for $L_\infty < \infty$, one has $k \leq L_\infty < L_n + \epsilon$ for some $\epsilon \in (0, 1)$ and $n \geq 1$, whereas under $L_\infty = \infty$, there exists $n \geq 1$ such that $k \leq L_n$; in both cases $T_k < \infty$. Define

$$p_{n,k}^* := \frac{\sum_{i \in \Pi_{n,k}} W_i}{\theta + \sum_{j=1}^n W_j} \cdot \mathbb{1}_{\{k \leq L_n\}}, \quad \text{and} \quad \theta_n := \frac{\theta}{\theta + \sum_{j=1}^n W_j}, \quad \text{for } n, k \geq 1,$$

where the sets $\Pi_{n,k} = \{j \in \{1, \dots, n\} : X_j = X_k^*\}$, for $n \geq 1$ and $k = 1, \dots, L_n$, denote the random partition of the first n observations according to their species. It follows that

$$P_n(\cdot) = \sum_{k=1}^{L_\infty} p_{n,k}^* \delta_{X_k^*}(\cdot) + \theta_n \nu(\cdot).$$

As X is \mathcal{F} -c.i.d. from Proposition 3.1.1, then $P_n(B) \xrightarrow{\text{a.s.}} \tilde{P}(B)$, for some $\tilde{P} \in \mathbb{K}_P(\Omega, \mathbb{X})$ and all $B \in \mathcal{X}$. Furthermore, $\mathbb{E}[P_n(B) | X_1, \dots, X_n] = \mathbb{E}[\tilde{P}(B) | X_1, \dots, X_n]$ (see Chapter I). Define

$$Q(\cdot) := \frac{1}{1 - \theta_\infty} (\tilde{P}(\cdot) - \theta_\infty \nu(\cdot)), \quad \text{and} \quad \mu(\cdot) := \mathbb{E}[Q(\cdot)],$$

where $\theta_\infty = \lim_{n \rightarrow \infty} \theta_n$. Then $Q(\mathbb{X}) = 1$, $Q(\emptyset) = 0$ and $\omega \mapsto Q(B)(\omega)$ is \mathcal{H} -measurable, for each $B \in \mathcal{X}$. In addition, $Q(B_1 \cup B_2) = Q(B_1) + Q(B_2)$, for disjoint $B_1, B_2 \in \mathcal{X}$, and $Q(B) \geq 0$ a.s.[\mathbb{P}] from predictive convergence. It follows that the map $B \mapsto \mu(B)$ is a probability measure on \mathbb{X} . Indeed, let $\{B_n\}_{n \in \mathbb{N}} \subseteq \mathcal{X}$ be disjoint. Then

$$\mu\left(\bigcup_{n=1}^{\infty} B_n\right) = \mathbb{E}\left[\frac{1}{1 - \theta_\infty} \sum_{n=1}^{\infty} (\tilde{P}(B_n) - \theta_\infty \nu(B_n))\right] = \sum_{n=1}^{\infty} \mu(B_n),$$

which is justified by dominated convergence. By a slight variant of Theorem 3.1 in Ghosal and van der Vaart (2017), there exists $\tilde{Q} \in \mathbb{K}_P(\Omega, \mathbb{X})$ such that $\tilde{Q}(B) = Q(B)$ a.s.[\mathbb{P}], for all $B \in \mathcal{X}$. As \mathbb{X} is separable, then $\mathcal{X} = \sigma(\mathcal{D})$, for some countable π -class $\mathcal{D} \subseteq \mathcal{X}$. Let $\Omega_0 \in \mathcal{X}$ be such that $\mathbb{P}(\Omega_0) = 1$ and $\tilde{P}(B)(\omega) =$

$(1 - \theta_\infty(\omega))\tilde{Q}(B)(\omega) + \theta_\infty(\omega)\nu(B)$, for $B \in \mathcal{D}$ and $\omega \in \Omega_0$. From standard results in measure theory, $\tilde{P} = \tilde{Q} + \theta_\infty\nu$ a.s.[\mathbb{P}]. Define

$$\begin{aligned}\tilde{\mu}_{1,n}(\cdot) &:= (\delta_{X_1, \dots, X_n} \otimes \delta_{X_{n+1}})(\cdot), & \tilde{\mu}_{2,n}(\cdot) &:= (\delta_{X_1, \dots, X_n} \otimes \tilde{P})(\cdot), \\ \mu_{1,n}(\cdot) &:= \mathbb{E}[\tilde{\mu}_{1,n}(\cdot)], & \mu_{2,n}(\cdot) &:= \mathbb{E}[\tilde{\mu}_{2,n}(\cdot)],\end{aligned}$$

and

$$H_n := \{(x_1, \dots, x_n, x_{n+1}) \in \mathbb{X}^{n+1} : x_{n+1} = x_i, \text{ for some } i \leq n\},$$

for $n \geq 1$. Then $\tilde{\mu}_{1,n}, \tilde{\mu}_{2,n} \in \mathbb{K}_P(\Omega, \mathbb{X}^{n+1})$, $\mu_{1,n}, \mu_{2,n} \in \mathbb{M}_P(\mathbb{X}^{n+1})$ and $H_n \in \mathcal{X}^{n+1}$, so that

$$\tilde{\mu}_{1,n}(H_n) = \delta_{X_{n+1}}(\{X_1, \dots, X_n\}), \quad \tilde{\mu}_{2,n}(H_n) = \tilde{P}(\{X_1, \dots, X_n\}).$$

Let $A \in \mathcal{X}^n$ and $B \in \mathcal{X}$. It follows that

$$\begin{aligned}\mu_{1,n}(A \times B) &= \mathbb{P}((X_1, \dots, X_n) \in A, X_{n+1} \in B) = \mathbb{E}[\mathbb{1}_A(X_1, \dots, X_n)\mathbb{P}(X_{n+1} \in B | X_1, \dots, X_n)] = \\ &= \mathbb{E}[\mathbb{1}_A(X_1, \dots, X_n)\mathbb{P}(\tilde{P}(B) | X_1, \dots, X_n)] = \mathbb{E}[\mathbb{1}_A(X_1, \dots, X_n)\tilde{P}(B)] = \mu_{2,n}(A \times B);\end{aligned}$$

thus, $\mu_{1,n} = \mu_{2,n}$ on \mathbb{X}^{n+1} . Define $\tilde{Q}^*(\cdot) := (1 - \theta_\infty)\tilde{Q}(\cdot)$ and the random sets $S := \bigcup_{k=1}^\infty \{X_k^*\} \setminus \{\vartheta\}$ and $S_n := \{X_1, \dots, X_n\}$, for $n \geq 1$. Then $\tilde{Q}^* \in \mathbb{K}_F(\Omega, \mathbb{X})$. As $S_n(\omega) \uparrow S(\omega)$, for $\omega \in \Omega$, one has from continuity and monotone convergence that

$$\begin{aligned}\mathbb{E}[\tilde{Q}^*(S)] &= \lim_{n \rightarrow \infty} \mathbb{E}[\tilde{Q}^*(S_n)] = \lim_{n \rightarrow \infty} \mu_{2,n}(H_n) = \\ &= \lim_{n \rightarrow \infty} \mathbb{P}(X_{n+1} \in \{X_1, \dots, X_n\}) = 1 - \lim_{n \rightarrow \infty} \mathbb{E}[\theta_n] = 1 - \mathbb{E}[\theta_\infty].\end{aligned}$$

But $\tilde{Q}^*(\mathbb{X}) = 1 - \theta_\infty$, so $\mathbb{E}[\tilde{Q}^*(\mathbb{X}) - \tilde{Q}^*(S)] = 0$ and $\tilde{Q}^*(\mathbb{X}) = \tilde{Q}^*(S)$ a.s.[\mathbb{P}]. Therefore, $\tilde{Q}^* = \sum_{k=1}^{L_\infty} p_k^* \delta_{X_k^*}$ a.s.[\mathbb{P}], for some $(p_k^*)_k \subseteq M_+(\mathcal{H})$ with $0 \leq p_k^* \leq 1$, and thus

$$\tilde{P} = \sum_{k=1}^{L_\infty} p_k^* \delta_{X_k^*} + \theta_\infty \nu \quad \text{a.s.}[\mathbb{P}],$$

say, on $\Omega_1 \in \mathcal{X}$ with $\mathbb{P}(\Omega_1) = 1$. Denote by $P_{n,d}(\cdot) := \sum_{k=1}^{L_\infty} p_{n,k}^* \delta_{X_k^*}(\cdot)$ and $\tilde{P}_d(\cdot) := \sum_{k=1}^{L_\infty} p_k^* \delta_{X_k^*}(\cdot)$, so that

$$P_n(\cdot) = P_{n,d}(\cdot) + \theta_n \nu(\cdot), \quad \text{and} \quad \tilde{P}(\cdot) = \tilde{P}_d(\cdot) + \theta_\infty \nu(\cdot) \quad \text{a.s.}[\mathbb{P}].$$

Let $\omega \in \Omega_1$. As ν is diffuse, one has $\tilde{P}_d(S(\omega))(\omega) = \tilde{P}(S(\omega))(\omega) = 1 - \theta_\infty(\omega)$ and $P_{n,d}(S(\omega))(\omega) = P_n(S(\omega))(\omega) = 1 - \theta_n(\omega)$, for $n \geq 1$; thus,

$$P_{n,d}(\mathbb{X})(\omega) = P_{n,d}(S(\omega))(\omega) \longrightarrow \tilde{P}_d(S(\omega))(\omega) = \tilde{P}_d(\mathbb{X})(\omega).$$

Let $k \geq 1$. By Lemma A.9 in the Appendix and as ν is diffuse,

$$\mathbb{P}(p_{n,k}^* = P_n(\{X_k^*\}) \rightarrow \tilde{P}(\{X_k^*\}) = p_k^* \mathbb{1}\{k \leq L_\infty\}) = 1. \quad (\text{III.2})$$

As a consequence,

$$\mathbb{P}\left(\{\omega \in \Omega : P_{n,d}(\{x\})(\omega) \rightarrow \tilde{P}_d(\{x\})(\omega), \text{ for all } x \in \mathbb{X}\}\right) =$$

$$\begin{aligned}
&= \mathbb{P}\left(\left\{\omega \in \Omega : P_{n,d}(\{x\})(\omega) \rightarrow \tilde{P}_d(\{x\})(\omega), \text{ for all } x \in S(\omega)\right\}\right) = \\
&= \mathbb{P}\left(\left\{\omega \in \Omega : p_{n,k}^*(\omega) \rightarrow p_k^*(\omega), \text{ for all } k \leq L_\infty(\omega)\right\}\right) = \\
&= \mathbb{P}\left(\left\{\omega \in \Omega : p_{n,k}^*(\omega) \cdot \mathbb{1}_{\{k \leq L_\infty\}}(\omega) \rightarrow p_k^*(\omega) \cdot \mathbb{1}_{\{k \leq L_\infty\}}(\omega), \text{ for all } k \geq 1\right\}\right) = \\
&= \mathbb{P}\left(\bigcap_{k=1}^{\infty} \left\{p_{n,k}^* \mathbb{1}_{\{k \leq L_\infty\}} \rightarrow p_k^* \mathbb{1}_{\{k \leq L_\infty\}}\right\}\right) \geq \\
&\geq 1 - \sum_{k=1}^{\infty} \mathbb{P}(p_{n,k}^* \cdot \mathbb{1}_{\{k \leq L_\infty\}} \not\rightarrow p_k^* \cdot \mathbb{1}_{\{k \leq L_\infty\}}) = \\
&= 1 - \sum_{k=1}^{\infty} \mathbb{P}(p_{n,k}^* \not\rightarrow p_k^* | \{k \leq L_\infty\}) \mathbb{P}(k \leq L_\infty) = 1,
\end{aligned}$$

say, on $\Omega_2 \in \mathcal{X}$ with $\mathbb{P}(\Omega_2) = 1$. Then $\mathbb{P}(\Omega_1 \cap \Omega_2) = 1$. Let $\omega \in \Omega_1 \cap \Omega_2$. It follows that

$$\|P_n(\omega) - \tilde{P}(\omega)\| \leq \sup_{x \in S(\omega)} |P_{n,d}(\{x\})(\omega) - \tilde{P}_d(\{x\})(\omega)| + \|\theta_n(\omega)\nu - \theta_\infty(\omega)\nu\| \rightarrow 0.$$

Finally, let $k \leq L_\infty$. By Lemma A.9 in the Appendix, one has that

$$\frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\{X_k^*\}) \rightarrow \tilde{P}(\{X_k^*\}) = p_k^* \quad \text{a.s.}[\mathbb{P}],$$

in the sense of (III.2).

Part II: Independence. Regarding the last part of the proof, let $B \in \mathcal{X}$ and $k \leq L_\infty$. It follows that

$$\begin{aligned}
\mathbb{P}(X_k^* \in B | X_1^*, \dots, X_{k-1}^*) &= \mathbb{E}[\mathbb{P}(X_k^* \in B | \mathcal{F}_{T_{k-1}}) | X_1^*, \dots, X_{k-1}^*] = \\
&= \mathbb{E}\left[\sum_{m=1}^{\infty} \mathbb{P}(X_{T_k} \in B | \mathcal{F}_{T_{k-1}}) \mathbb{1}_{\{T_{k-1}=m\}} \middle| X_1^*, \dots, X_{k-1}^*\right] = \\
&= \mathbb{E}\left[\sum_{m=1}^{\infty} \mathbb{P}(X_{T_k} \in B \cap \{X_1, \dots, X_{T_{k-1}}\}^c | \mathcal{F}_{T_{k-1}}) \mathbb{1}_{\{T_{k-1}=m\}} \middle| X_1^*, \dots, X_{k-1}^*\right] = \\
&= \mathbb{E}\left[\sum_{m=1}^{\infty} \mathbb{P}(X_{m+1} \in B \cap \{X_1, \dots, X_m\}^c | \mathcal{F}_m) \mathbb{1}_{\{T_{k-1}=m\}} \middle| X_1^*, \dots, X_{k-1}^*\right] = \\
&= \mathbb{E}\left[\sum_{m=1}^{\infty} \theta_m \nu(B) \cdot \mathbb{1}_{\{T_{k-1}=m\}} \middle| X_1^*, \dots, X_{k-1}^*\right] = \\
&= \nu(B) \cdot \mathbb{E}\left[\sum_{m=1}^{\infty} \mathbb{P}(X_{m+1} \notin \{X_1, \dots, X_m\} | \mathcal{F}_m) \mathbb{1}_{\{T_{k-1}=m\}} \middle| X_1^*, \dots, X_{k-1}^*\right] = \nu(B),
\end{aligned}$$

so $(X_k^*)_{k=1}^{L_\infty}$ is a sequence of i.i.d. (ν) random variables such that X_k^* is independent of $\mathcal{F}_{T_{k-1}}$. Let $\Pi = (\Pi_n)_{n \geq 1}$ be the random partition on \mathbb{N} generated by X , where $\Pi_n = \{\Pi_{n,1}, \dots, \Pi_{n,L_n}\}$ denotes the partition of $\{1, \dots, n\}$ by (X_1, \dots, X_n) into L_n clusters and takes values in the set \mathcal{P}_n . It follows that each p_k^* , being a weak limit, is a function of Π , so in order to have $(X_k^*)_{k \geq 1}$ and $(p_k^*)_{k \geq 1}$ independent it suffices to show that

$(X_k^*)_{k \geq 1}$ and $(\Pi_n)_{n \geq 1}$ are independent, for which it is enough¹ that

$$\mathbb{P}(\Pi_n = \pi_n, X_1^* \in B_1, \dots, X_{l_n}^* \in B_{l_n}) = \mathbb{P}(\Pi_n = \pi_n) \mathbb{P}(X_1^* \in B_1, \dots, X_{l_n}^* \in B_{l_n}), \quad (\text{III.3})$$

for each $n \geq 1$, $B_1, \dots, B_{l_n} \in \mathcal{X}$ and $\pi_n = \{\pi_{n,1}, \dots, \pi_{n,l_n}\} \in \mathcal{P}_n$, where l_n denotes the length of π_n . Let $c : \mathbb{N} \rightarrow \mathbb{N}$ be the mapping that shows the cluster membership of each X_n , i.e. $c(n) = k$ if and only if $X_n = X_k^*$. Denote by its "inverse" function, $c^{-1} : \mathbb{N} \rightarrow \mathbb{N}$, the first time that a particular distinct value has been observed, that is $c^{-1}(k) = \min\{m \in \{1, \dots, k\} : X_m = X_k^*\}$. The independence argument then proceeds in the following way,

$$\begin{aligned} \mathbb{P}(\Pi_n = \pi_n, X_1^* \in B_1, \dots, X_{l_n}^* \in B_{l_n}) &= \\ &= \mathbb{P}(X_1^* \in B_1, \dots, X_{l_n}^* \in B_{l_n}, X_1 = X_{c(1)}^*, \dots, X_{c^{-1}(2)} = X_2^*, \dots, X_{c^{-1}(l_n)} = X_{l_n}^*, \dots, X_n = X_{c(n)}^*) = \\ &= \mathbb{E}\left[\mathbb{1}_{\{X_1^* \in B_1, \dots, X_{l_n}^* \in B_{l_n}, X_1 = X_{c(1)}^*, \dots, X_{n-1} = X_{c(n-1)}^*\}} \mathbb{P}(X_n = X_{c(n)}^* | \mathcal{F}_{n-1})\right] = \\ &= \mathbb{E}\left[\mathbb{1}_{\{X_1^* \in B_1, \dots, X_{l_n}^* \in B_{l_n}, X_1 = X_{c(1)}^*, \dots, X_{n-1} = X_{c(n-1)}^*\}} \frac{\sum_{i \in \pi_{n-1, c(n)}} W_i}{\theta + \sum_{j=1}^{n-1} W_j}\right] = \\ &= \mathbb{E}\left[\mathbb{1}_{\{X_1^* \in B_1, \dots, X_{l_n}^* \in B_{l_n}, X_1 = X_{c(1)}^*, \dots, X_{n-2} = X_{c(n-2)}^*\}} \mathbb{E}\left[\frac{\sum_{i \in \pi_{n-1, c(n)}} W_i}{\theta + \sum_{j=1}^{n-1} W_j} \mathbb{1}_{\{X_{n-1} = X_{c(n-1)}^*\}} \middle| \mathcal{F}_{n-2}\right]\right] = \\ &= \mathbb{E}\left[\mathbb{1}_{\{X_1^* \in B_1, \dots, X_{l_n}^* \in B_{l_n}, X_1 = X_{c(1)}^*, \dots, X_{n-2} = X_{c(n-2)}^*\}} \mathbb{E}\left[\frac{\sum_{i \in \pi_{n-1, c(n)}} W_i}{\theta + \sum_{j=1}^{n-1} W_j} \middle| \mathcal{F}_{n-2}\right] \times \right. \\ &\quad \left. \times \mathbb{P}(X_{n-1} = X_{c(n-1)}^* | \mathcal{F}_{n-2})\right] = \\ &= \mathbb{E}\left[\mathbb{1}_{\{X_1^* \in B_1, \dots, X_{l_n}^* \in B_{l_n}, X_1 = X_{c(1)}^*, \dots, X_{n-2} = X_{c(n-2)}^*\}} \frac{\sum_{i \in \pi_{n-2, c(n-1)}} W_i}{\theta + \sum_{j=1}^{n-2} W_j} \times \right. \\ &\quad \left. \times \mathbb{E}\left[\frac{\sum_{i \in \pi_{n-1, c(n)}} W_i}{\theta + \sum_{j=1}^{n-1} W_j} \middle| \mathcal{F}_{n-2}\right]\right] = \\ &= \mathbb{E}\left[\mathbb{1}_{\{X_1^* \in B_1, \dots, X_{l_n}^* \in B_{l_n}, X_1 = X_{c(1)}^*, \dots, X_{n-2} = X_{c(n-2)}^*\}} \frac{\sum_{i \in \pi_{n-2, c(n-1)}} W_i}{\theta + \sum_{j=1}^{n-2} W_j} \frac{\sum_{i \in \pi_{n-1, c(n)}} W_i}{\theta + \sum_{j=1}^{n-1} W_j}\right] = \dots = \\ &= \mathbb{E}\left[\mathbb{1}_{\{X_1^* \in B_1, \dots, X_{l_n}^* \in B_{l_n}, X_1 = X_{c(1)}^*, \dots, X_{c^{-1}(l_n)} = X_{l_n}^*\}} \frac{\sum_{i \in \pi_{c^{-1}(l_n), c(c^{-1}(l_n)+1)}} W_i}{\theta + \sum_{j=1}^{c^{-1}(l_n)} W_j} \dots \frac{\sum_{i \in \pi_{n-1, c(n)}} W_i}{\theta + \sum_{j=1}^{n-1} W_j}\right] = \end{aligned}$$

¹Independence between $\bigcup_{n=1}^{\infty} \sigma(\Pi_1, \dots, \Pi_n)$ and $\bigcup_{n=1}^{\infty} \sigma(X_1^*, \dots, X_n^*)$ follows from (III.3) since, for any $j \geq 1$,

$$\begin{aligned} \mathbb{P}(\Pi_n = \pi_n, X_1^* \in B_1, \dots, X_{l_n}^* \in B_{l_n}, X_{l_n+1}^* \in B_{l_n+1}, \dots, X_{l_n+j}^* \in B_{l_n+j}) &= \\ &= \mathbb{P}(\Pi_n = \pi_n, X_1^* \in B_1, \dots, X_{l_n}^* \in B_{l_n}) \mathbb{P}(X_{l_n+1}^* \in B_{l_n+1}, \dots, X_{l_n+j}^* \in B_{l_n+j}), \end{aligned}$$

from before, whereas it follows for $i = 1, \dots, l_n - 1$ that

$$\mathbb{P}(\Pi_n = \pi_n, X_1^* \in B_1, \dots, X_{l_n-i}^* \in B_{l_n-i}) = \mathbb{P}(\Pi_n = \pi_n, X_1^* \in B_1, \dots, X_{l_n-i}^* \in B_{l_n-i}, X_{l_n-i+1}^* \in \mathbb{X}, \dots, X_{l_n}^* \in \mathbb{X}).$$

$$\begin{aligned}
&= \mathbb{E} \left[\mathbb{1}_{\{X_1^* \in B_1, \dots, X_{l_n-1}^* \in B_{l_n-1}, X_1 = X_{c(1)}^*, \dots, X_{c^{-1}(l_n)-1} = X_{c^{-1}(l_n)-1}^*\}} \mathbb{E} \left[\frac{\sum_{i \in \pi_{c^{-1}(l_n), c(c^{-1}(l_n)+1)}} W_i}{\theta + \sum_{j=1}^{c^{-1}(l_n)} W_j} \dots \right. \right. \\
&\quad \left. \left. \dots \frac{\sum_{i \in \pi_{n-1, c(n)}} W_i}{\theta + \sum_{j=1}^{n-1} W_j} \mathbb{1}_{\{X_{l_n}^* \in B_{l_n}, X_{c^{-1}(l_n)} = X_{l_n}^*\}} \middle| \mathcal{F}_{l_n-1} \right] \right] = \\
&= \mathbb{E} \left[\mathbb{1}_{\{X_1^* \in B_1, \dots, X_{l_n-1}^* \in B_{l_n-1}, X_1 = X_{c(1)}^*, \dots, X_{c^{-1}(l_n)-1} = X_{c^{-1}(l_n)-1}^*\}} \mathbb{E} \left[\frac{\sum_{i \in \pi_{c^{-1}(l_n), c(c^{-1}(l_n)+1)}} W_i}{\theta + \sum_{j=1}^{c^{-1}(l_n)} W_j} \dots \right. \right. \\
&\quad \left. \left. \dots \frac{\sum_{i \in \pi_{n-1, c(n)}} W_i}{\theta + \sum_{j=1}^{n-1} W_j} \middle| \mathcal{F}_{l_n-1} \right] \mathbb{P}(X_{c^{-1}(l_n)} \in B_{l_n} \cap \{X_1, \dots, X_{c^{-1}(l_n)-1}\}^c \middle| \mathcal{F}_{l_n-1}) \right] = \\
&= \mathbb{E} \left[\mathbb{1}_{\{X_1^* \in B_1, \dots, X_{l_n-1}^* \in B_{l_n-1}, X_1 = X_{c(1)}^*, \dots, X_{c^{-1}(l_n)-1} = X_{c^{-1}(l_n)-1}^*\}} \frac{\theta}{\theta + \sum_{j=1}^{c^{-1}(l_n)-1} W_j} \times \right. \\
&\quad \left. \times \frac{\sum_{i \in \pi_{c^{-1}(l_n), c(c^{-1}(l_n)+1)}} W_i}{\theta + \sum_{j=1}^{c^{-1}(l_n)} W_j} \dots \frac{\sum_{i \in \pi_{n-1, c(n)}} W_i}{\theta + \sum_{j=1}^{n-1} W_j} \right] \nu(B_{l_n}) = \dots = \\
&= \mathbb{E} \left[\mathbb{1} \cdot \frac{\sum_{i \in \pi_{1, c(2)}} W_i}{\theta + W_1} \dots \frac{\theta}{\theta + \sum_{j=1}^{c^{-1}(2)-1} W_j} \dots \frac{\theta}{\theta + \sum_{j=1}^{c^{-1}(l_n)-1} W_j} \dots \right. \\
&\quad \left. \dots \frac{\sum_{i \in \pi_{n-1, c(n)}} W_i}{\theta + \sum_{j=1}^{n-1} W_j} \right] \nu(B_1) \dots \nu(B_{l_n}) = \\
&= \mathbb{P}(\Pi_n = \pi_n) \prod_{k=1}^{l_n} \mathbb{P}(X_k^* \in B_k),
\end{aligned}$$

where $\pi_m = \pi_n \cap \{1, \dots, m\}$, for $m \leq n$, and we set $\sum_{i \in \pi_{l, k}} W_i / (\theta + \sum_{j=1}^l W_j) \equiv 1$, whenever $k > l$. Note that we have used assumption (A.1) repeatedly to split conditional expectations along the lines of conditional independence. \square

It follows that the RRPS from Theorem 3.2.1 is asymptotically exchangeable with directing measure the random limit of its predictive distributions. What is more, any random probability measure \tilde{P} that is structured as in (III.1) defines a species sampling sequence through de Finetti's representation theorem (see Pitman, 1996, Section 3); therefore, any RRPS satisfying assumption (A.1) is asymptotically equivalent in law to a species sampling model. This result is directly extendable to the larger class of generalized species sampling sequences (see Chapter II for definition) that observe a condition similar to (A.1).

Denote by

$$\theta_n := \frac{\theta}{\theta + \sum_{j=1}^n W_j},$$

the conditional probability of observing a new species at stage n of the experiment. Following Pitman (1996), we say that the model in Theorem 3.2.1 is *proper* if \tilde{P} is \mathbb{P} -a.s. discrete, i.e. $\sum_k p_k^* = 1$ a.s. [P], which is itself true if and only if $\theta_n \xrightarrow{\text{a.s.}} 0$. In that case, the a.s. discreteness of \tilde{P} can be deduced from Theorem 17 in Berti et al. (2019) without the part on the independence between atoms and masses. On the other hand,

it holds

$$\mathbb{P}(L_{n+1} = L_n + 1 | \mathcal{F}_n) = P_n(\{X_1, \dots, X_n\}^c) = \theta_n;$$

thus, we can derive discreteness conditions from the behavior of L_n .

Proposition 3.2.2. *Under the conditions of Theorem 3.2.1,*

$$\sum_k p_k^* = 1 \quad \text{a.s.}[\mathbb{P}], \quad \text{if and only if} \quad \frac{L_n}{n} \longrightarrow 0 \quad \text{a.s.}[\mathbb{P}].$$

Proof. Suppose $\sum_k p_k^* = 1$ a.s. $[\mathbb{P}]$. Define $U_n := L_n - L_{n-1}$, for $n \geq 1$, where $L_0 = 0$. Then $L_n = \sum_{i=1}^n U_i$ and $\theta_n = \mathbb{E}[U_{n+1} | \mathcal{F}_n]$ with $\theta_n \xrightarrow{\text{a.s.}} 0$ by hypothesis. As $U_n \in \{0, 1\}$, one has $\sum_{n=1}^{\infty} \mathbb{E}[U_n^2] / n^2 < \infty$, so by Lemma A.7 in the Appendix,

$$\frac{L_n}{n} = \frac{1}{n} \sum_{i=1}^n U_i \longrightarrow 0 \quad \text{a.s.}[\mathbb{P}].$$

Conversely, suppose $\frac{L_n}{n} \xrightarrow{\text{a.s.}} 0$. Since $L_n \leq n$, one has by the dominated convergence theorem that

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E}[\theta_{i-1}] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}[U_i] = \frac{1}{n} \mathbb{E}[L_n] \longrightarrow 0 \quad \text{a.s.}[\mathbb{P}],$$

where $\theta_0 = 1$. On the other hand,

$$\mathbb{E}[\theta_{n+1} | \mathcal{F}_n] = \mathbb{E} \left[\frac{\theta}{\theta + \sum_{j=1}^{n+1} W_j} \middle| \mathcal{F}_n \right] \leq \frac{\theta}{\theta + \sum_{j=1}^n W_j} = \theta_n;$$

thus, $(\theta_n)_{n \geq 0}$ is a bounded positive \mathcal{F} -supermartingale. From the martingale convergence theorem, $\theta_n \rightarrow \theta_\infty$ a.s. $[\mathbb{P}]$ and in L^1 , for some non-negative integrable random variable θ_∞ . In particular,

$$\mathbb{E}[\theta_n] \longrightarrow \mathbb{E}[\theta_\infty].$$

Suppose, by contradiction, $\mathbb{E}[\theta_\infty] > 0$. By Lemma A.3 in the Appendix,

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E}[\theta_{i-1}] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\theta_i] + \frac{1}{n} - \frac{\mathbb{E}[\theta_n]}{n} \longrightarrow \mathbb{E}[\theta_\infty] > 0,$$

absurd. As a consequence, $\mathbb{E}[\theta_\infty] = 0$, which implies that $\theta_\infty = \theta / (\theta + \sum_{n=1}^{\infty} W_n) = 0$ a.s. $[\mathbb{P}]$, and hence $\sum_{k=1}^{L_\infty} p_k^* = 1$ a.s. $[\mathbb{P}]$. \square

If $\mathbb{P}(L_{n+1} = L_n + 1 | \mathcal{F}_n)$ goes to zero at a fast enough rate that

$$\sum_{n=1}^{\infty} \frac{\theta}{\theta + \sum_{j=1}^n W_j} < \infty \quad \text{a.s.}[\mathbb{P}],$$

then one has from Proposition 2.1 in Bassetti et al. (2010) that

$$\mathbb{P}(L_{n+1} = L_n \text{ ult.}) = 1,$$

so that the number of distinct species L_n converges almost surely to some random variable L_∞ such that $\mathbb{P}(L_\infty < \infty) = 1$. Therefore, it is possible to come up with models that characterize a.s. discrete random probability measures having a finite number of atoms.

3.2.2 Random partition

We make a slight detour in order to state some facts about the random partition, which are derived from the existing theory on GOSs. To that end, let ν be diffuse. It follows from Airoidi et al. (2014, Appendix B) that the moment-generating function of L_{n+1} is given by

$$\mathbb{E}[e^{-tL_{n+1}}] = e^{-t} \sum_{m=0}^n (e^{-t} - 1)^m \phi_{n,m}, \quad \text{for } t \in \mathbb{R},$$

where $\phi_{n,0} := 1$ and $\phi_{n,m} := \sum_{1 \leq l_1 < l_2 < \dots < l_m \leq n} \mathbb{E}[\theta_{l_1} \dots \theta_{l_m}]$, for $n \geq 1$. Next, suppose that there exist $(h_n)_{n \geq 1} \subseteq \mathbb{R}_+$ such that $h_n \uparrow \infty$,

$$\sum_{n=1}^{\infty} \frac{1}{h_n^2} \mathbb{E}[\theta_n(1 - \theta_n)] < \infty, \quad \text{and} \quad \frac{1}{h_n} \sum_{i=1}^n \theta_{i-1} \longrightarrow L_{\infty} \quad \text{a.s.}[\mathbb{P}],$$

for some $L_{\infty} \in M_+(\mathcal{H})$, where $\theta_0 = 1$. Then Bassetti et al. (2010, Theorem 5.2 of technical report) show that

$$\frac{L_n}{h_n} \longrightarrow L_{\infty} \quad \text{a.s.}[\mathbb{P}];$$

thus, RRPSs have a wide range of clustering behavior, depending on the specification of the weights (see Bassetti et al., 2010, Example 5.4 and 5.5). If $(h_n)_{n \geq 1}$ are in fact such that

$$\Lambda_n := \frac{1}{h_n} \sum_{i=1}^n \theta_{i-1}(1 - \theta_{i-1}) \longrightarrow \Lambda \quad \text{a.s.}[\mathbb{P}],$$

for some $\Lambda \in M_+(\mathcal{H})$, then Theorem 5.1 in Bassetti et al. (2010) implies

$$\frac{L_n - \sum_{i=1}^n \theta_{i-1}}{\sqrt{h_n}} \xrightarrow{\text{stably}} \mathcal{N}(0, \Lambda).$$

3.2.3 C.i.d. RRPS with 0-1 weights

Complete probabilistic characterization of \tilde{P} is difficult to attain, save for some notable exchangeable cases. We make a first step in the direction of non-stationary c.i.d. models by investigating a RRPS that results from delaying the reinforcement in the classical Pólya sequence. For that purpose, let the sequence of weights $(W_n)_{n \geq 1}$ be binary, thereby achieving the phenomenon of randomly discarding part of the observations. For the model with $|\mathbb{X}| = 2$, Aletti et al. (2007) have shown that the distribution of the limiting proportion of balls of color 1 can be derived by considering only the non-zero part of the weights, in which case the model collapses to the two-color Pólya urn scheme and the random limit – to a Beta random variable. We extend this result to more general state spaces.

Proposition 3.2.3. *Let $X = (X_n)_{n \geq 1}$ be a randomly reinforced Pólya sequence with parameters θ , ν and $(W_n)_{n \geq 1}$ such that $W_n \in \{0, 1\}$, and directing measure $\tilde{P} \in \mathbb{K}_P(\Omega, \mathbb{X})$. Suppose X satisfies assumption (A.2). If $\mathbb{P}(W_n = 1 \text{ i.o.}) = 1$, then*

$$\tilde{P} \sim \text{DP}(\theta, \nu).$$

Proof. Define the sequence of random times $(\tau_n)_{n \geq 0}$ by $\tau_0 := 0$ and

$$\tau_n := \inf\{m \in \mathbb{N} : m > \tau_{n-1}, W_m = 1\}, \quad \text{for } n \geq 1,$$

that mark the time, at which reinforcement has taken place. It follows that $\{\tau_n \leq m\} = \{\sum_{i=1}^m W_j \geq n\} \in \sigma(W_1, \dots, W_m) \subseteq \mathcal{F}_m$, for $m, n \geq 1$. By hypothesis,

$$\mathbb{P}(\tau_n < \infty) = 1 - \mathbb{P}(\tau_n = \infty) \geq 1 - \mathbb{P}(W_n = 0 \text{ ult.}) = \mathbb{P}(W_n = 1 \text{ i.o.}) = 1,$$

so $(\tau_n)_{n \geq 0}$ forms a sequence of \mathbb{P} -a.s. finite \mathcal{F} -stopping times. Let $\Omega_0 \in \mathcal{H}$ with $\mathbb{P}(\Omega_0) = 1$ be the set on which all τ_n are real-valued simultaneously. Consider the infinite sequence $(X_{\tau_n})_{n \geq 1}$, which is set equal to some element not in \mathbb{X} outside of Ω_0 (see Section 1.3 of Chapter I for more information). Let $B \in \mathcal{X}$. Then

$$\begin{aligned} \mathbb{P}(X_{\tau_1} \in B) &= \sum_{n=1}^{\infty} \mathbb{P}(X_n \in B | \tau_1 = n) \mathbb{P}(\tau_1 = n) = \sum_{n=1}^{\infty} \int_{\{\tau_1=n\}} \mathbb{P}(X_n \in B | \tau_1 = n) \mathbb{P}(d\omega) = \\ &= \sum_{n=1}^{\infty} \int_{\{\tau_1=n\}} \mathbb{1}_{\{X_n \in B\}}(\omega) \mathbb{P}(d\omega) = \sum_{n=1}^{\infty} \int_{\{\tau_1=n\}} \mathbb{P}(X_n \in B | \mathcal{F}_{n-1} \vee \sigma(W_n))(\omega) \mathbb{P}(d\omega) = \\ &= \sum_{n=1}^{\infty} \int_{\{\tau_1=n\}} \mathbb{P}(X_n \in B | \mathcal{F}_{n-1})(\omega) \mathbb{P}(d\omega) = \sum_{n=1}^{\infty} \int_{\{\tau_1=n\}} \frac{\theta}{\theta + \sum_{j=1}^{n-1} W_j(\omega)} \nu(B) \mathbb{P}(d\omega) = \nu(B), \end{aligned}$$

where we have used assumption (A.2), the fact that $\{\tau_1 = n\} \in \mathcal{F}_{n-1} \vee \sigma(W_n)$, and $W_j = 0$ on $\{\tau_1 = n\}$, for $j = 1, \dots, n-1$. On the other hand, it holds \mathbb{P} -a.s. for $n \geq 1$ that

$$\begin{aligned} \mathbb{P}(X_{\tau_{n+1}} \in B | \mathcal{F}_{\tau_{n+1}-1}) &= \sum_{m=1}^{\infty} P_{m-1}(B) \mathbb{1}_{\{\tau_{n+1}=m\}} = \sum_{m=1}^{\infty} \frac{\theta \nu(B) + \sum_{i=1}^{m-1} W_i \delta_{X_i}(B)}{\theta + \sum_{j=1}^{m-1} W_j} \mathbb{1}_{\{\tau_{n+1}=m\}} = \\ &= \frac{\theta \nu(B) + \sum_{i=1}^{\tau_{n+1}-1} W_i \delta_{X_i}(B)}{\theta + \sum_{j=1}^{\tau_{n+1}-1} W_j} = \frac{\theta \nu(B) + \sum_{i=1}^{\tau_n} W_i \delta_{X_i}(B)}{\theta + \sum_{j=1}^{\tau_n} W_j} = \frac{\theta \nu(B) + \sum_{k=1}^n \delta_{X_{\tau_k}}(B)}{\theta + n}, \end{aligned}$$

where the second to last equality follows by noting that $W_j = 0$ for $j = \tau_n + 1, \dots, \tau_{n+1} - 1$, and the last one holds as $W_j = 1$ for $j = \tau_1, \dots, \tau_n$. Since $\sigma(X_{\tau_1}, \dots, X_{\tau_n}) \subseteq \mathcal{F}_{\tau_{n+1}-1}$, then

$$\mathbb{P}(X_{\tau_{n+1}} \in B | X_{\tau_1}, \dots, X_{\tau_n}) = \frac{\theta \nu(B) + \sum_{k=1}^n \delta_{X_{\tau_k}}(B)}{\theta + n} \quad \text{a.s.}[\mathbb{P}], \text{ for } n \geq 1 \text{ and } B \in \mathcal{X};$$

hence, $(X_{\tau_n})_{n \geq 1}$ is a Pólya sequence with parameters θ and ν . As a result,

$$\mathbb{P}(X_{\tau_{n+1}} \in \cdot | X_{\tau_1}, \dots, X_{\tau_n}) \xrightarrow{w} \tilde{P}(\cdot) \sim \text{DP}(\theta, \nu) \quad \text{a.s.}[\mathbb{P}].$$

Let $f \in C_b(\mathbb{X})$. By Proposition 3.1.1, X is \mathcal{F} -c.i.d., so $\mathbb{E}[f(X_{n+1}) | \mathcal{F}_n] = \frac{\theta \mathbb{E}[f(X_1)] + \sum_{i=1}^n W_i f(X_i)}{\theta + \sum_{j=1}^n W_j}$ converges a.s. $[\mathbb{P}]$. From above

$$\frac{\theta \mathbb{E}[f(X_1)] + \sum_{i=1}^{\tau_n} W_i \mathbb{E}[f(X_i)]}{\theta + \sum_{j=1}^{\tau_n} W_j} \longrightarrow \int_{\mathbb{X}} f(x) \tilde{P}(dx) \quad \text{a.s.}[\mathbb{P}];$$

therefore, $\mathbb{E}[f(X_{n+1}) | \mathcal{F}_n] \xrightarrow{\text{a.s.}} \int_{\mathbb{X}} f(x) \tilde{P}(dx)$ and, ultimately,

$$\mathbb{P}(X_{n+1} \in \cdot | \mathcal{F}_n) \xrightarrow{w} \tilde{P} \quad \text{a.s.}[\mathbb{P}].$$

□

The RRPS with binary weights is conceptually equivalent to the Pólya sequence of Blackwell and MacQueen (1973), provided one disregards the observations with a corresponding weight of zero, at which times the composition of the urn remains the same. In case \mathbb{X} is finite, the process coincides with the k -color Pólya urn and, as a result, the random limit of the vector of predictive distributions is Dirichlet distributed.

The sufficiency assumption in Theorem 3.2.3 requires the weights to attain a non-zero value occasionally, and a related condition can be derived from Lévy's extension to the Borel-Cantelli lemmas (see Williams, 1991, Section 12.5). To that end, suppose

$$\sum_{n=1}^{\infty} \mathbb{P}(W_n = 1 | \mathcal{F}_{n-1}) = \infty \quad \text{a.s.}[\mathbb{P}],$$

that is, learning on the sequence of weights does not inform us on it vanishing. The extended lemma then implies $\sum_{n=1}^{\infty} W_n = \sum_{n=1}^{\infty} \mathbb{1}_{\{W_n=1\}} = \infty$ a.s. $[\mathbb{P}]$, so that $\mathbb{P}(W_n = 1 \text{ i.o.}) = 1$. We give two examples next.

Example 3.2.4 (I.I.D. Weights). Let $(W_n)_{n \geq 1}$ be a sequence of independent and identically Bernoulli distributed random variables with parameter $\theta_0 \in (0, 1)$. It follows that

$$\sum_{n=1}^{\infty} \mathbb{P}(W_n = 1 | \mathcal{F}_{n-1}) = \sum_{n=1}^{\infty} \mathbb{P}(W_n = 1) = \sum_{n=1}^{\infty} \theta_0 = \infty \quad \text{a.s.}[\mathbb{P}].$$

□

Example 3.2.5 (Pólya Urn Weights). Let $(W_n)_{n \geq 1}$ be derived from a two-color Pólya urn scheme with initial composition $(\theta_0, \theta_1) \in \mathbb{N}^2$. It follows that

$$\sum_{n=1}^{\infty} \mathbb{P}(W_n = 1 | \mathcal{F}_{n-1}) = \frac{\theta_0}{\theta_0 + \theta_1} + \sum_{n=1}^{\infty} \frac{\theta_0 + \sum_{i=1}^n W_i}{\theta_0 + \theta_1 + n} \geq \sum_{n=1}^{\infty} \frac{\theta_0}{\theta_0 + \theta_1 + n} = \infty \quad \text{a.s.}[\mathbb{P}].$$

□

Each sequence $(W_n)_{n \geq 1}$ from the previous two examples is a special case of an exchangeable, and hence c.i.d. stochastic process. It turns out that there is a large class of c.i.d. sequences of weights such that Proposition 3.2.3 holds true without requiring explicitly that $\mathbb{P}(W_n = 1 \text{ i.o.}) = 1$. The crucial observation concerns the role of $\mathbb{P}(W_n = 1 \text{ i.o.}) = 1$ in proving that the discovery times are \mathbb{P} -a.s. finite.

Example 3.2.6 (C.I.D. Weights). Let $(W_n)_{n \geq 1}$ be an $\{0, 1\}$ -valued c.i.d. sequence of random variables such that $\frac{1}{n} \sum_{i=1}^n W_i \xrightarrow{\text{a.s.}} \theta^*$ for some absolutely continuous random variable θ^* . As before, denote by $(\tau_n)_{n \geq 0}$ the times that mark each new species discovery. It follows for each $j \geq 1$ that

$$\mathbb{P}(\tau_1 < \infty) = 1 - \lim_{n \rightarrow \infty} \mathbb{P}(\tau_1 > n + j) =$$

$$\begin{aligned}
&= 1 - \lim_{n \rightarrow \infty} \mathbb{P}(W_1 = 0, \dots, W_{n+j} = 0) \geq \\
&\geq 1 - \lim_{n \rightarrow \infty} \mathbb{P}(W_{n+1} = 0, \dots, W_{n+j} = 0) = 1 - \mathbb{E}[(1 - \theta^*)^j],
\end{aligned}$$

where the last equality is true from asymptotic exchangeability. But $\theta^* \neq 0$ a.s.[\mathbb{P}], so

$$\lim_{j \rightarrow \infty} \mathbb{E}[(1 - \theta^*)^j] = 0;$$

therefore, $\mathbb{P}(\tau_1 < \infty) = 1$. Consider the sequence $(W_{\tau_1+n})_{n \geq 1}$. Then

$$\tau_2 \equiv \inf\{m \in \mathbb{N} : W_{\tau_1+m} = 1\}.$$

It follows from Example 1.1 of Berti et al. (2004) that $(W_{\tau_1+n})_{n \geq 1}$ is c.i.d. Moreover, $\frac{1}{n} \sum_{i=1}^n W_{\tau_1+i} \xrightarrow{a.s.} \theta^*$, so the above argument extends to the whole sequence $(\tau_n)_{n \geq 0}$. \square

If it actually holds that $\mathbb{P}(W_n = 1 \text{ i.o.}) = 0$, and thus $\mathbb{P}(W_n = 0 \text{ ult.}) = 1$, then there exists an $N \geq 1$ such that $W_n = 0$ a.s.[\mathbb{P}], for all $n \geq N$. In that case $(X_{N+j})_{j \geq 1}$ is i.i.d. given \mathcal{F}_N . Indeed, one has for $k \geq 1$ and $B \in \mathcal{X}$ that

$$\begin{aligned}
\mathbb{P}(X_{N+k+1} \in B | \mathcal{F}_{N+k}) &= \sum_{i=1}^{N+k} \frac{W_i}{\theta + \sum_{l=1}^{N+k} W_l} \delta_{X_i}(B) + \frac{\theta}{\theta + \sum_{l=1}^{N+k} W_l} \nu(B) = \\
&= \sum_{i=1}^N \frac{W_i}{\theta + \sum_{l=1}^N W_l} \delta_{X_i}(B) + \frac{\theta}{\theta + \sum_{l=1}^N W_l} \nu(B) = \mathbb{P}(X_{N+k+1} \in B | \mathcal{F}_N) = \mathbb{P}(X_{N+1} \in B | \mathcal{F}_N),
\end{aligned}$$

where the last equality follows from the c.i.d. property.

3.3 Central limit theorem

Even if the exact distribution of \tilde{P} is generally unknown, we can try to approximate it through a central limit theorem. In this subsection we provide two such distributional results with respect to different centerings of the normal approximation. As a by-product, we give further details on the random masses p_k^* from (III.1).

3.3.1 Central limit theorem

Recall from Chapter II that RRPSs satisfying assumption (A.1) and GOSs generate the same stochastic processes. As a consequence, already existing results about the latter model could be rephrased in terms of the former. However, in some cases it is more natural to make assumptions directly on the sequence of weights $(W_n)_{n \geq 1}$ in the RRPS representation instead of on the functions $(r_n)_{n \geq 0}$ from the GOS equations. Finding sufficient conditions, under which

$$C_{n,f} := \sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n f(X_i) - \mathbb{E}[f(X_{n+1}) | \mathcal{F}_n] \right),$$

and/or

$$D_{n,f} := \sqrt{n}(\mathbb{E}[f(X_{n+1})|\mathcal{F}_n] - \tilde{P}_f),$$

converge in distribution, is one such example, where $\tilde{P}_f := \int_{\mathbb{X}} f(x)\tilde{P}(dx)$, for any $f \in M_b(\mathcal{X})$. With respect to the GOS reparameterization, Bassetti et al. (2010, Theorem 5.2, Theorem 5.3, Corollary 5.2) prove convergence under quite abstract conditions on the predictive and empirical means, whereas in the technical report to that paper they use a c.i.d. RRPS as an example (Example 4.2, Example 4.5) and put explicit distributional assumptions on the sequence $(W_n)_{n \geq 1}$. The following theorem, which is based on results and techniques derived from Crimaldi et al. (2007), Crimaldi (2009) and Berti et al. (2010), relaxes some of the conditions imposed on $(W_n)_{n \geq 1}$ and, in particular, dispenses with their restrictive $W_n > 0$ assumption.

The central limit result on $C_{n,f}$ is concerned with the rate of convergence of the empirical distribution to the predictive, and as such can be utilized for the estimation of $\mathbb{E}[f(X_{n+1})|\mathcal{F}_n]$ through $\frac{1}{n} \sum_{i=1}^n f(X_i)$ when the sequence of weights is non-observable. On the other hand, convergence of $D_{n,f}$ investigates the speed, with which the predictive distribution approaches its random limit \tilde{P}_f . The particular modes of convergence are discussed in detail in Chapter I, Section 1.3.

Theorem 3.3.1. *Let $X = (X_n)_{n \geq 1}$ be a randomly reinforced Pólya sequence with parameters θ, ν and $(W_n)_{n \geq 1}$. Suppose X satisfies assumption (A.2). If $(W_n)_{n \geq 1}$ are i.i.d. with $W_n \leq \beta < \infty$, then it holds for every $f \in M_b(\mathcal{X})$ that*

$$C_{n,f} \xrightarrow{\text{stably}} \mathcal{N}(0, U_f),$$

and

$$D_{n,f} \xrightarrow{\text{a.s. cond}} \mathcal{N}(0, V_f) \quad \text{w.r.t. } \mathcal{F},$$

where

$$V_f = \frac{\mathbb{E}[W_1^2]}{\mathbb{E}[W_1]^2} (\tilde{P}_{f^2} - (\tilde{P}_f)^2), \quad \text{and} \quad U_f = V_f - (\tilde{P}_{f^2} - (\tilde{P}_f)^2).$$

Remark. We make the implicit assumption that $\mathbb{E}[W_1] > 0$; else $(X_n)_{n \geq 1}$ is an i.i.d. (ν) sequence of random variables, in which case the classical central limit theorems apply.

Proof of Theorem 3.3.1. Let $f \in M_b(\mathcal{X})$, say, $|f| \leq c < \infty$. Denote by

$$P_n(f) := \mathbb{E}[f(X_{n+1})|\mathcal{F}_n], \quad \hat{P}_n(f) := \frac{1}{n} \sum_{i=1}^n f(X_i), \quad N_n(f) := \theta \mathbb{E}[f(X_1)] + \sum_{i=1}^n W_i f(X_i),$$

$$N_n := \theta + \sum_{i=1}^n W_i, \quad \text{and} \quad Q_n := W_{n+1}/N_{n+1},$$

for $n \geq 0$, where $\sum_{i=1}^0 a_i = 0$. It follows that

$$P_n(f) - P_{n+1}(f) = (P_n(f) - f(X_{n+1}))Q_n.$$

As X is \mathcal{F} -c.i.d., one has from Lemma 2.1 and Theorem 2.2 in Berti et al. (2004) that $(P_n(f))_{n \geq 0}$ is a bounded \mathcal{F} -martingale such that $P_n(f) \rightarrow \tilde{P}_f$ a.s. $[\mathbb{P}]$ and in L^1 , and $\hat{P}_n(f) \rightarrow \tilde{P}_f$ a.s. $[\mathbb{P}]$ and in L^1 . Define $H_n := \{2N_n \geq n\mathbb{E}[W_1]\}$, for $n \geq 1$. By the strong law of large numbers, $\frac{1}{n}N_n \xrightarrow{\text{a.s.}} \mathbb{E}[W_1]$; therefore, $\mathbb{P}(H_n^c \text{ i.o.}) = 0$. Finally, set $h := \mathbb{E}[W_1^2]/\mathbb{E}[W_1]^2$.

Part I: $D_{n,f} \xrightarrow{a.s.con\d} \mathcal{N}(0, V_f)$. The first part of the proof is based on a variant of Theorem 2.2 by Crimaldi (2009), whose proof can be found in the Appendix.

Theorem: Let $(M_n)_{n \geq 1}$ be a real-valued martingale w.r.t. a filtration $\mathcal{G} = (\mathcal{G}_n)_{n \geq 0}$ such that $M_n \rightarrow M$ in L^1 , for some $M \in M(\mathcal{H})$, and $H_n \in \mathcal{G}_n$ such that $\mathbb{P}(H_n^c \text{ i.o.}) = 0$. If it holds

- (i) $\mathbb{E}[\sup_{n \in \mathbb{N}} \sqrt{n} \cdot \mathbb{1}_{H_n} |M_n - M_{n+1}|] < \infty$;
- (ii) $n \cdot \sum_{m \geq n} (M_m - M_{m+1})^2 \xrightarrow{a.s.} V$, for some $V \in M_+(\mathcal{H})$;

then

$$\sqrt{n}(M_n - M) \xrightarrow{a.s.con\d} \mathcal{N}(0, V) \quad \text{w.r.t. } \mathcal{G}.$$

where in our case,

$$M_n = P_n(f), \quad M = \tilde{P}_f, \quad V = V_f, \quad \mathcal{G}_n = \mathcal{F}_n.$$

First note that $|P_n(f) - P_{n+1}(f)| \leq 2cQ_n$. Then $\sup_{n \in \mathbb{N}} n^2 \cdot \mathbb{1}_{H_n} |P_n(f) - P_{n+1}(f)|^4 \leq 16c^4 \sum_{n=1}^{\infty} n^2 \cdot \mathbb{1}_{H_n} Q_n^4$, which is integrable since

$$\sum_{n=1}^{\infty} n^2 \mathbb{E}[\mathbb{1}_{H_n} Q_n^4] = \sum_{n=1}^{\infty} n^2 \mathbb{E}\left[\mathbb{1}_{H_n} \left(\frac{W_{n+1}}{N_{n+1}}\right)^4\right] \leq \beta^4 \sum_{n=1}^{\infty} n^2 \mathbb{E}\left[\mathbb{1}_{H_n} \left(\frac{1}{N_n}\right)^4\right] \leq \frac{16\beta^4}{\mathbb{E}[W_1]^4} \sum_{n=1}^{\infty} \frac{1}{n^2} < \infty,$$

where we have used that $1/N_n \leq 2/(n\mathbb{E}[W_1])$ on H_n . Moreover,

$$\begin{aligned} n \cdot \sum_{m \geq n} (P_m(f) - P_{m+1}(f))^2 &= n \cdot \sum_{m \geq n} (P_m(f) - f(X_{m+1}))^2 Q_m^2 = \\ &= n \cdot \sum_{m \geq n} ((P_m(f))^2 + f^2(X_{m+1}) - 2P_m(f) \cdot f(X_{m+1})) Q_m^2. \end{aligned}$$

Define

$$Z_n := \frac{W_{n+1}^2 - \mathbb{E}[W_1^2]}{n}, \quad \text{for } n \geq 1.$$

Then $\mathbb{E}[Z_n] = 0$ and $\sum_{n=1}^{\infty} \text{Var}(Z_n) = \sum_{n=1}^{\infty} n^{-2} \text{Var}(W_1^2) < \infty$; hence, $\sum_{n=1}^{\infty} Z_n < \infty$ a.s.[\mathbb{P}] by Theorem 7.5 in Cinlar (2011). From Lemma A.5 in the Appendix, one has $\sum_{n=1}^{\infty} n^{-1} Z_n < \infty$ a.s.[\mathbb{P}], and thus $n \cdot \sum_{m \geq n} m^{-1} Z_m \xrightarrow{a.s.} 0$. But $n \cdot \sum_{m \geq n} m^{-2} \rightarrow 1$, so $n \cdot \sum_{m \geq n} m^{-2} W_{m+1}^2 \xrightarrow{a.s.} \mathbb{E}[W_1^2]$. As $Q_n^2 / \frac{W_{n+1}^2}{n^2 \mathbb{E}[W_1]^2} \xrightarrow{a.s.} 1$, it follows that

$$\lim_{n \rightarrow \infty} n \cdot \sum_{m \geq n} Q_m^2 = \lim_{n \rightarrow \infty} \frac{1}{\mathbb{E}[W_1]^2} n \cdot \sum_{m \geq n} m^{-2} W_{m+1}^2 = h \quad \text{a.s.}[\mathbb{P}];$$

therefore,

$$\begin{aligned} n \cdot \sum_{m \geq n} P_m(f) Q_m^2 &\longrightarrow h \tilde{P}_f \quad \text{a.s.}[\mathbb{P}], & n \cdot \sum_{m \geq n} (P_m(f))^2 Q_m^2 &\longrightarrow h \cdot (\tilde{P}_f)^2 \quad \text{a.s.}[\mathbb{P}], \\ n \cdot \sum_{m \geq n} P_m(f^2) Q_m^2 &\longrightarrow h \tilde{P}_{f^2} \quad \text{a.s.}[\mathbb{P}]. \end{aligned}$$

Define $Y_n := \sum_{k=1}^{n-1} k \cdot \mathbb{1}_{H_k} (f(X_{k+1}) - P_k(f)) Q_k^2$, for $n \geq 1$. Then $(Y_n)_{n \geq 1}$ is an \mathcal{F} -martingale such that

$$\mathbb{E}[Y_n^2] = \sum_{k=1}^{n-1} k^2 \mathbb{E}\left[\mathbb{1}_{H_k} (f(X_{k+1}) - P_k(f))^2 Q_k^4\right] \leq 4c^2 \sum_{k=1}^{\infty} n^2 \mathbb{E}[\mathbb{1}_{H_n} Q_n^4] < \infty;$$

hence, $\sum_{n=1}^{\infty} n \cdot \mathbb{1}_{H_n} (f(X_{n+1}) - P_n(f)) Q_n^2 < \infty$ a.s.[\mathbb{P}]. It follows from Lemma A.5 in the Appendix that $n \cdot \sum_{m \geq n} \mathbb{1}_{H_m} (f(X_{m+1}) - P_m(f)) Q_m^2 \xrightarrow{a.s.} 0$. As a consequence,

$$n \cdot \sum_{m \geq n} \mathbb{1}_{H_m} f(X_{m+1}) Q_m^2 \longrightarrow h \tilde{P}_f \quad \text{a.s.}[\mathbb{P}].$$

Similarly,

$$n \cdot \sum_{m \geq n} \mathbb{1}_{H_m} f^2(X_{m+1}) Q_m^2 \longrightarrow h \tilde{P}_{f^2} \quad \text{a.s.}[\mathbb{P}], \quad \text{and} \quad n \cdot \sum_{m \geq n} \mathbb{1}_{H_m} P_m(f) f(X_{m+1}) Q_m^2 \longrightarrow h \cdot (\tilde{P}_f)^2 \quad \text{a.s.}[\mathbb{P}];$$

therefore,

$$n \cdot \sum_{m \geq n} \mathbb{1}_{H_m} (P_m(f) - P_{m+1}(f))^2 \longrightarrow h \cdot (\tilde{P}_{f^2} - (\tilde{P}_f)^2) \quad \text{a.s.}[\mathbb{P}].$$

As $\mathbb{1}_{H_n} \xrightarrow{a.s.} 1$, then $n \cdot \sum_{m \geq n} (P_m(f) - P_{m+1}(f))^2 \xrightarrow{a.s.} h \cdot (\tilde{P}_{f^2} - (\tilde{P}_f)^2)$ and the result follows.

Part II: $C_{n,f} \xrightarrow{stably} \mathcal{N}(0, U_f)$. It follows that

$$C_{n,f} = \frac{1}{\sqrt{n}} \sum_{k=1}^n \left\{ f(X_k) - P_{k-1}(f) + k(P_{k-1}(f) - P_k(f)) \right\}.$$

Define

$$C_{n,f}^* := \frac{1}{\sqrt{n}} \sum_{k=1}^n \mathbb{1}_{H_{k-1}} \left\{ f(X_k) - P_{k-1}(f) + k(P_{k-1}(f) - P_k(f)) \right\}.$$

As $\mathbb{P}(\mathbb{1}_{H_n} = 1 \text{ ult.}) = 1$, then $C_{n,f} - C_{n,f}^* \xrightarrow{a.s.} 0$, so from the properties of stable convergence, it is enough that $C_{n,f}^* \xrightarrow{stably} \mathcal{N}(0, U_f)$ for the general result to hold. To that end, we show that the following proposition, which has been suggested by Berti et al. (2011) and is derived from Corollary 7 in Crimaldi et al. (2007), is true for our model.

Proposition: *Let $(\mathcal{G}_n)_{n \geq 0}$ be a filtration on (Ω, \mathcal{H}) , and $M_n = (M_{n,k})_{1 \leq k \leq n}$ be a martingale w.r.t. $(\mathcal{G}_k)_{1 \leq k \leq n}$ such that $M_{n,0} = 0$. Denote by \mathcal{U} the completion of \mathcal{G}_∞ in \mathcal{H} and*

$$Y_{n,k} := M_{n,k} - M_{n,k-1}.$$

If it holds

- (i) $\mathbb{E}[\max_{1 \leq k \leq n} |Y_{n,k}|] \longrightarrow 0$;
- (ii) $\sum_{k=1}^n Y_{n,k}^2 \xrightarrow{p} U$, for some $U \in M_+(\mathcal{U})$;

then

$$\sum_{k=1}^n Y_{n,k} \xrightarrow{stably} \mathcal{N}(0, U).$$

where in our case,

$$Y_{n,k} = \frac{1}{\sqrt{n}} \mathbb{1}_{H_{k-1}} \left\{ f(X_k) - P_{k-1}(f) + k(P_{k-1}(f) - P_k(f)) \right\}, \quad U = U_f, \quad \mathcal{G}_n = \mathcal{F}_n.$$

First note that $\mathbb{E}[Y_{n,k}|\mathcal{F}_{k-1}] = 0$. Regarding (i),

$$\sqrt{n} \max_{1 \leq k \leq n} |Y_{n,k}| \leq \max_{1 \leq k \leq n} \mathbb{1}_{H_{k-1}} |f(X_k) - P_{k-1}(f)| + \max_{1 \leq k \leq n} k \cdot \mathbb{1}_{H_{k-1}} |P_{k-1}(f) - P_k(f)|.$$

As f is bounded, one has $\frac{1}{\sqrt{n}} \mathbb{E}[\max_{1 \leq k \leq n} \mathbb{1}_{H_{k-1}} |f(X_k) - P_{k-1}(f)|] \rightarrow 0$. On the other hand, one has from the first part that $|P_{n-1}(f) - P_n(f)| \leq 2cQ_{n-1} \leq 2c\beta/N_{n-1}$, and hence

$$\frac{1}{\sqrt{n}} \mathbb{E} \left[\max_{1 \leq k \leq n} k \cdot \mathbb{1}_{H_{k-1}} |P_{k-1}(f) - P_k(f)| \right] \leq \frac{1}{\sqrt{n}} \mathbb{E} \left[\max_{1 \leq k \leq n} k \cdot \mathbb{1}_{H_{k-1}} \frac{2c\beta}{N_{k-1}} \right] \leq \frac{b}{\sqrt{n}} \rightarrow 0,$$

for some suitable constant $b \in \mathbb{R}_+$. As a consequence, $\mathbb{E}[\max_{1 \leq k \leq n} |Y_{n,k}|] \rightarrow 0$. Regarding (ii),

$$\begin{aligned} \sum_{k=1}^n Y_{n,k}^2 &= \frac{1}{n} \sum_{k=1}^n \mathbb{1}_{H_{k-1}} (f(X_k) - P_{k-1}(f))^2 + \frac{1}{n} \sum_{k=1}^n k^2 \cdot \mathbb{1}_{H_{k-1}} (P_{k-1}(f) - P_k(f))^2 + \\ &\quad + \frac{2}{n} \sum_{k=1}^n k \cdot \mathbb{1}_{H_{k-1}} (P_{k-1}(f) - P_k(f)) (f(X_k) - P_{k-1}(f)). \end{aligned}$$

From empirical convergence, $\mathbb{P}(\mathbb{1}_{H_n} = 1 \text{ ult.}) = 1$ and Lemma A.4 in the Appendix, it holds that

$$\frac{1}{n} \sum_{k=1}^n \mathbb{1}_{H_{k-1}} (f(X_k) - P_{k-1}(f))^2 \rightarrow \tilde{P}_{f^2} - (\tilde{P}_f)^2 \quad \text{a.s.}[\mathbb{P}].$$

The same arguments as in the first part imply $\sum_{n=1}^{\infty} n(f^2(X_n) - P_{n-1}(f^2))Q_{n-1}^2 < \infty$ a.s. $[\mathbb{P}]$, so one has by Kronecker's lemma that

$$\frac{1}{n} \sum_{k=1}^n k^2 (f^2(X_k) - P_{k-1}(f^2))Q_{k-1}^2 \rightarrow 0 \quad \text{a.s.}[\mathbb{P}].$$

But $\frac{1}{n} \sum_{k=1}^n k^2 P_{k-1}(f^2)Q_{k-1}^2 \xrightarrow{\text{a.s.}} h\tilde{P}_{f^2}$ from Lemma A.4 in the Appendix, so

$$\frac{1}{n} \sum_{k=1}^n k^2 \cdot \mathbb{1}_{H_{k-1}} f^2(X_k)Q_{k-1}^2 \rightarrow h\tilde{P}_{f^2} \quad \text{a.s.}[\mathbb{P}].$$

Similarly, $\frac{1}{n} \sum_{k=1}^n k^2 \cdot \mathbb{1}_{H_{k-1}} f(X_k)P_{k-1}(f)Q_{k-1}^2 \xrightarrow{\text{a.s.}} h \cdot (\tilde{P}_f)^2$ and $\frac{1}{n} \sum_{k=1}^n k^2 \cdot \mathbb{1}_{H_{k-1}} (P_{k-1}(f))^2 Q_{k-1}^2 \xrightarrow{\text{a.s.}} h \cdot (\tilde{P}_f)^2$; therefore,

$$\frac{1}{n} \sum_{k=1}^n k^2 \cdot \mathbb{1}_{H_{k-1}} (P_{k-1}(f) - P_k(f))^2 = \frac{1}{n} \sum_{k=1}^n k^2 \cdot \mathbb{1}_{H_{k-1}} (f(X_k) - P_{k-1}(f))^2 Q_{k-1}^2 \rightarrow V_f \quad \text{a.s.}[\mathbb{P}].$$

Regarding the last term in $\sum_{k=1}^n Y_{n,k}^2$, note that

$$(P_{n-1}(f) - P_n(f))(f(X_n) - P_{n-1}(f)) = -(f(X_n) - P_{n-1}(f))^2 Q_{n-1}.$$

Through the same reasoning as above we have that

$$\frac{2}{n} \sum_{k=1}^n k \cdot \mathbb{1}_{H_{k-1}} (P_{k-1}(f) - P_k(f))(f(X_k) - P_{k-1}(f)) \rightarrow -2(\tilde{P}_{f^2} - (\tilde{P}_f)^2) \quad \text{a.s.}[\mathbb{P}].$$

As a consequence, $\sum_{k=1}^n Y_{n,k}^2 \xrightarrow{\text{a.s.}} U_f$ and the conclusions of the Proposition follow.

It follows from Lemma 1 in Berti et al. (2011) that under the conditions of Theorem 3.3.1

$$(C_{n,f}, D_{n,f}) \xrightarrow{\text{stably}} \mathcal{N}(0, U_f) \otimes \mathcal{N}(0, V_f),$$

which in turn implies that

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n f(X_i) - \tilde{P}_f \right) = C_{n,f} + D_{n,f} \xrightarrow{\text{stably}} \mathcal{N}(0, U_f + V_f).$$

It should be noted that Berti et al. (2009, Section 4) study the stable limit of $\sqrt{n} \cdot \sup_{B \in \mathcal{D}} \|\hat{P}_n(B) - P_n(B)\|$ for the model with i.i.d. weights such that $0 < W_n < \beta$ and W_n is independent of (X_1, \dots, X_n) , where \mathcal{D} is a countable partition of \mathbb{X} in \mathcal{X} . Based on the techniques that were used in the proof of Theorem 3.3.1, we conjecture that their results would continue to hold in case the weights are not bounded from below. In fact, we hypothesize further that the convergence can be extended to $\sqrt{n} \cdot \sup_{B \in \mathcal{X}} \|\hat{P}_n(B) - P_n(B)\|$ since the model uses effectively a countable (random) number of states.

3.3.2 Almost sure discreteness - revisited

Let $(X_k^*)_{k \geq 1}$ denote the distinct species of X in order of appearance. By Lemma A.9 in the Appendix and Theorem 3.2.1, the sequence of predictive distributions P_n converges \mathbb{P} -a.s. on each random set $\{X_k^*\}$ to p_k^* . It is natural to ask whether the conclusions of Theorem 3.3.1 continue to hold with respect to

$$D_{n,k}^* := \sqrt{n} (P_n(\{X_k^*\}) - p_k^*).$$

In addition, such a result would provide additional information on the marginal distributions of the p_k^* 's. The proof of the next proposition follows closely that of Theorem 3.3.1, while taking into account the specifics of working on random sets as we did in (III.2) and Lemma A.9 in the Appendix.

Proposition 3.3.2. *Let $X = (X_n)_{n \geq 1}$ be a randomly reinforced Pólya sequence with parameters θ, ν and $(W_n)_{n \geq 1}$. Suppose X satisfies assumption (A.2) and ν is diffuse. If $(W_n)_{n \geq 1}$ are i.i.d. with $W_n \leq \beta < \infty$, then it holds for each $k \leq L_\infty$ that*

$$D_{n,k}^* \xrightarrow{\text{a.s. cond}} \mathcal{N}\left(0, \frac{\mathbb{E}[W_1^2]}{\mathbb{E}[W_1]^2} (p_k^* - (p_k^*)^2)\right) \quad \text{w.r.t. } \mathcal{F}.$$

Note that the RRPS from Proposition 3.3.2 is proper, which follows from an appeal to the strong law of large numbers with regards to the sequence $(W_n)_{n \geq 1}$. The next result shows, in particular, that the p_k^* 's are \mathbb{P} -a.s. distinct and non-zero. The proof itself is taken from Aletti et al. (2009, Theorem 3.2).

Proposition 3.3.3. *Under the conditions of Proposition 3.3.2, one has $\mathbb{P}(p_k^* = p) = 0$, for all $p \in [0, 1]$.*

Proof. Let $k \leq L_\infty$. Denote by K_n the conditional distribution of $\sqrt{n}(P_n(\{X_k^*\}) - p_k^*)$ given \mathcal{F}_n , and define $U := h(p_k^* - (p_k^*)^2)$, where $h := \mathbb{E}[W_1^2]/\mathbb{E}[W_1]^2$. From Proposition 3.3.2, $K_n \xrightarrow{w} \mathcal{N}(0, U)$ a.s.[\mathbb{P}] or, equivalently,

$$d_D(K_n, \mathcal{N}(0, U)) := \sup_{\text{closed balls } B} |K_n(B) - \mathcal{N}(0, U)(B)| \longrightarrow 0 \quad \text{a.s.}[\mathbb{P}],$$

where d_D is the discrepancy metric (see Gibbs and Su, 2002). Let $\Omega_0 \in \mathcal{H}$ be such that $\mathbb{P}(\Omega_0) = 1$ and $\lim_{n \rightarrow \infty} P_n(\{X_k^*(\omega)\})(\omega) = p_k^*(\omega)$ and $\lim_{n \rightarrow \infty} d_D(K_n(\omega), \mathcal{N}(0, U(\omega))) = 0$, for any $\omega \in \Omega_0$. Suppose, by contradiction, that there exists $p \in [0, 1]$ such that $\mathbb{P}(p_k^* = p) > 0$. As \tilde{P} is \mathcal{F}_∞ -measurable, then $\lim_{n \rightarrow \infty} \mathbb{P}(p_k^* = p | \mathcal{F}_n) = \mathbb{1}_{\{p\}}(p_k^*)$ a.s.[\mathbb{P}], so there is $F \in \mathcal{H} \cap \{p_k^* = p\} \cap \Omega_0$ such that $\mathbb{P}(F) > 0$ and

$$\lim_{n \rightarrow \infty} \mathbb{P}(p_k^* = p | \mathcal{F}_n)(\omega) = 1, \quad \text{for any } \omega \in F.$$

Fix $\omega \in F$. Define $p_n := \sqrt{n}(P_n(\{X_k^*\}) - p)$ and $B_n := \{p_n\}$, for $n \geq 1$. Then

$$d_D(K_n(\omega), \mathcal{N}(0, U(\omega))) \geq |K_n(\omega)(B_n(\omega)) - \mathcal{N}(0, U(\omega))(B_n(\omega))| = K_n(\omega)(B_n(\omega)) = \mathbb{P}(p_k^* = p | \mathcal{F}_n)(\omega),$$

for n large enough, so $\liminf_{n \rightarrow \infty} d_D(K_n(\omega), \mathcal{N}(0, U(\omega))) \geq 1$, absurd. \square

3.4 Inference

Let $X = (X_n)_{n \geq 1}$ be a c.i.d. RRPS with directing measure $\tilde{P} \in \mathbb{K}_P(\Omega, \mathbb{X})$. Fix $f \in M_b(\mathcal{X})$. It follows that $(\mathbb{E}[f(X_{n+1}) | X_1, \dots, X_n])_{n \geq 0}$ is an \mathcal{F}^X -martingale. In fact, under the conditions of Theorem 3.2.1,

$$\mathbb{E}[f(X_{n+1}) | X_1, \dots, X_n] = \mathbb{E}\left[\int_{\mathbb{X}} f(x) \tilde{P}(dx) | X_1, \dots, X_n\right] = \mathbb{E}\left[\sum_k p_k^* f(X_k^*) | X_1, \dots, X_n\right].$$

Therefore, if we truncate the series at some $K \geq 1$, we can approximate the predictive distributions of X for n large enough by

$$\mathbb{E}[f(X_{n+1}) | X_1, \dots, X_n] \approx \sum_{k=1}^K f(X_k^*) \mathbb{E}[p_k^* | X_1, \dots, X_n],$$

where the term on the right side is a weighted sum of posterior means. In this section we provide an asymptotic approximation of the posterior distribution of each p_k^* using Proposition 3.3.2. For that purpose, we consider the two situations of observable and non-observable weights.

3.4.1 Case 1: W_1, \dots, W_n are observable

The rather strong form of convergence in Proposition 3.3.2 can be utilized in a novel way to approximate the conditional distribution of p_k^* given the data. This suggestion originates from a recent paper by Fortini and Petrone (2019), even though a.s. conditional convergence has been around since at least the founding paper by Crimaldi (2009). The proposed approach differs from Bernstein-von-Mises types of theorems (see

Ghosal and van der Vaart, 2017, Chapter 12) as the approximation is given in terms of the model \mathbb{P} instead of some hypothetical "true" probability distribution.

It follows (see Chapter I, Section 1.3) under the conditions of Proposition 3.3.2 that

$$\mathbb{P}(D_{n,k}^* \in \cdot | \mathcal{F}_n) \xrightarrow{w} \mathcal{N}\left(0, \frac{\mathbb{E}[W_1^2]}{\mathbb{E}[W_1]^2} (p_k^* - (p_k^*)^2)\right)(\cdot) \quad \text{a.s.}[\mathbb{P}].$$

Define

$$V := \frac{\mathbb{E}[W_1^2]}{\mathbb{E}[W_1]^2} (p_k^* - (p_k^*)^2), \quad \text{and} \quad V_n := \frac{\sum_{i=1}^n W_i^2}{\sum_{i=1}^n W_i} \left(P_n(\{X_k^*\}) - (P_n(\{X_k^*\}))^2 \right), \quad \text{for } n \geq 1.$$

From the strong law of large numbers and Theorem 3.2.1, one has $V_n \xrightarrow{a.s.} V$. Then a variant of Theorem 4.2 in Fortini and Petrone (2019) implies

$$\mathbb{P}((D_{n,k}^*, V_n) \in \cdot | \mathcal{F}_n) \xrightarrow{w} (\mathcal{N}(0, V) \otimes \delta_V)(\cdot) \quad \text{a.s.}[\mathbb{P}].$$

Note that $V > 0$ a.s. $[\mathbb{P}]$ from Proposition 3.3.3, so $V_n > 0$ a.s. $[\mathbb{P}]$, for all but a finite number of n . As the mapping $(t, u) \mapsto tu$ from \mathbb{R}_+^2 to \mathbb{R}_+ is continuous, we have that

$$\mathbb{E} \left[f \left(\frac{D_{n,k}^*}{\sqrt{V_n}} \right) \middle| \mathcal{F}_n \right] \longrightarrow \int_{\mathbb{R}_+^2} f(tu) (\mathcal{N}(0, V) \otimes \delta_V)(dt, du) = \int_{\mathbb{R}_+} f(t \cdot V^{-1}) \mathcal{N}(0, V)(dt) = \int_{\mathbb{R}_+} f(s) \mathcal{N}(0, 1)(ds),$$

for each $f \in C_b(\mathbb{R}_+)$. Since the cumulative distribution function of the Normal distribution is continuous, it follows \mathbb{P} -a.s. that

$$\mathbb{P}(\sqrt{n}(P_n(\{X_k^*\}) - p_k^*) \leq t \cdot \sqrt{V_n} | \mathcal{F}_n) \longrightarrow \mathcal{N}(0, 1)((-\infty, t]), \quad \text{for } t \in \mathbb{R}.$$

This result allows us to obtain asymptotic credible intervals around p_k^* in the sense that

$$\mathbb{P}(P_n(\{X_k^*\}) - z_\alpha \sqrt{V_n/n} < p_k^* < P_n(\{X_k^*\}) + z_\alpha \sqrt{V_n/n} | \mathcal{F}_n) \approx 1 - \alpha,$$

for n large enough, where z_α is the appropriate critical value from the standard Normal distribution given $100(1 - \alpha)\%$ confidence, with $\alpha \in (0, 1)$.

3.4.2 Case 2: W_1, \dots, W_n are non-observable

When the weights are not observable, one is interested in the posterior distribution of p_k^* given a sample of (X_1, \dots, X_n) . An asymptotic approximation of

$$\mathbb{P}(p_k^* \in \cdot | X_1, \dots, X_n)$$

can be obtained by averaging out the conditional distributions from the previous subsection w.r.t. the posterior distribution of (W_1, \dots, W_n) given (X_1, \dots, X_n) . To that end, suppose that the prior distribution of the now latent (W_1, \dots, W_n) admits a density $p(w_{1:n})$ with $w_{1:n} = (w_1, \dots, w_n)$ and that ν is absolutely continuous w.r.t. the Lebesgue measure with density f_0 . It follows from Bayes' theorem that

$$p(w_{1:n} | x_{1:n}) \propto p(x_{1:n} | w_{1:n}) p(w_{1:n}), \quad (\text{III.4})$$

where we have used the same dominating measure as the one in Petrone and Raftery (1997) regarding $p(x_{1:n})$. Moreover, assumption (A.1) and Theorem 3.2.1 imply that

$$\begin{aligned}
p(x_{1:n}|w_{1:n}) &= p(x_n|x_{1:n-1}, w_{1:n})p(x_{n-1}|x_{1:n-2}, w_{1:n}) \cdots p(x_2|x_1, w_{1:n})p(x_1|w_{1:n}) = \\
&= p(x_n|x_{1:n-1}, w_{1:n-1})p(x_{n-1}|x_{1:n-2}, w_{1:n-2}) \cdots p(x_2|x_1, w_1)p(x_1) = \\
&= \frac{\sum_{i \in \pi_{1,c(2)}} w_i}{\theta + w_1} \cdots \frac{\theta}{\theta + \sum_{j=1}^{c^{-1}(2)-1} w_j} \cdots \frac{\theta}{\theta + \sum_{j=1}^{c^{-1}(l_n)-1} w_j} \cdots \frac{\sum_{i \in \pi_{n-1,c(n)}} w_i}{\theta + \sum_{j=1}^{n-1} w_j} \prod_{k=1}^{l_n} f_0(x_k^*),
\end{aligned}$$

where $\pi_m = \{\pi_{m,1}, \dots, \pi_{m,l_m}\}$ is the partition on $\{1, \dots, m\}$ with length l_m that is generated by $(x_{1:m})$, for $1 \leq m \leq n$, $x_1^*, \dots, x_{l_n}^*$ are the distinct values of $(x_{1:n})$, listed in order of appearance, $c : \mathbb{N} \rightarrow \mathbb{N}$ is the mapping that shows the cluster membership of each x_j , i.e. $c(j) = k$ if and only if $x_j = x_k^*$, $c^{-1}(k) := \min\{m \in \{1, \dots, k\} : x_m = x_k^*\}$, for $k = 1, \dots, l_n$, and we have used the convention that if $\pi_{l,k} = \emptyset$, then the corresponding term on the right hand side drops out. One can then use a Metropolis algorithm with (III.4) to draw samples from the probability distribution of (W_1, \dots, W_n) , with which to approximate the conditional distribution of p_k^* given a sample from (X_1, \dots, X_n) .

Chapter IV

Dominant Pólya sequence

4.1 Introduction

This chapter of the thesis discusses a randomly reinforced Pólya sequence (RRPS) that falls outside the framework of conditional identity in distribution (c.i.d.), but continues to exhibit elements of predictive convergence. On one hand, the suggested model is not constrained by the conditional independence assumption (A.1) from Chapter III, but allows weights to be contemporaneously dependent on observations. As a consequence, the properties implied by Proposition 3.1.1 are no longer guaranteed. On the other hand, the probabilistic structure of the weights is such as to tackle the concrete problem of dominance among species. For that purpose, we map dominance to larger average weight size, indicating, for example, higher species resilience net of random environmental factors. The working hypothesis is that the sequence of observations would be eventually comprised of exclusively dominant species, disturbed only by the occasional discovery of a new dominated species that will die out with time.

Randomly reinforced processes with an inherent dominance structure have been studied within a diverse range of scientific fields such as evolutionary biology, reinforcement learning, information science, operations research and neural networks (see e.g. Beggs, 2005; Martin and Ho, 2002). The same constructions have been considered in the development of urn models that describe randomized, response-adaptive designs of clinical trials, with the potential treatments being depicted as balls of different colors (see Rosenberger (2002), Rosenberger and Lachin (2002) and Hu and Rosenberger (2006) for a review on the topic). The protocol of one such design proceeds with the sequential assignment of treatments to patients, based on uniform draws from an urn containing an initial number balls, and the subsequent reinforcement of the contents of the urn with a certain number of additional balls, depending on the response of the subjects to the assigned treatment. In this context, the color that is associated to the treatment with the highest average response is said to be dominant. The scheme, for which only the color corresponding to the administered treatment is reinforced, has been called a *randomly reinforced urn* (RRU) by Muliere et al. (2006), with Li et al. (1996) and Durham et al. (1998) being the first to study different aspects of the model with dichotomous reinforcements. RRU's most favorable feature, proven by Beggs (2005), Muliere et al. (2006) and Aletti et al.

(2009), shows that the probability of picking the dominant color goes to 1, thereby demonstrating the urn's tendency to select the most favorable treatment as the experiment proceeds. In addition, May and Flournoy (2009) show that the proportion of patients assigned to the dominant treatment tends to one, even though the number of patients assigned to each treatment increases. The review paper by Flournoy et al. (2012) on RRUs provides a summary of other known results, as well as a discussion on statistical inference and a comparison with other urn models, used for response-adaptive designs.

One major shortcoming of the earlier studies on RRUs concerns the fact that they either assume a single best treatment or, alternatively, suppose that all colors are dominant. To that end, Berti et al. (2010) have proposed a k -color RRU scheme, with k finite, for which there could be more than one, but less than k dominant colors, with dominance being specified in terms of asymptotic averages. In that case they are able to show that treatment allocation is again asymptotically optimal. Moreover, Berti and coauthors provide central limit results, with which they draw inference on the composition of the dominant set. In the same setting, but under different specification of the reinforcement, Zhang et al. (2014) study the asymptotic properties of the proportion of balls of each color and the proportion of patients assigned to each treatment.

In this chapter we explore the RRPS that is the conceptual extension of the RRU to infinite colors, in which case we do not know the species/treatments/colors beforehand, but generate them when the need occurs. Such a generalization is non-trivial as the set of treatments is random and depends on the partition of the subjects. In fact, given that the cardinality of the sample space may be uncountable, a more suitable way to describe the observation process is to think of it as the sequential administering of doses from a single treatment (see Chapter V for more on the topic). We proceed by introducing the model first, which is a RRPS with a discriminating weighting process, and then look at its optimality properties under different specifications of the weights (Section 4.2). In Section 4.3 we derive central limit results, whereas in Section 4.4 we fit some of the aforementioned k -color RRUs to our framework. A section on statistical inference closes the chapter.

4.1.1 Model

Let $(X_n)_{n \geq 1}$ be a RRPS with parameters θ , ν and $(W_n)_{n \geq 1}$. Within a dominance setting, observations and weights can be no longer conditionally independent given the past. Instead, it makes sense to limit the influence that previous draws have on the current by assuming that the variability of each weight depends only on the particular dose it is attached to plus an exogenous factor. As a result, the weighting process of the RRPS becomes

$$W_n = h(X_n, U_n), \tag{IV.1}$$

for $n \geq 1$, where $(U_n)_{n \geq 1}$ is a sequence of independent real-valued random variables on $(\Omega, \mathcal{H}, \mathbb{P})$ such that U_n is independent of (X_1, \dots, X_n) , and h is some measurable function from $\mathbb{X} \times \mathbb{R}$ into \mathbb{R}_+ . It follows under (IV.1) that reinforcement is a function of a dose-specific component and an idiosyncratic error U_n (see Section 2.3 of Chapter II for a discussion on this particular form of the weights). We also require that

$$W_n \leq \beta < \infty.$$

Keeping in line with past research, we aim at expressing dominance through the conditional expectation of the weights given their associated doses. To that end, Lemma A.6 in the Appendix implies that

$$\mathbb{E}[W_n|X_n] = \mathbb{E}[h(X_n, U_n)|X_n] = \int_{\Omega} h(X_n, U_n(\omega))\mathbb{P}(d\omega) = w(X_n),$$

for some \mathcal{X} -measurable function $w : \mathbb{X} \rightarrow \mathbb{R}_+$, which we call the *dominance function* of the model. In order to induce a dominance structure through w , we ask for w to be a continuous function on \mathbb{X} such that

$$\{x \in \mathbb{X} : w(x) \geq \bar{w} - \eta\} \subseteq K, \quad (\text{IV.2})$$

for some compact subset $K \subseteq \mathbb{X}$ and $\eta > 0$, with $\bar{w} := \sup_{x \in \mathbb{X}} w(x)$. The last condition ensures the existence of dominant observations as understood by their image under w . To see this define

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

Under (IV.2), one has $\mathcal{D} \subseteq K$. As w is bounded by β , then $\bar{w} \in \mathbb{R}_+$ from the completeness axiom on the real numbers. But $\{\bar{w}\}$ is closed and w is continuous, so $\mathcal{D} = w^{-1}(\{\bar{w}\})$ is closed as well. It follows that \mathcal{D} is compact, which implies that $w(\mathcal{D})$ is compact; therefore, there exists a point $x_0 \in \mathcal{D}$ such that $w(x_0) = \max_{x \in \mathcal{D}} w(x) = \bar{w}$. As a result, the set \mathcal{D} is non-empty, so there is a family of so-called dominant doses that has a higher weight on average. We call a RRPS having weights as in (IV.1) and a dominance function satisfying (IV.2) a *dominant Pólya sequence* (DPS). Finally, we denote by

$$P_n(\cdot) := \mathbb{P}(X_{n+1} \in \cdot | \mathcal{F}_n), \quad \text{and} \quad \hat{P}_n(\cdot) := \frac{1}{n} \sum_{i=1}^n W_i \delta_{X_i}(\cdot),$$

where $\mathcal{F}_n \equiv \mathcal{F}_n^X \vee \mathcal{F}_n^U$, for $n \geq 0$.

4.2 First-order convergence results

The main research question concerning reinforcement schemes on a space with a dominant subset is on determining the minimal conditions, under which the model tends to promote these observations to the extend that they become prevalent in the system. After providing an answer in terms of both predictive and empirical convergence, we investigate the structure that emerges in the limit within the subsequence of dominant observations. Of particular interest are model specifications, for which the limit of the predictive distributions is a random probability measure \tilde{P} that is concentrated on \mathcal{D} . In that case, the stochastic process becomes asymptotically exchangeable with directing measure \tilde{P} , which brings about a sparse structure in the limit. As a final result, we show that the number of clusters in the population approaches a deterministic quantity.

4.2.1 DPS with strictly increasing w

Our first goal is to investigate whether the DPS tends to select the best doses with the accumulation of knowledge. The main requirement for that to happen turns out to be the capacity to sample doses that are

near-dominant with positive probability. In fact, since there is a continuum of doses, from which to select, the probability of picking a particular dose should, in principle, be zero. In order to get to the general optimality result, however, it is necessary to study first the particular family of models, for which w is a strictly increasing function and the space of observations is the interval $[0, 1]$. Under this model specification, the dominant dose is just the one with tag “1” and we require that at least “1” is in the support of ν .

Theorem 4.2.1. *Let $X = (X_n)_{n \geq 1}$ be a $[0, 1]$ -valued dominant Pólya sequence with parameters θ and ν , and a strictly increasing dominance function w . Suppose $1 \in \text{supp}(\nu)$. Then $\frac{1}{n} \sum_{i=1}^n W_i \xrightarrow{a.s.} w(1)$ and*

$$P_n \xrightarrow{w} \delta_1 \quad \text{a.s.}[\mathbb{P}], \quad \hat{P}_n \xrightarrow{w} \delta_1 \quad \text{a.s.}[\mathbb{P}], \quad \frac{1}{n} \sum_{i=1}^n W_i \delta_{X_i} \xrightarrow{w} w(1) \delta_1 \quad \text{a.s.}[\mathbb{P}].$$

Proof. Denote by $N_n(\cdot) := \theta \nu(\cdot) + \sum_{i=1}^n W_i \delta_{X_i}(\cdot)$ and $N_n := \theta + \sum_{i=1}^n W_i$, for $n \geq 1$. We will omit the parenthesis when working with intervals for the sake of clarity. The proof is divided into three parts. First, we demonstrate that $P_n \xrightarrow{w} \tilde{P}$ a.s. $[\mathbb{P}]$, for some $\tilde{P} \in \mathbb{K}_P(\Omega, \mathbb{X})$. Next, we prove that $\tilde{P} = \delta_Z$ a.s. $[\mathbb{P}]$, for some $Z \in M_+(\mathcal{H})$, and then we show that $Z = 1$ a.s. $[\mathbb{P}]$. In order to obtain these results, we need the following three preliminary lemmas.

Lemma 1: $N_n(t, 1] \rightarrow \infty$ a.s. $[\mathbb{P}]$, for each $t \in (0, 1)$.

Proof. Let $t \in (0, 1)$ and set $A = (t, 1]$. As it holds $P_n(A) \geq \frac{\theta \nu(A)}{\theta + n\beta} > 0$ from $\{1\} \in \text{supp}(\nu)$, then $\sum_{n=1}^{\infty} P_n(A) = \infty$ a.s. $[\mathbb{P}]$, and hence $\sum_{n=1}^{\infty} \delta_{X_n}(A) = \infty$ a.s. $[\mathbb{P}]$ by Levy’s extension to the Borel-Cantelli lemmas. Define $\mathcal{G}_n := \mathcal{F}_{n-1} \vee \sigma(X_n)$, for $n \geq 1$. By Lemma A.6 in the Appendix,

$$\sum_{n=1}^{\infty} \mathbb{E}[W_n | \mathcal{G}_n] \delta_{X_n}(A) = \sum_{n=1}^{\infty} w(X_n) \delta_{X_n}(A) \geq w(t) \sum_{n=1}^{\infty} \delta_{X_n}(A) = \infty \quad \text{a.s.}[\mathbb{P}].$$

Define

$$L_n := \sum_{i=1}^n W_i \delta_{X_i}(A) - \sum_{i=1}^n \mathbb{E}[W_i | \mathcal{G}_i] \delta_{X_i}(A) = \sum_{i=1}^n (W_i - \mathbb{E}[W_i | \mathcal{G}_i]) \delta_{X_i}(A), \quad \text{for } n \geq 1.$$

Then $(L_n)_{n \geq 1}$ is a martingale w.r.t. $(\mathcal{G}_n)_{n \geq 1}$, whose increments are bounded. By Theorem 5.3.1 in Durrett (2010), L_n either converges or oscillates between $-\infty$ and $+\infty$ a.s. $[\mathbb{P}]$; in both cases $\sum_{n=1}^{\infty} W_n \delta_{X_n}(A) = \infty$ a.s. $[\mathbb{P}]$, and thus $N_n(A) \xrightarrow{a.s.} \infty$. \square

Lemma 2: $\liminf_{n \rightarrow \infty} P_n(t, 1] > 0$ a.s. $[\mathbb{P}]$, for each $t \in (0, 1)$.

Proof. Let $t \in (0, 1)$. Suppose, without loss of generalization, that $\nu[0, t] > 0$; else, $P_n(t, 1] = 1$ a.s. $[\mathbb{P}]$. Fix $\epsilon \in (0, \frac{w(1)-w(t)}{2})$. Take $s \in (t, 1)$ such that $w(t) + \epsilon < w(s)$; such an s exists as $w([0, 1])$ is connected. Define $T := \inf\{n \in \mathbb{N} : \frac{N_n(s, 1]}{N_n(s, 1] + \beta} > 1 - \frac{\epsilon}{w(t) + \epsilon}\}$. By Lemma 1, $T < \infty$ a.s. $[\mathbb{P}]$. It follows for each $n \geq 1$ that

$$\mathbb{E}\left[\frac{N_{n+1}[0, t]}{N_{n+1}(s, 1]} \middle| \mathcal{F}_n\right] = \mathbb{E}\left[\frac{N_n[0, t] + W_{n+1} \mathbb{1}_{\{X_{n+1} \leq t\}}}{N_n(s, 1] + W_{n+1} \mathbb{1}_{\{X_{n+1} > s\}}} \middle| \mathcal{F}_n\right] =$$

$$\begin{aligned}
&= \mathbb{E} \left[\frac{N_n[0, t] + W_{n+1}}{N_n(s, 1)} \cdot \mathbb{1}_{\{X_{n+1} \leq t\}} \middle| \mathcal{F}_n \right] + \mathbb{E} \left[\frac{N_n[0, t]}{N_n(s, 1) + W_{n+1}} \cdot \mathbb{1}_{\{X_{n+1} > s\}} \middle| \mathcal{F}_n \right] + \\
&\quad + \mathbb{E} \left[\frac{N_n[0, t]}{N_n(s, 1)} \cdot \mathbb{1}_{\{t < X_{n+1} \leq s\}} \middle| \mathcal{F}_n \right] = \\
&= \frac{N_n[0, t]}{N_n(s, 1)} (P_n[0, t] + P_n(s, 1) + P_n(t, s)) + \frac{1}{N_n(s, 1)} \mathbb{E}[W_{n+1} \cdot \mathbb{1}_{\{X_{n+1} \leq t\}} | \mathcal{F}_n] - \\
&\quad - \frac{N_n[0, t]}{N_n(s, 1)} \mathbb{E} \left[\frac{W_{n+1}}{N_n(s, 1) + W_{n+1}} \cdot \mathbb{1}_{\{X_{n+1} > s\}} \middle| \mathcal{F}_n \right] \leq \\
&\leq \frac{N_n[0, t]}{N_n(s, 1)} + w(t) \frac{1}{N_n(s, 1)} P_n[0, t] - (w(t) + \epsilon) \frac{N_n[0, t]}{N_n(s, 1)} \frac{1}{N_n(s, 1) + \beta} P_n(s, 1) = \\
&= \frac{N_n[0, t]}{N_n(s, 1)} \left(1 + w(t) \frac{1}{N_n} - (w(t) + \epsilon) \frac{N_n(s, 1)}{N_n(s, 1) + \beta} \frac{1}{N_n} \right).
\end{aligned}$$

As a result, $\mathbb{E} \left[\frac{N_{T+n+1}[0, t]}{N_{T+n+1}(s, 1)} \middle| \mathcal{F}_{T+n} \right] \leq \frac{N_{T+n}[0, t]}{N_{T+n}(s, 1)}$, so $(\frac{N_{T+n}[0, t]}{N_{T+n}(s, 1)})_{n \geq 1}$ is a non-negative supermartingale w.r.t. $(\mathcal{F}_{T+n})_{n \geq 1}$ and, by Doob's martingale convergence theorem, $\frac{N_{T+n}[0, t]}{N_{T+n}(s, 1)}$ converges \mathbb{P} -a.s. to a finite limit. Since $\frac{P_n[0, t]}{P_n(t, 1)} \leq \frac{P_n[0, t]}{P_n(s, 1)}$, one has

$$\limsup_{n \rightarrow \infty} \frac{P_n[0, t]}{P_n(t, 1)} \leq \limsup_{n \rightarrow \infty} \frac{P_n[0, t]}{P_n(s, 1)} = \lim_{n \rightarrow \infty} \frac{N_n[0, t]}{N_n(s, 1)} = \lim_{n \rightarrow \infty} \frac{N_{T+n}[0, t]}{N_{T+n}(s, 1)} < \infty \quad \text{a.s.}[\mathbb{P}].$$

Then $\liminf_{n \rightarrow \infty} \frac{P_n(t, 1)}{1 - P_n(t, 1)} > 0$ a.s. $[\mathbb{P}]$, which implies $\liminf_{n \rightarrow \infty} P_n(t, 1) > 0$ a.s. $[\mathbb{P}]$. \square

Lemma 3: $P_n(t, 1)$ converges a.s. $[\mathbb{P}]$, for each $t \in (0, 1)$.

Proof. Let $t \in (0, 1)$ and set $A = (t, 1]$. Recall $\sum_{i=0}^k (-1)^i x^{i+1} \leq \frac{x}{1+x} \leq \sum_{i=0}^{k-1} (-1)^i x^{i+1}$, for $0 \leq x \leq 1$ and $k = 2j + 1$ with $j = 0, 1, 2, \dots$. It follows that

$$\begin{aligned}
\mathbb{E}[P_{n+1}(A) | \mathcal{F}_n] &= P_n(A) + \mathbb{E} \left[\frac{W_{n+1}}{N_n + W_{n+1}} (\delta_{X_{n+1}}(A) - P_n(A)) \middle| \mathcal{F}_n \right] = \\
&= P_n(A) + (1 - P_n(A)) \mathbb{E} \left[\frac{W_{n+1}}{N_n + W_{n+1}} \delta_{X_{n+1}}(A) \middle| \mathcal{F}_n \right] - P_n(A) \mathbb{E} \left[\frac{W_{n+1}}{N_n + W_{n+1}} \delta_{X_{n+1}}(A^c) \middle| \mathcal{F}_n \right];
\end{aligned}$$

therefore,

$$\begin{aligned}
\mathbb{E}[P_{n+1}(A) - P_n(A) | \mathcal{F}_n] &= \\
&= P_n(A)(1 - P_n(A)) \left\{ \frac{1}{P_n(A)} \mathbb{E} \left[\frac{W_{n+1} \delta_{X_{n+1}}(A)}{N_n + W_{n+1} \delta_{X_{n+1}}(A)} \middle| \mathcal{F}_n \right] - \right. \\
&\quad \left. - \frac{1}{1 - P_n(A)} \mathbb{E} \left[\frac{W_{n+1} \delta_{X_{n+1}}(A^c)}{N_n + W_{n+1} \delta_{X_{n+1}}(A^c)} \middle| \mathcal{F}_n \right] \right\} \geq \\
&\geq P_n(A)(1 - P_n(A)) \left\{ \frac{1}{P_n(A)} \mathbb{E} \left[\frac{W_{n+1} \delta_{X_{n+1}}(A)}{N_n} - \frac{W_{n+1}^2 \delta_{X_{n+1}}(A)}{N_n^2} \middle| \mathcal{F}_n \right] - \right. \\
&\quad \left. - \frac{1}{1 - P_n(A)} \mathbb{E} \left[\frac{W_{n+1} \delta_{X_{n+1}}(A^c)}{N_n} \middle| \mathcal{F}_n \right] \right\} \geq \\
&\geq \frac{P_n(A)(1 - P_n(A))}{N_n} \left\{ \frac{1}{P_n(A)} \mathbb{E}[W_{n+1} \delta_{X_{n+1}}(A) | \mathcal{F}_n] - w(t) - \frac{\beta^2}{N_n} \right\}.
\end{aligned}$$

Denote by $\hat{w} := \liminf_{n \rightarrow \infty} \frac{1}{P_n(A)} \mathbb{E}[W_{n+1} \delta_{X_{n+1}}(A) | \mathcal{F}_n] - w(t)$. Let $t' \in (t, 1)$. Then

$$\frac{1}{P_n(A)} \mathbb{E}[W_{n+1} \delta_{X_{n+1}}(A) | \mathcal{F}_n] \geq \frac{1}{P_n(A)} [w(t)P_n(t, t') + w(t')P_n(t', 1)] = w(t) + \frac{P_n(t', 1)}{P_n(A)} [w(t') - w(t)].$$

Since $P_n(t', 1) \leq P_n(A) \leq 1$, it holds from Lemma 2 that $\liminf_{n \rightarrow \infty} \frac{P_n(t', 1)}{P_n(A)} > 0$ a.s.[\mathbb{P}]. But w is strictly increasing by hypothesis, so $\hat{w} > 0$. As $N_n > \beta^2/\hat{w}$ eventually, then $(P_n(A))_{n \geq 1}$ becomes eventually a bounded \mathcal{F} -submartingale, and hence converges a.s.[\mathbb{P}]. \square

Part I. Let $g = \mathbb{1}_{(t, 1]}$, for $t \in [0, 1]$. By Lemma 3, $(\mathbb{E}[g(X_{n+1}) | \mathcal{F}_n])_{n \geq 1}$ converges a.s.[\mathbb{P}]. This convergence result extends immediately first to $\mathbb{1}_{(s, t]} = \mathbb{1}_{(t, 1]} - \mathbb{1}_{(s, 1]}$, for $s, t \in [0, 1]$ with $s < t$, and then to any $g \in M_0(\mathcal{B}[0, 1])$. Let $g \in C_b[0, 1]$. By Lemma A.2 in the Appendix, there exist sequences of simple functions $(g_{1,m})_{m \geq 1}, (g_{2,m})_{m \geq 1} \subseteq M_0(\mathcal{B}[0, 1])$ such that $g_{1,m} \leq g \leq g_{2,m}$ and $\sup_{x \in [0, 1]} (g_{2,m}(x) - g_{1,m}(x)) < 1/m$, for each $m \geq 1$. By monotonicity,

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{E}[g_{1,m}(X_{n+1}) | \mathcal{F}_n] &\leq \liminf_{n \rightarrow \infty} \mathbb{E}[g(X_{n+1}) | \mathcal{F}_n] \leq \\ &\leq \limsup_{n \rightarrow \infty} \mathbb{E}[g(X_{n+1}) | \mathcal{F}_n] \leq \lim_{n \rightarrow \infty} \mathbb{E}[g_{2,m}(X_{n+1}) | \mathcal{F}_n] \quad \text{a.s.}[\mathbb{P}], \text{ for } m \geq 1, \end{aligned}$$

where the end limit terms exist from before. We have by construction that

$$0 \leq \lim_{n \rightarrow \infty} \mathbb{E}[g_{2,m}(X_{n+1}) - g_{1,m}(X_{n+1}) | \mathcal{F}_n] \leq \frac{1}{m} \longrightarrow 0;$$

therefore, $\mathbb{E}[g(X_{n+1}) | \mathcal{F}_n]$ converges a.s.[\mathbb{P}]. By Lemma 2.4 in Berti et al. (2004), there exists $\tilde{P} \in \mathbb{K}_P(\Omega, [0, 1])$ such that

$$P_n \xrightarrow{w} \tilde{P} \quad \text{a.s.}[\mathbb{P}].$$

Part II. We show next that $\tilde{P} = \delta_Z$ a.s.[\mathbb{P}], for some $[0, 1]$ -valued random variable Z . Let $\epsilon \in (0, 1)$ and $t \in (0, 1)$ be such that $\mathbb{E}[\tilde{P}(\{t\})] = 0$. Then $\tilde{P}(\{t\}) = 0$ a.s.[\mathbb{P}]. Denote by $A = (t, 1]$. Define

$$T^{(1)} := \inf\{n \in \mathbb{N} : P_n(A)(1 - P_n(A)) \leq \epsilon\}, \quad \text{and} \quad T_n^{(1)} := T^{(1)} \wedge n, \quad \text{for } n \geq 1.$$

Then $T^{(1)}$ is an \mathcal{F} -stopping time, $\{T^{(1)} = \infty\} \subseteq \{T_n^{(1)} \geq k\} \in \mathcal{F}_{k-1}$, for $k = 1, \dots, n$, and

$$\mathbb{E}[P_n(A) | \mathcal{F}_{n-1}] = P_{n-1}(A) + \mathbb{E}\left[\frac{W_n}{N_n} (\delta_{X_n}(A) - P_{n-1}(A)) | \mathcal{F}_{n-1}\right].$$

It follows that

$$\begin{aligned} 1 &\geq \mathbb{E}[P_{T_n^{(1)}}(A)] = \mathbb{E}\left[\sum_{k=1}^n P_k(A) \cdot \mathbb{1}_{\{T_n^{(1)}=k\}}\right] \geq \\ &\geq \mathbb{E}\left[\sum_{k=1}^n (P_k(A) - P_{k-1}(A)) \cdot \mathbb{1}_{\{T_n^{(1)} \geq k\}}\right] = \\ &= \mathbb{E}\left[\sum_{k=1}^n \mathbb{E}[P_k(A) - P_{k-1}(A) | \mathcal{F}_{k-1}] \cdot \mathbb{1}_{\{T_n^{(1)} \geq k\}}\right] \geq \end{aligned}$$

$$\begin{aligned}
&\geq \mathbb{E} \left[\sum_{k=1}^n \mathbb{E}[P_k(A) - P_{k-1}(A) | \mathcal{F}_{k-1}] \cdot \mathbb{1}_{\{T^{(1)} = \infty\}} \right] = \\
&= \mathbb{E} \left[\sum_{k=1}^n \mathbb{E} \left[\frac{W_k}{N_k} (\delta_{X_k}(A) - P_{k-1}(A)) | \mathcal{F}_{k-1} \right] \cdot \mathbb{1}_{\{T^{(1)} = \infty\}} \right] = \\
&= \mathbb{E} \left[\sum_{k=1}^n \left\{ \mathbb{E} \left[\frac{W_k}{N_k} (1 - P_{k-1}(A)) \cdot \mathbb{1}_{\{X_k > t\}} | \mathcal{F}_{k-1} \right] - \right. \right. \\
&\quad \left. \left. - \mathbb{E} \left[\frac{W_k}{N_k} P_{k-1}(A) \cdot \mathbb{1}_{\{X_k \leq t\}} | \mathcal{F}_{k-1} \right] \right\} \cdot \mathbb{1}_{\{T^{(1)} = \infty\}} \right] = \\
&= \mathbb{E} \left[\sum_{k=1}^n P_{k-1}(A) (1 - P_{k-1}(A)) \left\{ \frac{1}{P_{k-1}(A)} \mathbb{E} \left[\frac{W_k}{N_k} \cdot \mathbb{1}_{\{X_k > t\}} | \mathcal{F}_{k-1} \right] - \right. \right. \\
&\quad \left. \left. - \frac{1}{1 - P_{k-1}(A)} \mathbb{E} \left[\frac{W_k}{N_k} \cdot \mathbb{1}_{\{X_k \leq t\}} | \mathcal{F}_{k-1} \right] \right\} \cdot \mathbb{1}_{\{T^{(1)} = \infty\}} \right] \geq \\
&\geq \epsilon \cdot \mathbb{E} \left[\sum_{k=1}^n \left\{ \frac{1}{P_{k-1}(A)} \mathbb{E} \left[\frac{W_k}{N_k} \cdot \mathbb{1}_{\{X_k > t\}} | \mathcal{F}_{k-1} \right] - \right. \right. \\
&\quad \left. \left. - \frac{1}{1 - P_{k-1}(A)} \mathbb{E} \left[\frac{W_k}{N_k} \cdot \mathbb{1}_{\{X_k \leq t\}} | \mathcal{F}_{k-1} \right] \right\} \cdot \mathbb{1}_{\{T^{(1)} = \infty\}} \right] \geq \\
&\geq \epsilon \cdot \mathbb{E} \left[\sum_{k=1}^n \frac{1}{N_{k-1} + \beta} \left\{ \frac{1}{P_{k-1}(A)} \mathbb{E} [w(X_k) \cdot \mathbb{1}_{\{X_k > t\}} | \mathcal{F}_{k-1}] - \right. \right. \\
&\quad \left. \left. - \frac{1}{1 - P_{k-1}(A)} \frac{N_{k-1} + \beta}{N_{k-1}} \mathbb{E} [w(X_k) \cdot \mathbb{1}_{\{X_k \leq t\}} | \mathcal{F}_{k-1}] \right\} \cdot \mathbb{1}_{\{T^{(1)} = \infty\}} \right] \geq \\
&\geq \epsilon \cdot \mathbb{E} \left[\sum_{k=1}^n \frac{1}{N_{k-1} + \beta} \left\{ \frac{1}{P_{k-1}(A)} \mathbb{E} [w(X_k) \cdot \mathbb{1}_{\{X_k > t\}} | \mathcal{F}_{k-1}] - w(t) \frac{N_{k-1} + \beta}{N_{k-1}} \right\} \cdot \mathbb{1}_{\{T^{(1)} = \infty\}} \right],
\end{aligned}$$

where we have used that $P_n(A)(1 - P_n(A)) > \epsilon$ on $\{T^{(1)} = \infty\}$ and $\mathbb{E}[W_n | \mathcal{F}_{n-1} \vee \sigma(X_n)] = w(X_n)$ from before. Let $\xi > 0$ and $m \geq 1$. By continuity, there exists $t_m > t$ such that $\mathbb{E}[\tilde{P}((t, t_m))] \leq \xi/2^{2m}$; hence, $\mathbb{P}(\tilde{P}((t, t_m)) > \xi/2^m) < 1/2^m$ from Markov's inequality. In other words,

for every $\delta, \delta' > 0$, there exists $t' > t$ such that $\mathbb{P}(\tilde{P}((t, t')) > \delta') < \delta$.

Let $\delta \in (0, 1)$ and take $t' \in (t, 1)$ such that $\mathbb{P}(\tilde{P}((t, t')) > \epsilon/2) < \delta$. Define $D_\delta := \{\tilde{P}((t, t')) \leq \epsilon/2\}$. On the other hand, $P_n \xrightarrow{w} \tilde{P}$ a.s.[\mathbb{P}] and $\tilde{P}(\partial A) = \tilde{P}(\{t\}) = 0$ a.s.[\mathbb{P}] both imply $P_n(A) \xrightarrow{a.s.} \tilde{P}(A)$, so we take $\Omega_0 \in \mathcal{H}$ with $\mathbb{P}(\Omega_0) = 1$ and $P_n(A)(\omega) \rightarrow \tilde{P}(A)(\omega)$, for $\omega \in \Omega_0$. Let $\omega \in D_\delta \cap \{T^{(1)} = \infty\} \cap \Omega_0$. Since w is non-decreasing,

$$\begin{aligned}
\frac{1}{\tilde{P}(A)(\omega)} \int_A w(x) \tilde{P}(dx)(\omega) &\geq \frac{1}{\tilde{P}(A)(\omega)} [w(t) \tilde{P}((t, t'))(\omega) + w(t') \tilde{P}((t', 1])(\omega)] = \\
&= w(t) + \frac{\tilde{P}((t', 1])(\omega)}{\tilde{P}(A)(\omega)} [w(t') - w(t)] = \\
&= w(t) + \frac{\tilde{P}(A)(\omega) - \tilde{P}((t, t'))(\omega)}{\tilde{P}(A)(\omega)} [w(t') - w(t)] \geq
\end{aligned}$$

$$\geq w(t) + \frac{1}{2}[w(t') - w(t)],$$

where the last inequality comes from the fact that $\{T^{(1)} = \infty\} \cap \Omega_0 \subseteq \{\tilde{P}(A) \geq \epsilon\}$ as $T^{(1)} = \infty$ implies $P_n(A)(1 - P_n(A)) > \epsilon$, so that $P_n(A) > \epsilon$, for each $n \geq 1$, and hence $\tilde{P}(A) \geq \epsilon$. Since w is continuous and bounded, one has¹ that

$$\mathbb{E}[w(X_n) \cdot \mathbb{1}_{\{X_n > t\}} | \mathcal{F}_{n-1}] \longrightarrow \int_A w(x) \tilde{P}(dx) \quad \text{a.s.}[\mathbb{P}],$$

say, on $\Omega_1 \in \mathcal{H}$ with $\mathbb{P}(\Omega_1) = 1$. As a consequence,

$$\liminf_{n \rightarrow \infty} \frac{1}{P_{n-1}(A)} \mathbb{E}[w(X_n) \cdot \mathbb{1}_{\{X_n > t\}} | \mathcal{F}_{n-1}] \cdot \mathbb{1}_{D_\delta \cap \{T^{(1)} = \infty\} \cap \Omega^*} \geq \left(w(t) + \frac{1}{2}[w(t') - w(t)] \right) \cdot \mathbb{1}_{D_\delta \cap \{T^{(1)} = \infty\} \cap \Omega^*},$$

where $\Omega^* = \Omega_0 \cap \Omega_1$. On the other hand,

$$\mathbb{E}[W_{n+1} | \mathcal{F}_n] = \mathbb{E}[w(X_{n+1}) | \mathcal{F}_n] \longrightarrow \int_{[0,1]} w(x) \tilde{P}(dx) := \tilde{w} \quad \text{a.s.}[\mathbb{P}].$$

As $\sum_{n=1}^{\infty} \mathbb{E}[W_n^2]/n^2 \leq \sum_{n=1}^{\infty} \beta^2/n^2 < \infty$, it holds from Lemma A.7 in the Appendix that

$$\frac{1}{n} \sum_{i=1}^n W_i \longrightarrow \tilde{w} \quad \text{a.s.}[\mathbb{P}],$$

say, on $\Omega_2 \in \mathcal{H}$ with $\mathbb{P}(\Omega_2) = 2$. Denote by $\tilde{\Omega} = \Omega_0 \cap \Omega_1 \cap \Omega_2$. Then $\mathbb{P}(\tilde{\Omega}) = 1$. As $\tilde{P}(A) > 0$ on $\{T^{(1)} = \infty\} \cap \Omega_0$ and w is strictly increasing, then $\tilde{w} > 0$, and thus $N_n/n \rightarrow \tilde{w}$ and $N_n \rightarrow \infty$ on $\{T^{(1)} = \infty\} \cap \tilde{\Omega}$. From the generalized Fatou's lemma,

$$\begin{aligned} \liminf_{n \rightarrow \infty} \mathbb{E} \left[\frac{n}{N_{n-1} + \beta} \left\{ \frac{1}{P_{n-1}(A)} \mathbb{E}[w(X_n) \cdot \mathbb{1}_{\{X_n > t\}} | \mathcal{F}_{n-1}] - w(t) \frac{N_{n-1} + \beta}{N_{n-1}} \right\} \cdot \mathbb{1}_{D_\delta \cap \{T^{(1)} = \infty\} \cap \tilde{\Omega}} \right] &\geq \\ &\geq \frac{1}{2\tilde{w}} \cdot [w(t') - w(t)] \mathbb{P}(D_\delta \cap \{T^{(1)} = \infty\} \cap \tilde{\Omega}). \end{aligned}$$

Suppose, by contradiction, that $\mathbb{P}(D_\delta \cap \{T^{(1)} = \infty\} \cap \tilde{\Omega}) > 0$. As w is strictly increasing, one has $w(t') - w(t) > 0$, so that

$$\begin{aligned} 1 &\geq \epsilon \cdot \mathbb{E} \left[\sum_{k=1}^n \frac{1}{N_{k-1} + \beta} \left\{ \frac{1}{P_{k-1}(A)} \mathbb{E}[w(X_k) \cdot \mathbb{1}_{\{X_k > t\}} | \mathcal{F}_{k-1}] - w(t) \frac{N_{k-1} + \beta}{N_{k-1}} \right\} \cdot \mathbb{1}_{\{T^{(1)} = \infty\}} \right] \geq \\ &\geq \epsilon \cdot \sum_{k=1}^n \frac{1}{k} \mathbb{E} \left[\frac{k}{N_{k-1} + \beta} \left\{ \frac{1}{P_{k-1}(A)} \mathbb{E}[w(X_k) \cdot \mathbb{1}_{\{X_k > t\}} | \mathcal{F}_{k-1}] - w(t) \frac{N_{k-1} + \beta}{N_{k-1}} \right\} \cdot \mathbb{1}_{D_\delta \cap \{T^{(1)} = \infty\} \cap \tilde{\Omega}} \right] \longrightarrow \infty, \end{aligned}$$

absurd, and thus $\mathbb{P}(D_\delta \cap \{T^{(1)} = \infty\} \cap \tilde{\Omega}) = 0$. But $\mathbb{P}(D_\delta^c) < \delta$, so it follows from Boole's inequality that $\mathbb{P}(T^{(1)} = \infty) < \delta$, which holds for arbitrary $\delta > 0$; therefore, $\mathbb{P}(T^{(1)} = \infty) = 0$, that is $T^{(1)}$ is an \mathbb{P} -a.s.

¹Define $P_n^A(\cdot) := P_n(\cdot \cap A)$, for $n \geq 1$, and $\tilde{P}^A(\cdot) := \tilde{P}(\cdot \cap A)$. Then $P_n^A \xrightarrow{w} \tilde{P}^A$ a.s. $[\mathbb{P}]$. Indeed, let $B \in \mathcal{B}[0, 1]$ be open. Denote by $\Omega_0 \in \mathcal{H}$ the \mathbb{P} -a.s. set, on which $P_n \xrightarrow{w} \tilde{P}$ and $\tilde{P}(\partial A) = 0$. Fix $\omega \in \Omega_0$. As $\partial A = \{t\} \subseteq [0, t] = A^c$, then A is open in $[0, 1]$ and so is $A \cap B$; thus,

$$\tilde{P}^A(B) = \tilde{P}(A \cap B) \leq \liminf_{n \rightarrow \infty} P_n(A \cap B) = \liminf_{n \rightarrow \infty} P_n^A(B).$$

finite \mathcal{F} -stopping time. As $(X_n, P_{n-1}, U_n)_{n \geq 1}$ is a Markov process², the strong Markov property implies that $(X_n^{(1)}, P_{n-1}^{(1)}, U_n^{(1)}) := (X_{T^{(1)}+n}, P_{T^{(1)}+n-1}, U_{T^{(1)}+n})$ is itself Markov with the same transition kernels, where $P_0^{(1)} := P_{T^{(1)}-1}$. Define

$$T^{(2)} := \inf\{n \in \mathbb{N} : P_n^{(1)}(A)(1 - P_n^{(1)}(A)) \leq \epsilon\} \equiv \inf\{n > T^{(1)} : P_n(A)(1 - P_n(A)) \leq \epsilon\}.$$

Repeating the same arguments shows that $T^{(2)}$ is finite a.s.[\mathbb{P}]; thus, iterating the procedure implies

$$\mathbb{P}(P_n(A)(1 - P_n(A)) < \epsilon \text{ i.o.}) = 1.$$

As this is true for any $\epsilon > 0$, one has $P_n(A)(1 - P_n(A)) \xrightarrow{p} 0$, and thus $\tilde{P}(A)(1 - \tilde{P}(A)) = 0$ a.s.[\mathbb{P}] since $P_n(A)(1 - P_n(A)) \xrightarrow{p} \tilde{P}(A)(1 - \tilde{P}(A))$ as well. But $\mathbb{E}[\tilde{P}(\cdot)]$ is a probability measure, so from standard results in measure theory $\mathbb{E}[\tilde{P}(\{t\})] = 0$ and, consequently, $\tilde{P}(\{t, 1\}) \in \{0, 1\}$ a.s.[\mathbb{P}], for all but countably many $t \in [0, 1]$. As these form a dense subset of $[0, 1]$, then $\tilde{P}(A) \in \{0, 1\}$ a.s.[\mathbb{P}], for every $A \in \mathcal{B}[0, 1]$. Therefore, by Lemma A.1 in the Appendix, \tilde{P} is \mathbb{P} -a.s. a Dirac measure, with $\tilde{P} = \delta_Z$ a.s.[\mathbb{P}], for some $Z \in M_+(\mathcal{H})$.

Part III. For the last part of the proof we show that $Z = 1$ a.s.[\mathbb{P}]. Let $t \in (0, 1)$ be such that $\mathbb{E}[\tilde{P}(\{t\})] = 0$. Fix $\epsilon \in (0, \frac{w(t)-w(0)}{2})$. Take $s \in (0, t)$ such that³ $\mathbb{E}[\tilde{P}(\{s\})] = 0$ and $w(s) < w(t) - \epsilon$. Then $\tilde{P}(\{t\}) = \tilde{P}(\{s\}) = 0$ a.s.[\mathbb{P}]. Define $T := \inf\{n \in \mathbb{N} : \frac{N_n[t, 1]}{N_n[t, 1] + \beta} > 1 - \frac{\epsilon}{w(s) + \epsilon}\}$. By Lemma 1, $N_n[t, 1] \xrightarrow{a.s.} \infty$, so $T < \infty$ a.s.[\mathbb{P}]. It follows for each $n \geq 1$ that

$$\begin{aligned} \mathbb{E}\left[\frac{N_{n+1}[0, s]}{N_{n+1}[t, 1]} \middle| \mathcal{F}_n\right] &= \mathbb{E}\left[\frac{N_n[0, s] + W_{n+1}\delta_{X_{n+1}}[0, s]}{N_n[t, 1] + W_{n+1}\delta_{X_{n+1}}[t, 1]} \middle| \mathcal{F}_n\right] = \\ &= \mathbb{E}\left[\frac{N_n[0, s] + W_{n+1}}{N_n[t, 1]} \cdot \mathbb{1}_{\{X_{n+1} < s\}} \middle| \mathcal{F}_n\right] + \mathbb{E}\left[\frac{N_n[0, s]}{N_n[t, 1] + W_{n+1}} \cdot \mathbb{1}_{\{X_{n+1} \geq t\}} \middle| \mathcal{F}_n\right] + \\ &\quad + \mathbb{E}\left[\frac{N_n[0, s]}{N_n[t, 1]} \cdot \mathbb{1}_{\{s \leq X_{n+1} < t\}} \middle| \mathcal{F}_n\right] = \\ &= \frac{N_n[0, s]}{N_n[t, 1]} (P_n[0, s] + P_n[t, 1] + P_n[s, t]) + \frac{1}{N_n[t, 1]} \mathbb{E}[W_{n+1} \cdot \mathbb{1}_{\{X_{n+1} < s\}} | \mathcal{F}_n] - \\ &\quad - \frac{N_n[0, s]}{N_n[t, 1]} \mathbb{E}\left[\frac{W_{n+1}}{N_n[t, 1] + W_{n+1}} \cdot \mathbb{1}_{\{X_{n+1} \geq t\}} \middle| \mathcal{F}_n\right] \leq \\ &\leq \frac{N_n[0, s]}{N_n[t, 1]} + w(s) \frac{1}{N_n[t, 1]} P_n[0, s] - (w(s) + \epsilon) \frac{N_n[0, s]}{N_n[t, 1]} \frac{1}{N_n[t, 1] + \beta} P_n[t, 1] = \\ &= \frac{N_n[0, s]}{N_n[t, 1]} \left(1 + w(s) \frac{1}{N_n} - (w(s) + \epsilon) \frac{N_n[t, 1]}{N_n[t, 1] + \beta} \frac{1}{N_n}\right). \end{aligned}$$

As a result, $\mathbb{E}\left[\frac{N_{T+n+1}[0, s]}{N_{T+n+1}[t, 1]} \middle| \mathcal{F}_{T+n}\right] \leq \frac{N_{T+n}[0, s]}{N_{T+n}[t, 1]}$, so $\left(\frac{N_{T+n}[0, s]}{N_{T+n}[t, 1]}\right)_{n \geq 1}$ is a non-negative supermartingale with respect to $(\mathcal{F}_{T+n})_{n \geq 1}$ and, by Doob's martingale convergence theorem, there exists a finite non-negative random variable Y such that $\frac{N_{T+n}[0, s]}{N_{T+n}[t, 1]} \xrightarrow{a.s.} Y$. In fact,

$$\lim_{n \rightarrow \infty} \frac{P_n[0, s]}{P_n[t, 1]} = \lim_{n \rightarrow \infty} \frac{N_n[0, s]}{N_n[t, 1]} = \lim_{n \rightarrow \infty} \frac{N_{T+n}[0, s]}{N_{T+n}[t, 1]} = Y \quad \text{a.s.}[\mathbb{P}].$$

²This follows from the observation that $(U_n)_{n \geq 1}$ is an independent sequence such that U_n is independent of (X_1, \dots, X_n) , P_n is a function of (P_{n-1}, X_n, U_n) , and the conditional distribution of X_n given $(X_1, U_1, \dots, X_{n-1}, U_{n-1})$ is P_{n-1} .

³Such an s indeed exists. It follows that $w([0, 1])$ is connected, so there is $r \in (0, t)$ such that $w(r) = w(t) - \epsilon$. On the other hand, $\mathbb{E}[\tilde{P}(\cdot)]$ is a probability measure, and hence, $\mathbb{E}[\tilde{P}(\{t\})] = 0$ for all but countably many $t \in [0, 1]$. As a result, there exists $s \in (0, r)$ such that $\mathbb{E}[\tilde{P}(\{s\})] = 0$. As w is strictly increasing by assumption, then $w(s) < w(r) = w(t) - \epsilon$.

On the other hand, $P_n[t, 1] \xrightarrow{a.s.} \delta_Z[t, 1]$ and $P_n[0, s] \xrightarrow{a.s.} \delta_Z[0, s]$; therefore,

$$\delta_Z[0, s] = \lim_{n \rightarrow \infty} P_n[0, s] = \lim_{n \rightarrow \infty} \frac{P_n[0, s]}{P_n[t, 1]} \lim_{n \rightarrow \infty} P_n[t, 1] = Y \cdot \delta_Z[t, 1].$$

But $[0, s] \cap [t, 1] = \emptyset$, so $\delta_Z[0, s] = 0$. Take $(r_n)_{n \geq 1}, (s_n)_{n \geq 1} \subseteq (0, 1)$ such that $w(r_n) = w(t) - \frac{w(t) - w(0)}{n}$, $w(r_{n-1}) \leq w(s_n) < w(r_n)$ and $\mathbb{E}[\tilde{P}(\{s_n\})] = 0$, for $n \geq 1$. As w is strictly increasing, then $[0, r_{n-1}] \subseteq [0, s_n] \subseteq [0, r_n]$. Moreover,

$$\bigcup_{n=1}^{\infty} [0, r_n] = \bigcup_{n=1}^{\infty} w^{-1}([w(0), w(r_n)]) = w^{-1}\left(\bigcup_{n=1}^{\infty} [w(0), w(r_n)]\right) = w^{-1}([w(0), w(t)]) = [0, t];$$

thus,

$$\mathbb{P}(Z \geq t) = \delta_Z[t, 1] = 1 - \delta_Z[0, t] = 1 - \lim_{n \rightarrow \infty} \delta_Z[0, s_n] = 1 \quad \text{a.s.}[\mathbb{P}],$$

and $\frac{N_n[0, t]}{N_n[t, 1]} \xrightarrow{a.s.} \frac{\delta_Z[0, t]}{\delta_Z[t, 1]} = 0$. But $\mathbb{E}[\tilde{P}(\{t\})] = 0$ and, consequently, $\frac{N_n[0, t]}{N_n[t, 1]} \xrightarrow{a.s.} 0$ and $\mathbb{P}(Z \geq t) = 1$, for all but countably many $t \in (0, 1)$, which form a dense subset of $[0, 1]$. Therefore, for each $t' \in (0, 1)$, there is $t \in (0, 1)$ such that $t' < t$, and it holds $0 \leq \frac{P_n[0, t']}{P_n[t', 1]} \leq \frac{P_n[0, t]}{P_n[t, 1]} \xrightarrow{a.s.} 0$ and $1 = \mathbb{P}(Z \geq t) \leq \mathbb{P}(Z \geq t') \leq 1$. As a result, $\mathbb{P}(Z \geq t) = 1$, for all $t \in (0, 1)$, so $Z = 1$ a.s. $[\mathbb{P}]$. It follows that $\tilde{w} = \int_{[0, 1]} w(x) \tilde{P}(dx) = w(1)$, and hence $\frac{1}{n} \sum_{i=1}^n W_i \xrightarrow{a.s.} w(1)$. Let $t \in (0, 1)$. Then

$$P_n[t, 1] = \frac{1}{1 + \frac{N_n[0, t]}{N_n[t, 1]}} \longrightarrow 1 = \delta_1[t, 1] \quad \text{a.s.}[\mathbb{P}], \quad \text{and} \quad P_n[0, t] \longrightarrow 0 = \delta_1[0, t] \quad \text{a.s.}[\mathbb{P}];$$

thus,

$$\frac{1}{n} \sum_{i=1}^n W_i \delta_{X_i}[t, 1] = \frac{\theta + \sum_{i=1}^n W_i}{n} \frac{N_n[t, 1]}{N_n[0, t] + N_n[t, 1]} \frac{\sum_{i=1}^n W_i \delta_{X_i}[t, 1]}{N_n[t, 1]} \longrightarrow w(1) = w(1) \delta_1[t, 1] \quad \text{a.s.}[\mathbb{P}],$$

and $\frac{1}{n} \sum_{i=1}^n W_i \delta_{X_i}[0, t] \xrightarrow{a.s.} 0 = w(1) \delta_1[0, t]$. As these form a countable convergence-determining family of subsets, it holds

$$\frac{1}{n} \sum_{i=1}^n W_i \delta_{X_i} \xrightarrow{w} w(1) \delta_1 \quad \text{a.s.}[\mathbb{P}].$$

Let $g \in C_b[0, 1]$. As $\mathbb{E}[g(X_{n+1}) | \mathcal{F}_n] \xrightarrow{a.s.} g(1)$, one has from Lemma A.7 in the Appendix that $\frac{1}{n} \sum_{i=1}^n g(X_i) \xrightarrow{a.s.} g(1)$, and thus

$$\frac{1}{n} \sum_{i=1}^n \delta_{X_i} \xrightarrow{w} \delta_1 \quad \text{a.s.}[\mathbb{P}].$$

□

In particular, Theorem 4.2.1 implies (see Chapter I, Section 1.3) that

$$X_n \xrightarrow{p} 1.$$

This convergence cannot be strengthened to hold \mathbb{P} -a.s. as one is bound to pick dominated doses infinitely often. Indeed, let $t \in (0, 1)$ be such that $\nu((t, 1]) < 1$. Unless $\nu(\{1\}) = 1$, such a t exists from the regularity of probability measures on metric spaces. It follows that

$$P_n((0, t]) \geq \frac{\theta\nu([0, t])}{\theta + n\beta} > 0, \quad \text{for } n \geq 1;$$

thus, $\sum_{n=1}^{\infty} \mathbb{P}(X_{n+1} \leq t | \mathcal{F}_n) = \infty$. By Lévy's extension to the Borel-Cantelli lemmas,

$$\mathbb{P}(X_{n+1} \leq t \text{ i.o.}) = 1;$$

therefore, there would be a non-dominant observation occasionally. Nevertheless, Theorem 4.2.1 shows that the number of observed, near-dominant doses increases with rate n , which means that suboptimal doses will be relatively rarely administered.

4.2.2 DPS with general w

Given the results in the previous subsection, we are ready to tackle the optimal allocation problem for the DPS with an arbitrary dominance function. Under the same hypothesis as above, we show that the predictive and empirical distributions concentrate asymptotically around the set of dominant doses. For that purpose, define by

$$\mathcal{D}_\delta := \{x \in \mathbb{X} : d(x, \mathcal{D}) < \delta\},$$

the δ -neighborhood of \mathcal{D} , for $\delta > 0$, where $d(x, \mathcal{D}) := \inf\{d(x, y) : y \in \mathcal{D}\}$ denotes the distance from the dominant set and d is the inherent metric on \mathbb{X} .

Theorem 4.2.2. *Let $X = (X_n)_{n \geq 1}$ be a dominant Pólya sequence with parameters θ , ν and w . Suppose $\mathcal{D} \subseteq \text{supp}(\nu)$. Then $\frac{1}{n} \sum_{i=1}^n W_i \xrightarrow{\text{a.s.}} \bar{w}$ and, for each $\delta > 0$, it holds*

$$P_n(\mathcal{D}_\delta) \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}], \quad \hat{P}_n(\mathcal{D}_\delta) \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}], \quad \frac{1}{n} \sum_{i=1}^n W_i \delta_{X_i}(\mathcal{D}_\delta) \longrightarrow \bar{w} \quad \text{a.s.}[\mathbb{P}].$$

Remark. As long as it holds $\mathcal{D} \neq \emptyset$ and $\nu(\mathcal{D}) > 0$, the continuity of w is non-essential for proving the above results. In that case, existence of K and η from (IV.2) is similarly unnecessary, whenever $\bar{w} > \sup_{x \in \mathcal{D}^c} w(x)$. In fact, under those stronger conditions, more can be said about the predictive and empirical convergence of the process than what is stated in the theorem (see Section 4.2.4).

Proof of Theorem 4.2.2. Define

$$X_n^* := \frac{w(X_n)}{\bar{w}}, \quad W_n^* := W_n, \quad P_n^*(B) := P_n(w^{-1}(\bar{w}B)), \quad \nu^*(B) := \nu(w^{-1}(\bar{w}B)),$$

for $n \geq 1$ and $B \in \mathcal{B}[0, 1]$, where $w^{-1}(\bar{w}B) = \{x \in \mathbb{X} : w(x)/\bar{w} \in B\}$. As w is continuous, then $1 \in \text{supp}(\nu^*)$ and

$$\mathbb{E}[W_n^* | X_n^*] = \mathbb{E}[\mathbb{E}[W_n | X_n] | X_n^*] = \bar{w} X_n^*.$$

Moreover,

$$P_n^*(B) = \frac{\theta\nu(w^{-1}(\bar{w}B)) + \sum_{i=1}^n W_i \delta_{X_i}(w^{-1}(\bar{w}B))}{\theta + \sum_{j=1}^n W_j} = \frac{\theta\nu^*(B) + \sum_{i=1}^n W_i^* \delta_{X_i^*}(B)}{\theta + \sum_{j=1}^n W_j^*},$$

so $X^* = (X_n^*)_{n \geq 1}$ is⁴ a $[0, 1]$ -valued DPS with parameters θ and ν^* , and the continuous and strictly increasing dominance function $w^* : [0, 1] \rightarrow [0, \bar{w}]$, given by $w^*(t) := \bar{w}t$, for $t \in [0, 1]$. Let $\delta > 0$. From standard topological results, \mathcal{D}_δ is open. Suppose, by contradiction, that $\sup_{x \in \mathcal{D}_\delta^c} w(x) \geq \bar{w}$. Then there exists $(x_n)_{n \geq 1} \subseteq \mathcal{D}_\delta^c$ such that $w(x_n) > \bar{w} - \frac{1}{n}$, for $n \geq 1$. From (IV.2), it follows that the sequence should be ultimately in K . But K is compact, so there exists by sequential compactness $(n_k) \subseteq (n)$ such that $x_{n_k} \rightarrow \bar{x}$, for some $\bar{x} \in K$. In fact, as \mathcal{D}_δ^c is closed, it is $\bar{x} \in \mathcal{D}_\delta^c \subseteq \mathcal{D}^c$. On the other hand, it holds $w(\bar{x}) = \bar{w}$ from the assumptions on $(x_n)_{n \geq 1}$, and thus, $\bar{x} \in \mathcal{D}$, absurd. Let $\epsilon = (\bar{w} - \sup_{x \in \mathcal{D}_\delta^c} w(x))/2$. Then $\mathcal{D}_\delta^c \subseteq \{x \in \mathbb{X} : w(x) < \bar{w} - \epsilon\}$. From monotonicity and by Theorem 4.2.1,

$$P_n(\mathcal{D}_\delta^c) \leq \mathbb{P}(w(X_{n+1}) < \bar{w} - \epsilon | \mathcal{F}_n) = P_n^*([0, 1 - \epsilon/\bar{w}]) \longrightarrow 0 \quad \text{a.s.}[\mathbb{P}];$$

thus, $P_n(\mathcal{D}_\delta) \xrightarrow{\text{a.s.}} 1$, for any $\delta > 0$. Moreover,

$$\frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\mathcal{D}_\delta) \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}],$$

from Lemma A.7 in the Appendix. Finally,

$$\frac{1}{n} \sum_{i=1}^n W_i = \frac{1}{n} \sum_{i=1}^n W_i^* \longrightarrow w^*(1) = \bar{w} \quad \text{a.s.}[\mathbb{P}];$$

therefore, using that $0 \leq \frac{1}{n} \sum_{i=1}^n W_i \delta_{X_i}(\mathcal{D}_\delta^c) \leq \beta \frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\mathcal{D}_\delta^c) \xrightarrow{\text{a.s.}} 0$, one has

$$\frac{1}{n} \sum_{i=1}^n W_i \delta_{X_i}(\mathcal{D}_\delta) = \frac{1}{n} \sum_{i=1}^n W_i - \frac{1}{n} \sum_{i=1}^n W_i \delta_{X_i}(\mathcal{D}_\delta^c) \longrightarrow \bar{w} \quad \text{a.s.}[\mathbb{P}].$$

□

An immediate consequence of Theorem 4.2.2 is a sort of convergence in probability towards the dominant set. Let $\delta > 0$. By the dominated convergence theorem,

$$\mathbb{P}(d(X_n, \mathcal{D}) \geq \delta) = \mathbb{P}(X_n \in \mathcal{D}_\delta^c) = \mathbb{E}[\mathbb{P}(X_n \in \mathcal{D}_\delta^c | \mathcal{F}_{n-1})] \longrightarrow 0;$$

therefore,

$$d(X_n, \mathcal{D}) \xrightarrow{P} 0.$$

⁴Strictly speaking, X^* differs from the definition of a DPS in that W_n^* is a function of (X_n, U_n) , but not of (X_n^*, U_n) , and P_n^* is the conditional distribution of X_{n+1}^* given \mathcal{F}_n , instead of $\mathcal{F}_n^* = \mathcal{F}_n^{X^*} \vee \mathcal{F}_n^U$. Nevertheless, the conclusions of Theorem 4.2.1 continue to hold for X^* even after changing the filtration since

$$\mathbb{E}[W_n^* | \mathcal{F}_{n-1} \vee \sigma(X_n)] = w^*(X_n^*),$$

and as $(X_n, P_{n-1}, U_n, X_n^*, P_{n-1}^*)_{n \geq 1}$ is a Markov process.

Curiously, if the dominant dose is only one, i.e. $\mathcal{D} = \{\bar{x}\}$ for some $\bar{x} \in \mathbb{X}$, then

$$X_n \xrightarrow{P} \bar{x}.$$

In either case, one cannot extend the convergence in probability to almost sure. Indeed, under the stronger assumption $0 < \nu(\mathcal{D}) < 1$, one has $\sum_{n=1}^{\infty} P_n(\mathcal{D}^c) = \infty$ a.s.[\mathbb{P}], and thus $\sum_{n=1}^{\infty} \delta_{X_n}(\mathcal{D}^c) = \infty$ a.s.[\mathbb{P}] by Lévy's extension to the Borel-Cantelli lemmas, so one is certain to see infinitely many dominated doses. Yet, $\frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\mathcal{D}_\delta^c) \xrightarrow{a.s.} 0$; hence, $\sum_{i=1}^n \delta_{X_i}(\mathcal{D}_\delta) / \sum_{i=1}^n \delta_{X_i}(\mathcal{D}_\delta^c) \xrightarrow{a.s.} \infty$, for any $\delta > 0$, so the rate with which one discovers near-dominant relative to non-dominant doses increases.

In fact, the following result sheds some light on the rate of convergence. For that purpose, let $w_0 \in (\bar{w} - \eta, \bar{w})$, where η was introduced in (IV.2). Take $x_0 \in w^{-1}(\{w_0\})$ and define $\delta_0 := d(x_0, \mathcal{D})$, $\bar{w}_0 := \sup_{x \in \mathcal{D}_{\delta_0}^c} w(x)$ and

$$\mathcal{D}_0 := \{x \in \mathcal{D}_{\delta_0}^c : w(x) = \bar{w}_0\}.$$

Then $\mathcal{D}_0 \neq \emptyset$ collects the doses that are dominant within $\mathcal{D}_{\delta_0}^c$ and have average weight $\bar{w}_0 < \bar{w}$. We show next that $\sum_{i=1}^n \delta_{X_i}(\mathcal{D}_{\delta_0}^c)$ is \mathbb{P} -a.s. eventually greater than any power of n less than \bar{w}_0/\bar{w} .

Proposition 4.2.3. *Suppose $\mathcal{D}_0 \subseteq \text{supp}(\nu(\cdot \cap \mathcal{D}_{\delta_0}^c))$. Under the conditions of Theorem 4.2.2, one has*

$$\frac{1}{n^\gamma} \sum_{i=1}^n \delta_{X_i}(\mathcal{D}_{\delta_0}^c) \longrightarrow \begin{cases} 0 & \gamma > \bar{w}_0/\bar{w}, \\ \infty & \gamma < \bar{w}_0/\bar{w} \end{cases} \quad \text{a.s.}[\mathbb{P}].$$

Proof. Denote by

$$M_n(\cdot) := \sum_{i=1}^n \delta_{X_i}(\cdot) + 1, \quad N_n(\cdot) := \theta \nu(\cdot) + \sum_{i=1}^n W_i \delta_{X_i}(\cdot), \quad N_n = \theta + \sum_{i=1}^n W_i,$$

for $n \geq 1$. As $\sum_{i=1}^n \delta_{X_i}(\mathcal{D}_{\delta_0})$ and $\sum_{i=1}^n \delta_{X_i}(\mathcal{D}_{\delta_0}^c)$ both diverge, we can work with $M_n(\mathcal{D}_{\delta_0})$ and $M_n(\mathcal{D}_{\delta_0}^c)$ instead. From Theorem 4.2.2,

$$\frac{N_n(\mathcal{D}_{\delta_0})}{M_n(\mathcal{D}_{\delta_0})} = \frac{n}{\sum_{i=1}^n \delta_{X_i}(\mathcal{D}_{\delta_0}) + 1} \frac{1}{n} \left(\theta \nu(\mathcal{D}_{\delta_0}) + \sum_{i=1}^n W_i \delta_{X_i}(\mathcal{D}_{\delta_0}) \right) \longrightarrow \bar{w} \quad \text{a.s.}[\mathbb{P}].$$

On the other hand, Theorem 4.2.2 as applied to the subsequence of $\mathcal{D}_{\delta_0}^c$ -valued observations, which is itself a DPS, implies

$$\frac{N_n(\mathcal{D}_{\delta_0}^c)}{M_n(\mathcal{D}_{\delta_0}^c)} \longrightarrow \bar{w}_0 \quad \text{a.s.}[\mathbb{P}].$$

Let $\psi > \bar{w}/\bar{w}_0$. It follows for each $n \geq 1$ that

$$\begin{aligned} \mathbb{E} \left[\frac{M_{n+1}(\mathcal{D}_{\delta_0})}{M_{n+1}(\mathcal{D}_{\delta_0}^c)^\psi} - \frac{M_n(\mathcal{D}_{\delta_0})}{M_n(\mathcal{D}_{\delta_0}^c)^\psi} \middle| \mathcal{F}_n \right] &= \frac{P_n(\mathcal{D}_{\delta_0})}{M_n(\mathcal{D}_{\delta_0}^c)^\psi} + \left(\frac{M_n(\mathcal{D}_{\delta_0})}{(M_n(\mathcal{D}_{\delta_0}^c) + 1)^\psi} - \frac{M_n(\mathcal{D}_{\delta_0})}{M_n(\mathcal{D}_{\delta_0}^c)^\psi} \right) P_n(\mathcal{D}_{\delta_0}^c) \leq \\ &\leq \frac{P_n(\mathcal{D}_{\delta_0})}{M_n(\mathcal{D}_{\delta_0}^c)^\psi} + P_n(\mathcal{D}_{\delta_0}^c) M_n(\mathcal{D}_{\delta_0}) \left(-\frac{\psi}{M_n(\mathcal{D}_{\delta_0}^c)^{\psi+1}} + \frac{c}{M_n(\mathcal{D}_{\delta_0}^c)^{\psi+2}} \right) = \\ &= \frac{P_n(\mathcal{D}_{\delta_0})}{M_n(\mathcal{D}_{\delta_0}^c)^\psi} \left(1 - \psi \frac{M_n(\mathcal{D}_{\delta_0})}{N_n(\mathcal{D}_{\delta_0})} \frac{N_n(\mathcal{D}_{\delta_0}^c)}{M_n(\mathcal{D}_{\delta_0}^c)} + c \frac{M_n(\mathcal{D}_{\delta_0})}{N_n(\mathcal{D}_{\delta_0})} \frac{N_n(\mathcal{D}_{\delta_0}^c)}{M_n(\mathcal{D}_{\delta_0}^c)^2} \right), \end{aligned}$$

where the inequality is derived from Taylor's expansion of the function $f(x) = (a+x)^{-\psi}$ with $a = M_n(\mathcal{D}_{\delta_0}^c)$ and $x = 1$, for a suitable constant $c \in \mathbb{R}_+$. But

$$\limsup_{n \rightarrow \infty} \left(1 - \psi \frac{M_n(\mathcal{D}_{\delta_0})}{N_n(\mathcal{D}_{\delta_0})} \frac{N_n(\mathcal{D}_{\delta_0}^c)}{M_n(\mathcal{D}_{\delta_0}^c)} + c \frac{M_n(\mathcal{D}_{\delta_0})}{N_n(\mathcal{D}_{\delta_0})} \frac{N_n(\mathcal{D}_{\delta_0}^c)}{M_n(\mathcal{D}_{\delta_0}^c)^2} \right) < 0 \quad \text{a.s.}[\mathbb{P}];$$

therefore, $(M_n(\mathcal{D}_{\delta_0})/M_n(\mathcal{D}_{\delta_0}^c)^\psi)_{n \geq 1}$ is eventually a positive supermartingale and converges \mathbb{P} -a.s. to a finite limit. From Theorem 4.2.2, one has $\frac{M_n(\mathcal{D}_{\delta_0})}{n} \xrightarrow{a.s.} 1$, and thus $\limsup_{n \rightarrow \infty} \frac{n^{1/\psi}}{M_n(\mathcal{D}_{\delta_0}^c)} < \infty$ a.s. $[\mathbb{P}]$, so $\limsup_{n \rightarrow \infty} \frac{n^{1/(\psi+\epsilon)}}{M_n(\mathcal{D}_{\delta_0}^c)} = 0$ a.s. $[\mathbb{P}]$, for any $\epsilon > 0$. Repeating the same steps, one has $\limsup_{n \rightarrow \infty} \frac{n^{1/(\psi'+\epsilon)}}{M_n(\mathcal{D}_{\delta_0}^c)} = 0$ a.s. $[\mathbb{P}]$, for any $\bar{w}/\bar{w}_0 < \psi' < \psi$ and, in particular, with $\epsilon = \psi - \psi' > 0$. As a consequence,

$$\frac{1}{n^{1/\psi}} \sum_{i=1}^n \delta_{X_i}(\mathcal{D}_{\delta_0}^c) \longrightarrow \infty \quad \text{a.s.}[\mathbb{P}],$$

for any $\psi > \bar{w}/\bar{w}_0$. Analogously, $(M_n(\mathcal{D}_{\delta_0}^c)/M_n(\mathcal{D}_{\delta_0})^{1/\psi})_{n \geq 1}$ is eventually a positive supermartingale, for any $\psi < \bar{w}/\bar{w}_0$, and hence converges \mathbb{P} -a.s. to a finite limit. Then $\limsup_{n \rightarrow \infty} \frac{M_n(\mathcal{D}_{\delta_0}^c)}{n^{1/\psi}} < \infty$ a.s. $[\mathbb{P}]$, so $\limsup_{n \rightarrow \infty} \frac{M_n(\mathcal{D}_{\delta_0}^c)}{n^{1/(\psi-\epsilon)}} = 0$ a.s. $[\mathbb{P}]$, for any $0 < \epsilon < \psi$. Repeating the same steps, one has $\limsup_{n \rightarrow \infty} \frac{M_n(\mathcal{D}_{\delta_0}^c)}{n^{1/(\psi'-\epsilon)}} = 0$ a.s. $[\mathbb{P}]$, for any $\psi < \psi' < \bar{w}/\bar{w}_0$ and, in particular, with $\epsilon = \psi' - \psi > 0$. As a result,

$$\frac{1}{n^{1/\psi}} \sum_{i=1}^n \delta_{X_i}(\mathcal{D}_{\delta_0}^c) \longrightarrow 0 \quad \text{a.s.}[\mathbb{P}],$$

for any $\psi < \bar{w}/\bar{w}_0$. □

The behavior of $\frac{1}{n^\gamma} \sum_{i=1}^n \delta_{X_i}(\mathcal{D}_{\delta_0}^c)$, for $\gamma = \bar{w}_0/\bar{w}$, remains an open question in the general case, with Zhang et al. (2014) showing that the limit is random on $(0, \infty)$ for k colors. In fact, they derive its exact distribution under additional assumptions. Unfortunately, the complexity of using an infinite state space prevents their results to be immediately adapted to the whole class of DPSs.

4.2.3 Convergence on the dominant subsequence

As the DPS with an arbitrary dominance function concentrates all of its mass around the dominant set, one could be interested in knowing the behavior of that part of the process, which is in \mathcal{D} . A technical prerequisite is the adoption of the stronger assumption of having non-zero probability of observing dominant doses, $\nu(\mathcal{D}) > 0$, which may be unrealistic depending on the application. Under this condition, one can prove the a.s. weak convergence of the predictive and empirical distributions of the *dominant subsequence*. To that end, define the sequence of random variables $(T_n)_{n \geq 0}$ by $T_0 := 0$ and

$$T_n := \inf\{n \in \mathbb{N} : n > T_{n-1}, X_n \in \mathcal{D}\}, \quad \text{for } n \geq 1,$$

that mark the time, at which a dominant dose has been administered. Then the sequence of dominant doses, $\tilde{X} = (\tilde{X}_n)_{n \geq 1}$, listed in order of delivery, and the corresponding weights, $\tilde{W} = (\tilde{W}_n)_{n \geq 1}$, are given by

$$\tilde{X}_n := X_{T_n}, \quad \text{and} \quad \tilde{W}_n := W_{T_n}, \quad \text{for } n \geq 1,$$

respectively. The accumulated information at those times is the filtration $(\tilde{\mathcal{F}}_n)_{n \geq 0}$, where $\tilde{\mathcal{F}}_0 := \{\emptyset, \Omega\}$ and $\tilde{\mathcal{F}}_n := \mathcal{F}_n^{\tilde{X}} \vee \mathcal{F}_n^{\tilde{W}}$, for $n \geq 1$. The next theorem shows that the predictive and empirical distributions of \tilde{X} converge to one and the same random probability measure, whose mass is distributed over \mathcal{D} . In fact, one can perceive these results as if the model is a DPS with $\mathbb{X} = \mathcal{D}$.

Theorem 4.2.4. *Let $X = (X_n)_{n \geq 1}$ be a dominant Pólya sequence with parameters θ, ν and w . Suppose $\nu(\mathcal{D}) > 0$. Then $\frac{1}{n} \sum_{i=1}^n \tilde{W}_i \xrightarrow{a.s.} \bar{w}$ and there exists an $\tilde{P} \in \mathbb{K}_P(\Omega, \mathcal{D})$ such that*

$$\mathbb{P}(\tilde{X}_{n+1} \in \cdot | \tilde{\mathcal{F}}_n) \xrightarrow{w} \tilde{P}(\cdot) \quad \text{a.s.}[\mathbb{P}], \quad \frac{1}{n} \sum_{i=1}^n \delta_{\tilde{X}_i}(\cdot) \xrightarrow{w} \tilde{P}(\cdot) \quad \text{a.s.}[\mathbb{P}], \quad \frac{1}{n} \sum_{i=1}^n \tilde{W}_i \delta_{\tilde{X}_i}(\cdot) \xrightarrow{w} \bar{w} \tilde{P}(\cdot) \quad \text{a.s.}[\mathbb{P}].$$

Proof. It follows that $P_n(\mathcal{D}) \geq \frac{\theta\nu(\mathcal{D})}{\theta + n\beta} > 0$ a.s. $[\mathbb{P}]$, so $\sum_{n=1}^{\infty} P_n(\mathcal{D}) = \infty$ a.s. $[\mathbb{P}]$, and thus $\mathbb{P}(X_n \in \mathcal{D} \text{ i.o.}) = 1$ by Lévy's extension to the Borel-Cantelli lemmas. As a result, $T_n < \infty$ a.s. $[\mathbb{P}]$, for $n \geq 1$. Define

$$P_n(\cdot | \mathcal{D}) := \frac{P_n(\cdot \cap \mathcal{D})}{P_n(\mathcal{D})}, \quad \text{and} \quad \tilde{P}_n(\cdot) := P_{T_{n+1}-1}(\cdot | \mathcal{D}), \quad \text{for } n \geq 1.$$

It follows \mathbb{P} -a.s. set-wise that

$$\begin{aligned} \tilde{P}_n(\cdot) &= \frac{\theta\nu(\cdot \cap \mathcal{D}) + \sum_{i=1}^{T_{n+1}-1} W_i \delta_{X_i}(\cdot \cap \mathcal{D})}{\theta\nu(\mathcal{D}) + \sum_{i=1}^{T_{n+1}-1} W_i \delta_{X_i}(\mathcal{D})} = \\ &= \frac{\theta\nu(\cdot \cap \mathcal{D}) + \sum_{k=1}^n W_{T_k} \delta_{X_{T_k}}(\cdot \cap \mathcal{D})}{\theta\nu(\mathcal{D}) + \sum_{k=1}^n W_{T_k} \delta_{X_{T_k}}(\mathcal{D})} = \\ &= \frac{\theta\nu(\mathcal{D})\nu(\cdot | \mathcal{D}) + \sum_{k=1}^n \tilde{W}_k \delta_{\tilde{X}_k}(\cdot \cap \mathcal{D})}{\theta\nu(\mathcal{D}) + \sum_{k=1}^n \tilde{W}_k \delta_{\tilde{X}_k}(\mathcal{D})} = \frac{\tilde{\theta}\tilde{\nu}(\cdot) + \sum_{k=1}^n \tilde{W}_k \delta_{\tilde{X}_k}(\cdot)}{\tilde{\theta} + \sum_{k=1}^n \tilde{W}_k}, \end{aligned}$$

where $\tilde{\theta} := \theta\nu(\mathcal{D})$ and $\tilde{\nu}(\cdot) := \nu(\cdot | \mathcal{D})$. Define $\tilde{P}_n^0(\cdot) := \frac{\tilde{\theta}\tilde{\nu}(\cdot) + \sum_{k=1}^n \tilde{W}_k \delta_{\tilde{X}_k}(\cdot)}{\tilde{\theta} + \sum_{k=1}^n \tilde{W}_k}$. Then $\tilde{P}_n^0(B)$ is $\tilde{\mathcal{F}}_n$ -measurable and $\tilde{P}_n^0(B) = \tilde{P}_n(B)$ a.s. $[\mathbb{P}]$, for $B \in \mathcal{X}$. Let $A \in \tilde{\mathcal{F}}_n$ and $B \in \mathcal{X}$. Since $\tilde{\mathcal{F}}_n \subseteq \mathcal{F}_{T_n} \subseteq \mathcal{F}_{T_{n+1}-1}$, one has

$$\begin{aligned} \mathbb{P}(A \cap \{\tilde{X}_{n+1} \in B\}) &= \mathbb{E}[\mathbb{1}_A \cdot P_{T_{n+1}-1}(B)] = \\ &= \mathbb{E}[\mathbb{1}_A \cdot P_{T_{n+1}-1}(B \cap \mathcal{D})] = \\ &= \mathbb{E}[\mathbb{1}_A \cdot \tilde{P}_n(B) P_{T_{n+1}-1}(\mathcal{D})] = \\ &= \mathbb{E}[\mathbb{1}_A \cdot \tilde{P}_n(B) \cdot \mathbb{1}_{\{X_{T_{n+1}} \in \mathcal{D}\}}] = \mathbb{E}[\mathbb{1}_A \cdot \tilde{P}_n(B)] = \mathbb{E}[\mathbb{1}_A \cdot \tilde{P}_n^0(B)]; \end{aligned}$$

thus, \tilde{P}_n^0 is a version of the conditional distribution of \tilde{X}_{n+1} given $\tilde{\mathcal{F}}_n$. Denote by $\tilde{N}_n := \tilde{\theta} + \sum_{i=1}^n \tilde{W}_i$, for $n \geq 1$. As $\mathbb{E}[\tilde{W}_{n+1} | \tilde{\mathcal{F}}_n] = \bar{w}$, for each $n \geq 1$, and $\sum_{n=1}^{\infty} \mathbb{E}[\tilde{W}_n^2]/n^2 < \infty$, then

$$\frac{1}{n} \sum_{i=1}^n \tilde{W}_i \xrightarrow{a.s.} \bar{w} \quad \text{a.s.}[\mathbb{P}],$$

by Lemma A.7 in the Appendix; hence, $\tilde{N}_n/n \xrightarrow{a.s.} \bar{w}$. Let $A \in \mathcal{X} \cap \mathcal{D}$. From $x - x^2 \leq \frac{x}{1+x} \leq x$, for $0 \leq x \leq 1$, one has that

$$\mathbb{E}[\tilde{P}_{n+1}^0(A) - \tilde{P}_n^0(A) | \tilde{\mathcal{F}}_n] = \mathbb{E}\left[\frac{\tilde{W}_{n+1}}{\tilde{N}_{n+1}} (\delta_{\tilde{X}_{n+1}}(A) - \tilde{P}_n^0(A)) \middle| \tilde{\mathcal{F}}_n\right] =$$

$$\begin{aligned}
&= \mathbb{E} \left[\frac{\tilde{W}_{n+1}}{\tilde{N}_{n+1}} (1 - \tilde{P}_n^0(A)) \delta_{\tilde{X}_{n+1}}(A) \middle| \tilde{\mathcal{F}}_n \right] + \mathbb{E} \left[\frac{\tilde{W}_{n+1}}{\tilde{N}_{n+1}} (-\tilde{P}_n^0(A)) \delta_{\tilde{X}_{n+1}}(A^c) \middle| \tilde{\mathcal{F}}_n \right] = \\
&= \tilde{P}_n^0(A^c) \mathbb{E} \left[\frac{\tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A)}{\tilde{N}_n + \tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A)} \middle| \tilde{\mathcal{F}}_n \right] - \tilde{P}_n^0(A) \mathbb{E} \left[\frac{\tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A^c)}{\tilde{N}_n + \tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A^c)} \middle| \tilde{\mathcal{F}}_n \right] = \\
&= \tilde{P}_n^0(A) \tilde{P}_n^0(A^c) \left\{ \frac{1}{\tilde{P}_n^0(A)} \mathbb{E} \left[\frac{\tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A) / \tilde{N}_n}{1 + \tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A) / \tilde{N}_n} \middle| \tilde{\mathcal{F}}_n \right] - \right. \\
&\quad \left. - \frac{1}{\tilde{P}_n^0(A^c)} \mathbb{E} \left[\frac{\tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A^c) / \tilde{N}_n}{1 + \tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A^c) / \tilde{N}_n} \middle| \tilde{\mathcal{F}}_n \right] \right\} \leq \\
&\leq \tilde{P}_n^0(A) \tilde{P}_n^0(A^c) \left\{ \frac{1}{\tilde{P}_n^0(A)} \mathbb{E} \left[\frac{\tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A)}{\tilde{N}_n} \middle| \tilde{\mathcal{F}}_n \right] - \right. \\
&\quad \left. - \frac{1}{\tilde{P}_n^0(A^c)} \mathbb{E} \left[\frac{\tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A^c)}{\tilde{N}_n} - \frac{\tilde{W}_{n+1}^2 \delta_{\tilde{X}_{n+1}}(A^c)}{\tilde{N}_n^2} \middle| \tilde{\mathcal{F}}_n \right] \right\} = \\
&= \tilde{P}_n^0(A) \tilde{P}_n^0(A^c) \left\{ \frac{\bar{w}}{\tilde{P}_n^0(A) \tilde{N}_n} \mathbb{P}(\tilde{X}_{n+1} \in A | \tilde{\mathcal{F}}_n) - \frac{\bar{w}}{\tilde{P}_n^0(A^c) \tilde{N}_n} \mathbb{P}(\tilde{X}_{n+1} \in A^c | \tilde{\mathcal{F}}_n) + \right. \\
&\quad \left. + \frac{1}{\tilde{P}_n^0(A^c)} \mathbb{E} \left[\frac{\tilde{W}_{n+1}^2 \delta_{\tilde{X}_{n+1}}(A^c)}{\tilde{N}_n^2} \middle| \tilde{\mathcal{F}}_n \right] \right\} \leq \\
&\leq \tilde{P}_n^0(A) \tilde{P}_n^0(A^c) \frac{\beta^2}{\tilde{N}_n^2} \leq \frac{\beta^2}{4\tilde{N}_n^2},
\end{aligned}$$

where

$$\mathbb{E}[\tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A) | \tilde{\mathcal{F}}_n] = \mathbb{E}[\mathbb{E}[\tilde{W}_{n+1} | \tilde{\mathcal{F}}_n \vee \sigma(\tilde{X}_{n+1})] \delta_{\tilde{X}_{n+1}}(A) | \tilde{\mathcal{F}}_n] = \mathbb{E}[w(\tilde{X}_{n+1}) \delta_{\tilde{X}_{n+1}}(A) | \tilde{\mathcal{F}}_n] = \bar{w} \tilde{P}_n^0(A),$$

using that \tilde{U}_n is independent of $(\tilde{X}_1, \tilde{U}_1, \dots, \tilde{X}_{n-1}, \tilde{U}_{n-1}, \tilde{X}_n)$, and

$$\mathbb{E}[\tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A^c) | \tilde{\mathcal{F}}_n] = \mathbb{E}[w(\tilde{X}_{n+1}) \delta_{\tilde{X}_{n+1}}(A^c) | \tilde{\mathcal{F}}_n] = \bar{w} \tilde{P}_n^0(A^c).$$

As a result,

$$\sum_{n=1}^{\infty} \mathbb{E}[\tilde{P}_{n+1}^0(A) - \tilde{P}_n^0(A) | \tilde{\mathcal{F}}_n] \leq \frac{\beta^2}{4} \sum_{n=1}^{\infty} \frac{1}{n^2} \left(\frac{n}{\tilde{N}_n} \right)^2 < \infty \quad \text{a.s.}[\mathbb{P}].$$

On the other hand, it holds \mathbb{P} -a.s. that

$$\sum_{n=1}^{\infty} \mathbb{E}[(\tilde{P}_{n+1}^0(A) - \tilde{P}_n^0(A))^2 | \tilde{\mathcal{F}}_n] = \sum_{n=1}^{\infty} \mathbb{E} \left[\frac{\tilde{W}_{n+1}^2}{\tilde{N}_{n+1}^2} (\delta_{\tilde{X}_{n+1}}(A) - \tilde{P}_n^0(A))^2 \middle| \tilde{\mathcal{F}}_n \right] \leq 2\beta^2 \sum_{n=1}^{\infty} \frac{1}{n^2} \left(\frac{n}{\tilde{N}_n} \right)^2 < \infty;$$

thus, by Lemma A.8 in the Appendix, $(\tilde{P}_n^0(A))_{n \geq 1}$ converges a.s.[\mathbb{P}]. As a consequence, $(\mathbb{E}[f(\tilde{X}_{n+1}) | \tilde{\mathcal{F}}_n])_{n \geq 1}$ converges a.s.[\mathbb{P}], for each $f \in C_b(\mathcal{D})$ (see proof to Theorem 4.2.1). By Lemma 2.4 in Berti et al. (2004), there exists an $\tilde{P} \in \mathbb{K}_P(\Omega, \mathcal{D})$ such that $\tilde{P}_n^0 \xrightarrow{w} \tilde{P}$ a.s.[\mathbb{P}]. By Lemma A.7 in the Appendix, one has that $\frac{1}{n} \sum_{i=1}^n \delta_{\tilde{X}_i} \xrightarrow{w} \tilde{P}$ a.s.[\mathbb{P}]. Let $f \in C_b(\mathcal{D})$. It follows that

$$\frac{1}{n} \sum_{i=1}^n \tilde{W}_i f(\tilde{X}_i) = \frac{\theta + \sum_{i=1}^n \tilde{W}_i}{n} \frac{\sum_{i=1}^n \tilde{W}_i f(\tilde{X}_i)}{\theta \mathbb{E}[f(\tilde{X}_1)] + \sum_{i=1}^n \tilde{W}_i f(\tilde{X}_i)} \mathbb{E}[f(\tilde{X}_{n+1}) | \tilde{\mathcal{F}}_n] \xrightarrow{w} \bar{w} \int_{\mathbb{X}} f(x) \tilde{P}(dx) \quad \text{a.s.}[\mathbb{P}];$$

therefore, $\frac{1}{n} \sum_{i=1}^n \tilde{W}_i \delta_{\tilde{X}_i} \xrightarrow{w} \bar{w} \tilde{P}$ a.s.[\mathbb{P}]. \square

It follows from the results in Chapter I, Section 1.3 that

$$\mathbb{P}(\tilde{X}_{n+1} \in \cdot | \tilde{X}_1, \dots, \tilde{X}_n) \xrightarrow{w} \tilde{P}(\cdot) \quad \text{a.s.}[\mathbb{P}];$$

thus, \tilde{X} is asymptotically exchangeable with directing measure \tilde{P} . As \tilde{X} is generally not conditionally identically distributed (c.i.d.), we cannot apply the same reasoning as in Theorem 3.2.1 to show that \tilde{P} is a mixture between ν and a discrete measure. However, in the particular case that not only the mean but the whole distribution of the \tilde{W}_n 's coincides, the process \tilde{X} becomes trivially a c.i.d. RRPS and satisfies assumption (A.1) from Chapter III. Then $\tilde{\theta}/(\tilde{\theta} + \sum_{n=1}^{\infty} \tilde{W}_n) = 0$ a.s.[\mathbb{P}], so if one assumes further that ν is diffuse, it follows $\tilde{P} = \sum_k p_k^* \delta_{\tilde{X}_k^*}$ a.s.[\mathbb{P}] by Theorem 3.2.1, where the \tilde{X}_k^* 's are the distinct values of \tilde{X} in order of appearance and

$$p_k^* = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \delta_{\tilde{X}_i}(\{\tilde{X}_k^*\}) \quad \text{a.s.}[\mathbb{P}].$$

In the extreme case that the dominant weights are deterministic, i.e. $\tilde{W}_n = M$, for some $M \in \mathbb{R}_+$, the sequence $(\mathbb{P}(\tilde{X}_{n+1} \in \cdot | \tilde{\mathcal{F}}_n))_{n \geq 1}$ corresponds to the predictive distributions of a Pólya sequence up to a multiplicative constant, and hence

$$\tilde{P} \sim \text{DP}(M^{-1}\tilde{\theta}, \tilde{\nu}).$$

4.2.4 DPS with a discontinuity at \mathcal{D}

Predictive and empirical a.s. weak convergence for the original sequence can be realized at least when $P_n(\mathcal{D}) \xrightarrow{a.s.} 1$, which is not granted under the most general conditions as \tilde{P} might allocate some of its mass on the boundary of \mathcal{D} . Indeed, in the notation of Theorem 4.2.4,

$$\mathbb{E}[f(X_{n+1}) | \mathcal{F}_n] = P_n(\mathcal{D}) \int_{\mathbb{X}} f(x) P_n(dx | \mathcal{D}) + P_n(\mathcal{D}^c) \int_{\mathbb{X}} f(x) P_n(dx | \mathcal{D}^c) \longrightarrow \int_{\mathbb{X}} f(x) \tilde{P}(dx) \quad \text{a.s.}[\mathbb{P}],$$

for each $f \in C_b(\mathbb{X})$, provided $P_n(\mathcal{D}) \xrightarrow{a.s.} 1$, where the second term vanishes because f is bounded, and $\lim_{n \rightarrow \infty} \int_{\mathbb{X}} f(x) P_n(dx | \mathcal{D}) = \lim_{n \rightarrow \infty} \int_{\mathbb{X}} f(x) \tilde{P}_n(dx)$ a.s.[\mathbb{P}] since $(P_n(\cdot | \mathcal{D}))_{n \geq 1}$ is \mathbb{P} -a.s. a sequence of the form

$$P_{T_1-1}(\cdot | \mathcal{D}), \dots, P_{T_1-1}(\cdot | \mathcal{D}), P_{T_2-1}(\cdot | \mathcal{D}), \dots, P_{T_{n-1}-1}(\cdot | \mathcal{D}), P_{T_{n-1}-1}(\cdot | \mathcal{D}), \dots, P_{T_{n-1}-1}(\cdot | \mathcal{D}), P_{T_{n+1}-1}(\cdot | \mathcal{D}), \dots,$$

where each $P_{T_n-1}(\cdot | \mathcal{D})$ term appears $T_n - T_{n-1}$ times, which is \mathbb{P} -a.s. finite. One can then show that

$$P_n \xrightarrow{w} \tilde{P} \quad \text{a.s.}[\mathbb{P}], \quad \hat{P}_n \xrightarrow{w} \tilde{P} \quad \text{a.s.}[\mathbb{P}], \quad \frac{1}{n} \sum_{i=1}^n W_i \delta_{X_i} \xrightarrow{w} \bar{w} \tilde{P} \quad \text{a.s.}[\mathbb{P}].$$

As a consequence, the sequence $(X_n)_{n \geq 1}$ is directed by a random probability measure that is concentrated on a small subset of all possible doses, thereby achieving a sparse structure in the limit.

A trivial example when that occurs is if $\mathbb{X} = \mathcal{D}$, so that $P_n(\mathcal{D}) = 1$, for $n \geq 1$. A not so trivial case involves the violation of the continuity assumption on w at the dominant set \mathcal{D} . Note that continuity of w was used with (IV.2) to show that \mathcal{D} is non-empty (and closed) and was needed by the auxillary process from the

proof of Theorem 4.2.2, which could, nonetheless, be constructed with a continuous dominance function. With these considerations in mind, we have the following result.

Proposition 4.2.5. *Let $X = (X_n)_{n \geq 1}$ be a dominant Pólya sequence with parameters θ and ν , dominant set $\mathcal{D} \neq \emptyset$, and dominance function w such that $\bar{w} > \bar{w}^c := \sup_{x \in \mathcal{D}^c} w(x)$. Suppose $\nu(\mathcal{D}) > 0$. Then*

$$P_n(\mathcal{D}) \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}].$$

Proof. Denote by $N_n(\cdot) := \theta\nu(\cdot) + \sum_{i=1}^n W_i \delta_{X_i}(\cdot)$ and $N_n := \theta + \sum_{i=1}^n W_i$, for $n \geq 1$. As seen in the proof of Theorem 4.2.2,

$$\begin{aligned} \mathbb{E}[P_{n+1}(\mathcal{D}) - P_n(\mathcal{D}) | \mathcal{F}_n] &= \mathbb{E} \left[\frac{W_{n+1}}{N_{n+1}} (\delta_{X_{n+1}}(\mathcal{D}) - P_n(\mathcal{D})) \middle| \mathcal{F}_n \right] \geq \\ &\geq P_n(\mathcal{D}) P_n(\mathcal{D}^c) \left\{ \frac{1}{P_n(\mathcal{D})} \mathbb{E} \left[\frac{W_{n+1} \delta_{X_{n+1}}(\mathcal{D})}{N_n} \middle| \mathcal{F}_n \right] - \frac{1}{P_n(\mathcal{D}^c)} \mathbb{E} \left[\frac{W_{n+1} \delta_{X_{n+1}}(\mathcal{D}^c)}{N_n} \middle| \mathcal{F}_n \right] - \right. \\ &\quad \left. - \frac{1}{P_n(\mathcal{D})} \mathbb{E} \left[\frac{W_{n+1}^2 \delta_{X_{n+1}}(\mathcal{D})}{N_n^2} \middle| \mathcal{F}_n \right] \right\} \geq \\ &\geq P_n(\mathcal{D}) P_n(\mathcal{D}^c) \left\{ \frac{\bar{w} - \bar{w}^c}{N_n} - \frac{\beta^2}{N_n^2} \right\}; \end{aligned}$$

therefore, $(P_n(\mathcal{D}))_{n \geq 1}$ is eventually a bounded \mathcal{F} -submartingale and converges a.s. $[\mathbb{P}]$ and in L^1 to some finite random variable, say, $\tilde{p}_{\mathcal{D}}$. On the other hand,

$$\begin{aligned} \mathbb{E} \left[\frac{N_{n+1}(\mathcal{D}^c)}{N_{n+1}(\mathcal{D})} \middle| \mathcal{F}_n \right] &= \mathbb{E} \left[\frac{N_n(\mathcal{D}^c) + W_{n+1} \delta_{X_{n+1}}(\mathcal{D}^c)}{N_n(\mathcal{D}) + W_{n+1}} \middle| \mathcal{F}_n \right] + \mathbb{E} \left[\frac{N_n(\mathcal{D}^c)}{N_n(\mathcal{D}) + W_{n+1}} \delta_{X_{n+1}}(\mathcal{D}) \middle| \mathcal{F}_n \right] = \\ &= \frac{N_n(\mathcal{D}^c)}{N_n(\mathcal{D})} + \frac{\mathbb{E}[W_{n+1} \delta_{X_{n+1}}(\mathcal{D}^c) | \mathcal{F}_n]}{N_n(\mathcal{D})} - \frac{N_n(\mathcal{D}^c)}{N_n(\mathcal{D})} \mathbb{E} \left[\frac{W_{n+1}}{N_n(\mathcal{D}) + W_{n+1}} \delta_{X_{n+1}}(\mathcal{D}) \middle| \mathcal{F}_n \right] \leq \\ &\leq \frac{N_n(\mathcal{D}^c)}{N_n(\mathcal{D})} \left\{ 1 + \frac{1}{N_n} \left\{ \bar{w}^c - \bar{w} \frac{N_n(\mathcal{D})}{N_n(\mathcal{D}) + \beta} \right\} \right\}. \end{aligned}$$

Since $P_n(\mathcal{D}) \geq \frac{\theta\nu(\mathcal{D})}{\theta+n\beta} > 0$, then $\sum_{n=1}^{\infty} P_n(\mathcal{D}) = \infty$ a.s. $[\mathbb{P}]$, so $\sum_{n=1}^{\infty} \delta_{X_n}(\mathcal{D}) = \infty$ a.s. $[\mathbb{P}]$ by Levy's extension to the Borel-Cantelli lemmas. It follows that $N_n(\mathcal{D}) \xrightarrow{\text{a.s.}} \infty$, and thus $(\frac{N_n(\mathcal{D}^c)}{N_n(\mathcal{D})})_{n \geq 1}$ is eventually a non-negative supermartingale with $\lim_{n \rightarrow \infty} \frac{N_n(\mathcal{D}^c)}{N_n(\mathcal{D})} < \infty$ a.s. $[\mathbb{P}]$. As a result,

$$\limsup_{n \rightarrow \infty} \frac{P_n(\mathcal{D}^c)}{P_n(\mathcal{D})} < \infty \quad \text{a.s.}[\mathbb{P}];$$

thus, $\liminf_{n \rightarrow \infty} \frac{P_n(\mathcal{D})}{P_n(\mathcal{D}^c)} > 0$ a.s. $[\mathbb{P}]$, so $\tilde{p}_{\mathcal{D}} > 0$ a.s. $[\mathbb{P}]$. Let $\epsilon \in (0, 1)$. Define

$$T^{(1)} := \inf \{ n \in \mathbb{N} : P_n(\mathcal{D})(1 - P_n(\mathcal{D})) \leq \epsilon \}, \quad \text{and} \quad T_n^{(1)} := T^{(1)} \wedge n, \quad \text{for } n \geq 1.$$

Then $T^{(1)}$ is an \mathcal{F} -stopping time and $\{T^{(1)} = \infty\} \subseteq \{T_n^{(1)} \geq k\} \in \mathcal{F}_{k-1}$, for $k = 1, \dots, n$. It follows as in the proof of Theorem 4.2.1 and from the above results that

$$1 \geq \mathbb{E}[P_{T_n^{(1)}}(\mathcal{D})] \geq \mathbb{E} \left[\sum_{k=1}^n \mathbb{E}[P_k(\mathcal{D}) - P_{k-1}(\mathcal{D}) | \mathcal{F}_{k-1}] \cdot \mathbb{1}_{\{T^{(1)} = \infty\}} \right] \geq$$

$$\geq \epsilon \cdot \sum_{k=1}^n \frac{1}{k} \mathbb{E} \left[\frac{k}{N_{k-1}} \left\{ \bar{w} - \bar{w}^c - \frac{\beta^2}{N_{k-1}} \right\} \cdot \mathbb{1}_{\{T^{(1)}=\infty\}} \right].$$

But $N_n \xrightarrow{a.s.} \infty$ and $\frac{n}{N_{n-1}} \geq \frac{1}{\theta/n+\beta}$, so from the generalized Fatou's lemma,

$$\liminf_{n \rightarrow \infty} \mathbb{E} \left[\frac{n}{N_{n-1}} \left\{ \bar{w} - \bar{w}^c - \frac{\beta^2}{N_{n-1}} \right\} \cdot \mathbb{1}_{\{T^{(1)}=\infty\}} \right] \geq \frac{\bar{w} - \bar{w}^c}{\beta} \cdot \mathbb{P}(T^{(1)} = \infty).$$

Suppose, by contradiction, that $\mathbb{P}(T^{(1)} = \infty) > 0$. Then

$$1 \geq \epsilon \cdot \sum_{k=1}^n \frac{1}{k} \mathbb{E} \left[\frac{k}{N_{k-1}} \left\{ \bar{w} - \bar{w}^c - \frac{\beta^2}{N_{k-1}} \right\} \cdot \mathbb{1}_{\{T^{(1)}=\infty\}} \right] \rightarrow \infty,$$

absurd; hence, $\mathbb{P}(T^{(1)} = \infty) = 0$. Arguing as in Theorem 4.2.1 through the strong Markov property, one has

$$\mathbb{P}(P_n(\mathcal{D})(1 - P_n(\mathcal{D})) < \epsilon \text{ i.o.}) = 1.$$

As this is true for any $\epsilon > 0$, it follows that $P_n(\mathcal{D})(1 - P_n(\mathcal{D})) \xrightarrow{p} 0$, and thus $\tilde{p}_{\mathcal{D}}(1 - \tilde{p}_{\mathcal{D}}) = 0$ a.s.[\mathbb{P}]. But $\tilde{p}_{\mathcal{D}} > 0$ a.s.[\mathbb{P}]; therefore, $\tilde{p}_{\mathcal{D}} = 1$ a.s.[\mathbb{P}]. \square

As the conclusions of Theorem 4.2.4 are true, irrespective of the continuity of w , it follows under the conditions of Proposition 4.2.5 that

$$P_n \xrightarrow{w} \tilde{P} \quad \text{a.s.}[\mathbb{P}], \quad \hat{P}_n \xrightarrow{w} \tilde{P} \quad \text{a.s.}[\mathbb{P}], \quad \frac{1}{n} \sum_{i=1}^n W_i \delta_{X_i} \xrightarrow{w} \bar{w} \tilde{P} \quad \text{a.s.}[\mathbb{P}].$$

From Lemma A.7 in the Appendix and Theorem 4.2.2, one has further that

$$\frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\mathcal{D}) \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}], \quad \frac{1}{n} \sum_{i=1}^n W_i \longrightarrow \bar{w} \quad \text{a.s.}[\mathbb{P}], \quad \frac{1}{n} \sum_{i=1}^n W_i \delta_{X_i}(\mathcal{D}) \longrightarrow \bar{w} \quad \text{a.s.}[\mathbb{P}].$$

Denote by $\mathcal{D}(\mathcal{D}^c) := \{x \in \mathcal{D}^c : w(x) = \bar{w}^c\}$. If $\mathcal{D}(\mathcal{D}^c) \neq \emptyset$ and $\nu(\mathcal{D}(\mathcal{D}^c)) > 0$, then Theorem 4.2.2, as applied to the subsequence in \mathcal{D}^c , implies $\frac{1}{\sum_{i=1}^n \delta_{X_i}(\mathcal{D}^c)} \sum_{i=1}^n W_i \delta_{X_i}(\mathcal{D}^c) \xrightarrow{a.s.} \bar{w}^c$. As a consequence, the following result, whose proof is identical to that of Proposition 4.2.3, is true

$$\frac{1}{n^\gamma} \sum_{i=1}^n \delta_{X_i}(\mathcal{D}^c) \longrightarrow \begin{cases} 0 & \gamma > \bar{w}^c/\bar{w}, \\ \infty & \gamma < \bar{w}^c/\bar{w} \end{cases} \quad \text{a.s.}[\mathbb{P}].$$

4.2.5 Random partition

Given that non-dominant doses are administered less frequently over time, it would be informative to investigate the clustering behavior of the process as a whole. For that purpose, define the sequence $(L_n)_{n \geq 1}$ by $L_1 := 1$ and

$$L_n := \max \{k \in \{1, \dots, n\} : X_k \notin \{X_1, \dots, X_{k-1}\}\}, \quad \text{for } n \geq 2,$$

that counts the number of distinct doses that have been already tried. In fact, $(L_n)_{n \geq 1}$ denotes the length of the random partition, induced by the observations at each stage n of the experiment. It follows that the clustering structure implied by a DPS has the same behavior as that of the classical Pólya sequence.

Proposition 4.2.6. *Let $X = (X_n)_{n \geq 1}$ be a dominant Pólya sequence with parameters θ , ν and w . Suppose $\mathcal{D} \subseteq \text{supp}(\nu)$ and ν is diffuse. Then*

$$\frac{L_n}{\log n} \longrightarrow \frac{\theta}{\bar{w}} \quad \text{a.s.}[\mathbb{P}].$$

Proof. Define $\theta_0 := 1$, $Z_1 := L_1 = 1$ and

$$\theta_n := \frac{\theta}{\theta + \sum_{i=1}^n W_i}, \quad \text{for } n \geq 1, \quad \text{and} \quad Z_n := L_n - L_{n-1}, \quad \text{for } n \geq 2.$$

Then $L_n = \sum_{i=1}^n Z_i = \sum_{i=1}^n \mathbb{1}_{\{Z_i=1\}}$ and $\theta_n = \mathbb{P}(L_{n+1} = L_n + 1 | \mathcal{F}_n) = \mathbb{P}(Z_{n+1} = 1 | \mathcal{F}_n)$, for $n \geq 1$. As $\theta_n \geq \theta/(\theta + \beta n)$, then $\sum_{n=1}^{\infty} \mathbb{P}(Z_{n+1} = 1 | \mathcal{F}_n) = \sum_{n=1}^{\infty} \theta_n = \infty$ a.s. $[\mathbb{P}]$, so from Lévy's extension to the Borel-Cantelli lemmas,

$$\frac{L_n}{\sum_{k=1}^n \theta_{k-1}} = \frac{\sum_{i=1}^n \mathbb{1}_{\{Z_i=1\}}}{\sum_{k=1}^n \mathbb{P}(Z_k = 1 | \mathcal{F}_{k-1})} \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}].$$

On the other hand, Lemma A.3 in the Appendix and Theorem 4.2.2 imply

$$\frac{1}{\log n} \sum_{k=1}^n \theta_{k-1} = \frac{1}{\log n} + \frac{\theta}{\log n} \sum_{k=1}^{n-1} \frac{1}{k} \left(\frac{k}{\theta + \sum_{j=1}^k W_j} \right) \longrightarrow \frac{\theta}{\bar{w}} \quad \text{a.s.}[\mathbb{P}],$$

where we have used that $\frac{1}{\log n} \sum_{k=1}^n \frac{1}{k} \rightarrow 1$. Combining both results,

$$\frac{L_n}{\log n} \longrightarrow \frac{\theta}{\bar{w}} \quad \text{a.s.}[\mathbb{P}].$$

□

It follows for large n that the number of clusters is approximately $\bar{w}^{-1} \theta \log n$; thus, regardless of the weighting process, the clustering behavior of any DPS is ultimately governed by the parameters θ and \bar{w} . Moreover, we have from Lemma A.3 in the Appendix that

$$\begin{aligned} \frac{1}{\log n} \sum_{k=1}^n \theta_{k-1} (1 - \theta_{k-1}) &= \frac{\theta}{\log n} \sum_{k=1}^{n-1} \frac{\sum_{i=1}^k W_i}{(\theta + \sum_{i=1}^k W_i)^2} = \\ &= \frac{\theta}{n} \sum_{k=1}^{n-1} \left(\frac{\sum_{i=1}^k W_i}{\theta + \sum_{i=1}^k W_i} \right)^2 \frac{k}{\sum_{i=1}^k W_i} \frac{1}{k} \longrightarrow \frac{\theta}{\bar{w}} \quad \text{a.s.}[\mathbb{P}]. \end{aligned}$$

In that case, if we assume further that $W_n = \bar{h}(U_n)$, for some $h \in M_+(\mathcal{B}[0, 1])$, then $(X_n)_{n \geq 1}$ becomes c.i.d., and thus Theorem 5.1 in Bassetti et al. (2010) implies (see also the discussion in Section 3.2.2, Chapter III)

$$\frac{L_n - \sum_{k=1}^n \theta_{k-1}}{\log n} \xrightarrow{\text{stably}} \mathcal{N}\left(0, \frac{\theta}{\bar{w}}\right).$$

4.3 Central limit theorems

In this section we present second-order convergence properties of DPSs in terms of stable and a.s. conditional convergence (see Chapter I, Section 1.3), which would inform us on the set-wise rates of convergence of the predictive and empirical distributions to the common random limit. The first theorem, which is based on the findings in Section 4.2.3, discusses convergence with respect to the dominant subsequence, $\tilde{X} = (\tilde{X}_n)_{n \geq 1}$, and implies, in particular, that the random limit has a non-atomic distribution. Extending these results to the original sequence of observations is shown to be generally not possible, save for some peculiar cases.

4.3.1 Central limit theorem for dominant subsequence

The proof of Theorem 4.2.4 contained within itself the fact that

$$\mathbb{P}(\tilde{X}_{n+1} \in A | \tilde{\mathcal{F}}_n) \longrightarrow \tilde{p}_A \quad \text{a.s.}[\mathbb{P}],$$

for each $A \in \mathcal{X} \cap \mathcal{D}$ and some $\tilde{p}_A \in M(\mathcal{H})$, where $\tilde{\mathcal{F}} = (\tilde{\mathcal{F}}_n)_{n \geq 0}$ is the filtration generated by (\tilde{X}, \tilde{W}) , with $\tilde{W} = (\tilde{W}_n)_{n \geq 1}$ the weighting process of the dominant subsequence. It follows from Lemma A.7 in the Appendix that

$$\frac{1}{n} \sum_{i=1}^n \delta_{\tilde{X}_i}(A) \longrightarrow \tilde{p}_A \quad \text{a.s.}[\mathbb{P}].$$

As the empirical and predictive distributions of \tilde{X} converge to the same random limit, it makes sense to investigate the rate, with which they approach each other; for that purpose, we study the asymptotic behavior of

$$C_n(A) := \sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n \delta_{\tilde{X}_i}(A) - \mathbb{P}(\tilde{X}_{n+1} \in A | \tilde{\mathcal{F}}_n) \right),$$

and

$$D_n(A) := \sqrt{n} (\mathbb{P}(\tilde{X}_{n+1} \in A | \tilde{\mathcal{F}}_n) - \tilde{p}_A).$$

The next theorem, whose proof incorporates ideas and techniques from Berti et al. (2010), Berti et al. (2011), Crimaldi et al. (2007) and Crimaldi (2009), requires the existence of second-order predictive weight limits in order for $C_n(A)$ and $D_n(A)$ to converge. As a by-product, we show that $(\mathbb{P}(\tilde{X}_{n+1} \in A | \tilde{\mathcal{F}}_n))_{n \geq 1}$ is, in fact, a uniformly integrable quasi- $\tilde{\mathcal{F}}$ -martingale.

Theorem 4.3.1. *Let $X = (X_n)_{n \geq 1}$ be a dominant Pólya sequence with parameters θ , ν and w . Suppose $\nu(\mathcal{D}) > 0$. If $\lim_{n \rightarrow \infty} \mathbb{E}[\tilde{W}_{n+1}^2 \delta_{\tilde{X}_{n+1}}(A_j) | \tilde{\mathcal{F}}_n] = q_{A_j}$ a.s. $[\mathbb{P}]$, for $j = 1, 2$ with $A_1 = A$ and $A_2 = A^c$, then*

$$C_n(A) \xrightarrow{\text{stably}} \mathcal{N}(0, U(A)),$$

and

$$D_n(A) \xrightarrow{\text{a.s. cond.}} \mathcal{N}(0, V(A)) \quad \text{w.r.t. } \tilde{\mathcal{F}},$$

where

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

Proof. Denote by

$$\tilde{P}_n(\cdot) := \mathbb{P}(\tilde{X}_{n+1} \in \cdot | \tilde{\mathcal{F}}_n), \quad \tilde{N}_n(\cdot) := \tilde{\theta} \tilde{\nu}(\cdot) + \sum_{i=1}^n \tilde{W}_i \delta_{\tilde{X}_i}(\cdot), \quad \tilde{N}_n := \tilde{\theta} + \sum_{i=1}^n \tilde{W}_i \quad H_n := \{2\tilde{N}_n \geq n\bar{w}\},$$

for $n \geq 0$, where $\sum_{i=1}^0 a_i = 0$. It follows from Theorem 4.2.4 that $\frac{1}{n}\tilde{N}_n \xrightarrow{a.s.} \bar{w}$, and thus $\mathbb{P}(H_n^c \text{ i.o.}) = 0$.

Part I: $D_n(A) \xrightarrow{a.s.\text{cond.}} \mathcal{N}(0, V(A))$. The first part of the proof is based on a variant of Proposition 1 in (Berti et al., 2011).

Proposition: *Let $(Y_n)_{n \geq 1}$ be a sequence of real-valued integrable random variables on $(\Omega, \mathcal{H}, \mathbb{P})$, $\mathcal{F} = (\mathcal{F}_n)_{n \geq 0}$ be a filtration on (Ω, \mathcal{H}) , and $H_n \in \mathcal{F}_n$ be such that $\mathbb{P}(H_n^c \text{ i.o.}) = 0$. Denote by $Z_n := \mathbb{E}[Y_{n+1} | \mathcal{F}_n]$, for $n \geq 1$. If $(Z_n)_{n \geq 1}$ is uniformly integrable,*

- (i) $\sum_{n=1}^{\infty} \sqrt{n} \cdot \mathbb{E}[\mathbb{1}_{H_{n-1}} \mathbb{E}[Z_n - Z_{n-1} | \mathcal{F}_{n-1}]] < \infty$;
- (ii) $\mathbb{E}[\sup_{n \in \mathbb{N}} \sqrt{n} \cdot \mathbb{1}_{H_{n-1}} |Z_{n-1} - Z_n|] < \infty$;
- (iii) $n \cdot \sum_{k \geq n} (Z_{k-1} - Z_k)^2 \xrightarrow{a.s.} V$, for some $V \in M_+(\mathcal{H})$;

then $Z_n \rightarrow Z$ a.s. $[\mathbb{P}]$ and in L^1 , for some $Z \in M(\mathcal{H})$, and

$$\sqrt{n}(Z_n - Z) \xrightarrow{a.s.\text{cond.}} \mathcal{N}(0, V) \quad \text{w.r.t. } \mathcal{F}.$$

where in our case,

$$Y_n = \mathbb{1}_A(\tilde{X}_n), \quad Z_n = \tilde{P}_n(A), \quad V = V(A), \quad \mathcal{F}_n = \tilde{\mathcal{F}}_n.$$

It was shown in the proof of Theorem 4.2.4 that

$$\mathbb{E}[\tilde{P}_{n+1}(A) - \tilde{P}_n(A) | \tilde{\mathcal{F}}_n] = \tilde{P}_n(A) \tilde{P}_n(A^c) \left\{ \frac{1}{\tilde{P}_n(A)} \mathbb{E} \left[\frac{\tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A) / \tilde{N}_n}{1 + \tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A) / \tilde{N}_n} \middle| \tilde{\mathcal{F}}_n \right] - \frac{1}{\tilde{P}_n(A^c)} \mathbb{E} \left[\frac{\tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A^c) / \tilde{N}_n}{1 + \tilde{W}_{n+1} \delta_{\tilde{X}_{n+1}}(A^c) / \tilde{N}_n} \middle| \tilde{\mathcal{F}}_n \right] \right\}.$$

As $x - x^2 \leq \frac{x}{1+x} \leq x$, for $0 \leq x \leq 1$, one has $|\mathbb{E}[\tilde{P}_{n+1}(A) - \tilde{P}_n(A) | \tilde{\mathcal{F}}_n]| \leq \frac{\beta^2}{4\tilde{N}_n^2}$. Then

$$\sum_{n=1}^{\infty} \sqrt{n} \cdot \mathbb{E}[\mathbb{1}_{H_{n-1}} |\tilde{P}_n(A) - \tilde{P}_{n-1}(A) | \tilde{\mathcal{F}}_{n-1}] \leq \frac{\beta^2}{4} \sum_{n=1}^{\infty} \frac{1}{n^{3/2}} \mathbb{E} \left[\mathbb{1}_{H_{n-1}} \left(\frac{n}{\tilde{N}_{n-1}} \right)^2 \right] < \infty,$$

since $n/\tilde{N}_n \leq 2/\bar{w}$ on H_n . On the other hand,

$$|\tilde{P}_n(A) - \tilde{P}_{n-1}(A)| = \left| \frac{\tilde{W}_n}{\tilde{N}_n} (\delta_{\tilde{X}_n}(A) - \tilde{P}_{n-1}(A)) \right| \leq \frac{2\beta}{\tilde{N}_n};$$

hence,

$$\mathbb{E} \left[\sup_{n \in \mathbb{N}} n^2 \cdot \mathbb{1}_{H_{n-1}} (\tilde{P}_n(A) - \tilde{P}_{n-1}(A))^4 \right] \leq 16\beta^4 \sum_{n=1}^{\infty} \frac{1}{n^2} \mathbb{E} \left[\mathbb{1}_{H_{n-1}} \left(\frac{n}{\tilde{N}_{n-1}} \right)^4 \right] < \infty,$$

by the same token. Since $\delta_{X_n}(A) \cdot \delta_{X_n}(A^c) = 0$, it follows

$$\mathbb{E}[(\tilde{P}_n(A) - \tilde{P}_{n-1}(A))^2 | \tilde{\mathcal{F}}_{n-1}] = (\tilde{P}_{n-1}(A^c))^2 \mathbb{E} \left[\frac{\tilde{W}_n^2 \delta_{\tilde{X}_n}(A)}{\tilde{N}_n^2} \middle| \tilde{\mathcal{F}}_{n-1} \right] + (\tilde{P}_{n-1}(A))^2 \mathbb{E} \left[\frac{\tilde{W}_n^2 \delta_{\tilde{X}_n}(A^c)}{\tilde{N}_n^2} \middle| \tilde{\mathcal{F}}_{n-1} \right].$$

As $\frac{n^2}{(\tilde{N}_{n-1} + \beta)^2} \mathbb{E}[\tilde{W}_n^2 \delta_{\tilde{X}_n}(A_i) | \tilde{\mathcal{F}}_{n-1}] \leq n^2 \cdot \mathbb{E}\left[\frac{\tilde{W}_n^2 \delta_{\tilde{X}_n}(A_i)}{\tilde{N}_n^2} | \tilde{\mathcal{F}}_{n-1}\right] \leq \frac{n^2}{\tilde{N}_{n-1}^2} \mathbb{E}[\tilde{W}_n^2 \delta_{\tilde{X}_n}(A_i) | \tilde{\mathcal{F}}_{n-1}]$, for $i = 1, 2$, then

$$n^2 \cdot \mathbb{E}[(\tilde{P}_n(A) - \tilde{P}_{n-1}(A))^2 | \tilde{\mathcal{F}}_{n-1}] \xrightarrow{a.s.} V(A).$$

Define $R_n := n^2 \cdot \mathbb{1}_{H_{n-1}} (\tilde{P}_n(A) - \tilde{P}_{n-1}(A))^2$, for $n \geq 1$. Given that $H_{n-1} \in \tilde{\mathcal{F}}_{n-1}$ and $\mathbb{P}(\mathbb{1}_{H_n} = 1 \text{ ult.}) = 1$, one has $\mathbb{E}[R_n | \tilde{\mathcal{F}}_{n-1}] \xrightarrow{a.s.} V(A)$. Moreover,

$$\frac{\mathbb{E}[R_n^2]}{n^2} \leq n^2 \cdot \mathbb{E}\left[\mathbb{1}_{H_{n-1}} \frac{16\beta^4}{\tilde{N}_{n-1}^4}\right] \leq \frac{16^2 \beta^4 n^2}{\bar{w}^4 (n-1)^4};$$

thus, $\sum_{n=1}^{\infty} \mathbb{E}[R_n^2]/n^2 < \infty$, so $n \cdot \sum_{k \geq n} R_k/k^2 \xrightarrow{a.s.} V(A)$ by Lemma A.7 in the Appendix. As a result,

$$n \cdot \sum_{k \geq n} \mathbb{1}_{H_{k-1}} (\tilde{P}_k(A) - \tilde{P}_{k-1}(A))^2 = n \cdot \sum_{k \geq n} \frac{R_k}{k^2} \longrightarrow V(A) \quad \text{a.s.}[\mathbb{P}].$$

But $\mathbb{P}(\mathbb{1}_{H_n} = 1 \text{ ult.}) = 1$, so it holds

$$n \cdot \sum_{k \geq n} (\tilde{P}_k(A) - \tilde{P}_{k-1}(A))^2 \longrightarrow V(A) \quad \text{a.s.}[\mathbb{P}].$$

Part II: $C_n(A) \xrightarrow{\text{stably}} \mathcal{N}(0, U(A))$. It follows that

$$C_n(A) = \frac{1}{\sqrt{n}} \sum_{k=1}^n \left\{ \delta_{\tilde{X}_k}(A) - \tilde{P}_{k-1}(A) + k(\tilde{P}_{k-1}(A) - \tilde{P}_k(A)) \right\}.$$

Define

$$C_n^*(A) := \frac{1}{\sqrt{n}} \sum_{k=1}^n \mathbb{1}_{H_{k-1}} \left\{ \delta_{\tilde{X}_k}(A) - \tilde{P}_{k-1}(A) + k(\tilde{P}_{k-1}(A) - \tilde{P}_k(A)) \right\}.$$

As $\mathbb{P}(\mathbb{1}_{H_n} = 1 \text{ ult.}) = 1$, then $C_n(A) - C_n^*(A) \xrightarrow{a.s.} 0$, so from the properties of stable convergence, it is enough that $C_n^*(A) \xrightarrow{\text{stably}} \mathcal{N}(0, U(A))$ for the general result to hold. Note that

$$C_n^*(A) = \frac{1}{\sqrt{n}} \sum_{k=1}^n \mathbb{1}_{H_{k-1}} \left\{ \delta_{\tilde{X}_k}(A) - \tilde{P}_{k-1}(A) + k(\mathbb{E}[\tilde{P}_k(A) | \tilde{\mathcal{F}}_{k-1}] - \tilde{P}_k(A)) \right\} + Q_n,$$

where

$$Q_n := \frac{1}{\sqrt{n}} \sum_{k=1}^n k \cdot \mathbb{1}_{H_{k-1}} (\tilde{P}_{k-1}(A) - \mathbb{E}[\tilde{P}_k(A) | \tilde{\mathcal{F}}_{k-1}]).$$

Since $|\mathbb{E}[\tilde{P}_k(A) | \tilde{\mathcal{F}}_{k-1}] - \tilde{P}_{k-1}(A)| \leq \frac{\beta^2}{4\tilde{N}_{k-1}^2}$, one has

$$\begin{aligned} \mathbb{E}|Q_n| &\leq \frac{1}{\sqrt{n}} \sum_{k=1}^n k \cdot \mathbb{E}\left[\mathbb{1}_{H_{k-1}} |\mathbb{E}[\tilde{P}_k(A) | \tilde{\mathcal{F}}_{k-1}] - \tilde{P}_{k-1}(A)|\right] \leq \\ &\leq \frac{1}{\sqrt{n}} \sum_{k=1}^n k \cdot \mathbb{E}\left[\mathbb{1}_{H_{k-1}} \frac{\beta^2}{4\tilde{N}_{k-1}^2}\right] \leq \frac{\beta^2}{4\sqrt{n}} + \frac{1}{\sqrt{n}} \sum_{k=2}^n \frac{\beta^2 k}{\bar{w}^2 (k-1)^2} \approx \frac{\log n}{\sqrt{n}} \longrightarrow 0. \end{aligned}$$

Therefore, we need only to show that the following result, which has been suggested by Berti et al. (2011) and is derived from Corollary 7 in Crimaldi et al. (2007), holds.

Proposition: Let $(\mathcal{G}_n)_{n \geq 0}$ be a filtration on (Ω, \mathcal{H}) , and $M_n = (M_{n,k})_{1 \leq k \leq n}$ be a martingale w.r.t. $(\mathcal{G}_k)_{1 \leq k \leq n}$ such that $M_{n,0} = 0$. Denote by \mathcal{U} the completion of \mathcal{G}_∞ in \mathcal{H} and

$$Y_{n,k} = M_{n,k} - M_{n,k-1}.$$

If it holds

- (i) $\mathbb{E}[\max_{1 \leq k \leq n} |Y_{n,k}|] \rightarrow 0$;
- (ii) $\sum_{k=1}^n Y_{n,k}^2 \xrightarrow{p} U$, for some $U \in M_+(\mathcal{U})$;

then

$$\sum_{k=1}^n Y_{n,k} \xrightarrow{\text{stably}} \mathcal{N}(0, U).$$

where in our case,

$$Y_{n,k} = \frac{1}{\sqrt{n}} \mathbb{1}_{H_{k-1}} \left\{ \delta_{\tilde{X}_k}(A) - \tilde{P}_{k-1}(A) + k(\mathbb{E}[\tilde{P}_k(A)|\tilde{\mathcal{F}}_{k-1}] - \tilde{P}_k(A)) \right\}, \quad \mathcal{G}_n = \tilde{\mathcal{F}}_n, \quad U = U(A).$$

First note that $\mathbb{E}[Y_{n,k}|\tilde{\mathcal{F}}_{k-1}] = 0$. Regarding (i),

$$\begin{aligned} \max_{1 \leq k \leq n} |Y_{n,k}| &\leq \frac{1}{\sqrt{n}} \cdot \max_{1 \leq k \leq n} \mathbb{1}_{H_{k-1}} |\delta_{\tilde{X}_k}(A) - \tilde{P}_{k-1}(A)| + \frac{1}{\sqrt{n}} \cdot \sum_{k=1}^n k \cdot \mathbb{1}_{H_{k-1}} |\mathbb{E}[\tilde{P}_k(A)|\tilde{\mathcal{F}}_{k-1}] - \tilde{P}_{k-1}(A)| + \\ &\quad + \frac{1}{\sqrt{n}} \cdot \max_{1 \leq k \leq n} k \cdot \mathbb{1}_{H_{k-1}} |\tilde{P}_{k-1}(A) - \tilde{P}_k(A)|. \end{aligned}$$

Then $\frac{1}{\sqrt{n}} \mathbb{E}[\max_{1 \leq k \leq n} \mathbb{1}_{H_{k-1}} |\delta_{\tilde{X}_k}(A) - \tilde{P}_{k-1}(A)|] \rightarrow 0$ and $\frac{1}{\sqrt{n}} \sum_{k=1}^n k \mathbb{E}[\mathbb{1}_{H_{k-1}} |\mathbb{E}[\tilde{P}_k(A)|\tilde{\mathcal{F}}_{k-1}] - \tilde{P}_{k-1}(A)|] \rightarrow 0$ from above. In addition,

$$\frac{1}{\sqrt{n}} \cdot \mathbb{E} \left[\max_{1 \leq k \leq n} k \cdot \mathbb{1}_{H_{k-1}} |\tilde{P}_{k-1}(A) - \tilde{P}_k(A)| \right] \leq \frac{1}{\sqrt{n}} \cdot \mathbb{E} \left[\max_{1 \leq k \leq n} k \cdot \mathbb{1}_{H_{k-1}} \frac{\beta}{\tilde{N}_{k-1}} \right] \leq \frac{c}{\sqrt{n}} \rightarrow 0,$$

for some suitable constant $c \in \mathbb{R}_+$; therefore, $\mathbb{E}[\max_{1 \leq k \leq n} |Y_{n,k}|] \rightarrow 0$. Regarding (ii), write

$$\sum_{k=1}^n Y_{n,k}^2 = F_n + G_n + K_n,$$

where

$$\begin{aligned} F_n &= \frac{1}{n} \sum_{k=1}^n \mathbb{1}_{H_{k-1}} \left\{ \delta_{\tilde{X}_k}(A) - \tilde{P}_{k-1}(A) + k(\tilde{P}_{k-1}(A) - \tilde{P}_k(A)) \right\}^2, \\ G_n &= \frac{1}{n} \sum_{k=1}^n k^2 \cdot \mathbb{1}_{H_{k-1}} (\mathbb{E}[\tilde{P}_k(A)|\tilde{\mathcal{F}}_{k-1}] - \tilde{P}_{k-1}(A))^2, \end{aligned}$$

and

$$K_n = \frac{2}{n} \sum_{k=1}^n k \cdot \mathbb{1}_{H_{k-1}} \left\{ \delta_{\tilde{X}_k}(A) - \tilde{P}_{k-1}(A) + k(\tilde{P}_{k-1}(A) - \tilde{P}_k(A)) \right\} (\mathbb{E}[\tilde{P}_k(A)|\tilde{\mathcal{F}}_{k-1}] - \tilde{P}_{k-1}(A)).$$

Now,

$$F_n = \frac{1}{n} \sum_{k=1}^n \mathbb{1}_{H_{k-1}} (\delta_{\tilde{X}_k}(A) - \tilde{P}_{k-1}(A))^2 + \frac{1}{n} \sum_{k=1}^n k^2 \cdot \mathbb{1}_{H_{k-1}} (\tilde{P}_{k-1}(A) - \tilde{P}_k(A))^2 +$$

$$+ \frac{2}{n} \sum_{k=1}^n k \cdot \mathbb{1}_{H_{k-1}} (\tilde{P}_{k-1}(A) - \tilde{P}_k(A)) (\delta_{\tilde{X}_k}(A) - \tilde{P}_{k-1}(A)).$$

From Part I, $n^2 \cdot \mathbb{1}_{H_{n-1}} \mathbb{E}[(\tilde{P}_n(A) - \tilde{P}_{n-1}(A))^2 | \tilde{\mathcal{F}}_{n-1}] \xrightarrow{a.s.} V(A)$, and thus, by Lemma A.7 in the Appendix,

$$\frac{1}{n} \sum_{k=1}^n k^2 \cdot \mathbb{1}_{H_{k-1}} (\tilde{P}_{k-1}(A) - \tilde{P}_k(A))^2 \longrightarrow V(A) \quad \text{a.s.}[\mathbb{P}].$$

In addition, $\mathbb{P}(\mathbb{1}_{H_n} = 1 \text{ ult.}) = 1$ implies that

$$\frac{1}{n} \sum_{k=1}^n \mathbb{1}_{H_{k-1}} (\delta_{\tilde{X}_k}(A) - \tilde{P}_{k-1}(A))^2 \longrightarrow \tilde{p}_A(1 - \tilde{p}_A) \quad \text{a.s.}[\mathbb{P}].$$

On the other hand,

$$\begin{aligned} (\tilde{P}_{n-1}(A) - \tilde{P}_n(A)) (\delta_{\tilde{X}_n}(A) - \tilde{P}_{n-1}(A)) &= -\tilde{P}_{n-1}(A) (\tilde{P}_{n-1}(A) - \tilde{P}_n(A)) - \left(\frac{\tilde{N}_n(A)}{\tilde{N}_n} - \frac{\tilde{N}_{n-1}(A)}{\tilde{N}_{n-1}} \right) \delta_{\tilde{X}_n}(A) = \\ &= -\tilde{P}_{n-1}(A) (\tilde{P}_{n-1}(A) - \tilde{P}_n(A)) - (1 - \tilde{P}_{n-1}(A)) \frac{\tilde{W}_n \delta_{\tilde{X}_n}(A)}{\tilde{N}_n}. \end{aligned}$$

As $\frac{n}{\tilde{N}_{n-1} + \beta} \mathbb{E}[\tilde{W}_n \delta_{\tilde{X}_n}(A) | \tilde{\mathcal{F}}_{n-1}] \leq n \cdot \mathbb{E}[\tilde{W}_n \delta_{\tilde{X}_n}(A) / \tilde{N}_n | \tilde{\mathcal{F}}_{n-1}] \leq \frac{n}{\tilde{N}_{n-1}} \mathbb{E}[\tilde{W}_n \delta_{\tilde{X}_n}(A) | \tilde{\mathcal{F}}_{n-1}]$, it follows that $n \cdot \mathbb{E}[\tilde{W}_n \delta_{\tilde{X}_n}(A) / \tilde{N}_n | \tilde{\mathcal{F}}_{n-1}] \xrightarrow{a.s.} \tilde{p}_A$ by Theorem 4.2.4. Moreover, $n \cdot \tilde{P}_{n-1}(A) \mathbb{1}_{H_{n-1}} |\mathbb{E}[\tilde{P}_n(A) - \tilde{P}_{n-1}(A) | \tilde{\mathcal{F}}_{n-1}]| \leq \frac{cn}{(n-1)^2} \xrightarrow{a.s.} 0$ for some suitable constant $c \in \mathbb{R}_+$. As a result,

$$n \cdot \mathbb{1}_{H_{n-1}} \mathbb{E}[(\tilde{P}_{n-1}(A) - \tilde{P}_n(A)) (\delta_{\tilde{X}_n}(A) - \tilde{P}_{n-1}(A)) | \tilde{\mathcal{F}}_{n-1}] \longrightarrow -\tilde{p}_A(1 - \tilde{p}_A) \quad \text{a.s.}[\mathbb{P}];$$

hence, by Lemma A.7 in the Appendix,

$$\frac{2}{n} \sum_{k=1}^n k \cdot \mathbb{1}_{H_{k-1}} (\tilde{P}_{k-1}(A) - \tilde{P}_k(A)) (\delta_{\tilde{X}_k}(A) - \tilde{P}_{k-1}(A)) \longrightarrow -2\tilde{p}_A(1 - \tilde{p}_A) \quad \text{a.s.}[\mathbb{P}],$$

and, ultimately, $F_n \xrightarrow{a.s.} U(A)$. As for G_n and K_n , we have from before

$$\mathbb{E}|G_n| = \frac{1}{n} \sum_{k=1}^n k^2 \cdot \mathbb{E}[\mathbb{1}_{H_{k-1}} (\mathbb{E}[\tilde{P}_k(A) | \tilde{\mathcal{F}}_{k-1}] - \tilde{P}_{k-1}(A))^2] \leq \frac{1}{n} \sum_{k=1}^n o(k^{-1}) \longrightarrow 0,$$

that is $G_n \rightarrow 0$ in L^1 , so $G_n \xrightarrow{p} 0$, whereas $K_n^2/4 \leq F_n G_n \xrightarrow{p} 0$, and thus $K_n \xrightarrow{p} 0$. As a consequence, $\sum_{k=1}^n Y_{n,k}^2 \xrightarrow{p} U(A)$ and the conclusions of the proposition follow. □

It follows from Lemma 1 in Berti et al. (2011) that under the conditions of Theorem 4.3.1

$$(C_n(A), D_n(A)) \xrightarrow{\text{stably}} \mathcal{N}(0, U(A)) \otimes \mathcal{N}(0, V(A)),$$

which in turn implies

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n \delta_{\tilde{X}_i}(A) - \tilde{p}_A \right) = C_n(A) + D_n(A) \xrightarrow{\text{stably}} \mathcal{N}(0, U(A) + V(A)).$$

If $(\tilde{W}_n)_{n \geq 1}$ are further i.i.d., then $(\tilde{X}_n)_{n \geq 1}$ is a c.i.d. RRPS that satisfies the assumptions in Theorem 3.3.1 of Chapter III, so we can recover a central limit result for the predictive and empirical means, evaluated at any $f \in M_b(\mathcal{X})$.

Another consequence of the a.s. conditional convergence in Theorem 4.3.1 is the following proposition, which states that the distribution of \tilde{p}_A has no point masses. In particular, this implies $0 < \tilde{p}_A < 1$ a.s.[\mathbb{P}], so that $\mathcal{N}(0, V(A))$ and $\mathcal{N}(0, U(A))$ are guaranteed to be non-degenerate kernels. The proof is identical to the one that accompanies Proposition 3.3.3 of Chapter III, which is itself derived from Theorem 3.2 in Aletti et al. (2009).

Proposition 4.3.2. *Under the conditions of Theorem 4.3.1, one has $\mathbb{P}(\tilde{p}_A = p) = 0$, for all $p \in [0, 1]$.*

4.3.2 Central limit theorem for original sequence

It is more natural to have central limit results for the predictive and empirical distributions of the original sequence $(X_n)_{n \geq 1}$, yet convergence of, say, $\sqrt{n}(P_n(A) - \tilde{p}_A)$ is not always guaranteed (a counterexample follows later). However, Proposition 4.2.3 and the discussion following Proposition 4.2.5 suggest at least one case, for which the weak limit exists. To that end, define by $\mathcal{D}(\mathcal{D}^c) := \{x \in \mathcal{D} : w(x) = \bar{w}^c\}$ the dominant subset within \mathcal{D}^c , with $\bar{w}^c := \sup_{x \in \mathcal{D}^c} w(x)$.

Corollary 4.3.3. *Under the conditions of Proposition 4.2.5, if it holds further that $\bar{w} > 2\bar{w}^c$, $\mathcal{D}(\mathcal{D}^c) \neq \emptyset$ and $\nu(\mathcal{D}(\mathcal{D}^c)) > 0$, then*

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

and

$$\sqrt{n}(P_n(A) - \tilde{p}_A) \xrightarrow{\text{a.s. cond.}} \mathcal{N}(0, V(A)) \quad \text{w.r.t. } \mathcal{F},$$

where $U(A)$ and $V(A)$ are as in Theorem 4.3.1.

Remark. Note that $U(A)$ and $V(A)$ are evaluated as in Theorem 4.3.1, which means that q_{A^c} is the \mathbb{P} -a.s. limit of $\mathbb{E}[\tilde{W}_{n+1}^2 \delta_{\tilde{X}_{n+1}}(A^c \cap \mathcal{D}) | \tilde{\mathcal{F}}_n]$ and \tilde{p}_{A^c} is, in fact, $\tilde{p}_{A^c \cap \mathcal{D}}$.

Proof of Corollary 4.3.3. Denote by

$$\begin{aligned} N_n(\cdot) &:= \theta \nu(\cdot) + \sum_{i=1}^n \delta_{X_i}(\cdot), & N_n &:= \theta + \sum_{i=1}^n W_i, & M_n(\cdot) &:= \sum_{i=1}^n \delta_{X_i}(\cdot) + 1, \\ P_n(\cdot | \mathcal{D}) &:= \frac{P_n(\cdot \cap \mathcal{D})}{P_n(\mathcal{D})}, & \text{and} & & \hat{P}_n(\cdot | \mathcal{D}) &:= n \frac{\hat{P}_n(\cdot \cap \mathcal{D})}{M_n(\mathcal{D})}, \end{aligned}$$

for $n \geq 1$. Without loss of generalization, we will use M_n instead of $\sum_{i=1}^n \delta_{X_i}$. Let $A \in \mathcal{X} \cap \mathcal{D}$ satisfy the conditions of Theorem 4.3.1. It follows from Theorem 4.3.1 and the discussion prior to Proposition 4.2.5 that

$$\sqrt{M_n(\mathcal{D})}(P_n(A|\mathcal{D}) - \tilde{p}_A) \xrightarrow{a.s.cond.} \mathcal{N}(0, V(A)) \quad \text{w.r.t. } \tilde{\mathcal{F}}^*,$$

where $\tilde{\mathcal{F}}^*$ is a filtration of the form

$$\tilde{\mathcal{F}}_{T_1-1}, \dots, \tilde{\mathcal{F}}_{T_1-1}, \tilde{\mathcal{F}}_{T_2-1}, \dots, \tilde{\mathcal{F}}_{T_{n-1}-1}, \tilde{\mathcal{F}}_{T_n-1}, \dots, \tilde{\mathcal{F}}_{T_n-1}, \tilde{\mathcal{F}}_{T_{n+1}-1}, \dots,$$

where each $\tilde{\mathcal{F}}_{T_n}$ term appears $T_n - T_{n-1}$ times, which is \mathbb{P} -a.s. finite. Moreover, $M_n(\mathcal{D})/n \xrightarrow{a.s.} 1$ by Proposition 4.2.5 and Lemma A.7 in the Appendix, so from a variant of Theorem 4.2 in Fortini and Petrone (2019),

$$\sqrt{n}(P_n(A|\mathcal{D}) - \tilde{p}_A) \xrightarrow{a.s.cond.} \mathcal{N}(0, V(A)) \quad \text{w.r.t. } \tilde{\mathcal{F}}^*.$$

Using the results in the proof of Theorem 4.2.4, one can show that the dominant process is conditionally independent of what happens at non-dominant times, that is

$$\mathbb{P}(\tilde{X}_{n+m} \in \cdot | \mathcal{F}_{T_{n+1}-1+k}) = \mathbb{P}(\tilde{X}_{n+m} \in \cdot | \mathcal{F}_{T_{n+1}-1}) = \mathbb{P}(\tilde{X}_{n+m} \in \cdot | \tilde{\mathcal{F}}_n),$$

for $m \geq 1$ and $k = 0, \dots, T_{n+1} - T_n - 1$. As it holds $\tilde{p}_A = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \delta_{\tilde{X}_i}(A)$ a.s.[\mathbb{P}] from Lemma A.7 in the Appendix, then

$$\sqrt{n}(P_n(A|\mathcal{D}) - \tilde{p}_A) \xrightarrow{a.s.cond.} \mathcal{N}(0, V(A)) \quad \text{w.r.t. } \mathcal{F}.$$

On the other hand, Theorem 4.2.4 and Proposition 4.2.5 imply

$$P_n(A) = P_n(\mathcal{D})P_n(A|\mathcal{D}) + P_n(\mathcal{D}^c)P_n(A|\mathcal{D}) \longrightarrow \tilde{p}_A \quad \text{a.s.}[\mathbb{P}].$$

As $A \subseteq \mathcal{D}$, it follows that

$$\begin{aligned} \left| \sqrt{n}(P_n(A|\mathcal{D}) - \tilde{p}_A) - \sqrt{n}(P_n(A) - \tilde{p}_A) \right| &= \sqrt{n} \left| \frac{P_n(A)}{P_n(\mathcal{D})} - P_n(A) \right| = \\ &= \sqrt{n} \frac{P_n(A)}{P_n(\mathcal{D})} P_n(\mathcal{D}^c) \leq \frac{n}{N_n} \frac{N_n(\mathcal{D}^c)}{M_n(\mathcal{D}^c)} \frac{M_n(\mathcal{D}^c)}{\sqrt{n}} \longrightarrow 0 \quad \text{a.s.}[\mathbb{P}], \end{aligned}$$

where $N_n/n \xrightarrow{a.s.} \bar{w}$ from Theorem 4.2.2, $N_n(\mathcal{D}^c)/M_n(\mathcal{D}^c) \xrightarrow{a.s.} \bar{w}^c$ from Theorem 4.2.2 as applied to a \mathcal{D}^c -valued DPS, and $M_n(\mathcal{D}^c)/\sqrt{n} \xrightarrow{a.s.} 0$ from Proposition 4.2.3, the discussion after Proposition 4.2.5 and as $\bar{w}^c/\bar{w} < 1/2$ by hypothesis. Using again Theorem 4.2 in Fortini and Petrone (2019),

$$\sqrt{n}(P_n(A) - \tilde{p}_A) \xrightarrow{a.s.cond.} \mathcal{N}(0, V(A)) \quad \text{w.r.t. } \mathcal{F}.$$

In the same way, Theorem 4.3.1 implies

$$\sqrt{n}(\hat{P}_n(A|\mathcal{D}) - P_n(A|\mathcal{D})) \xrightarrow{stably} \mathcal{N}(0, U(A)).$$

It follows that

$$\begin{aligned} \sqrt{n}(\hat{P}_n(A|\mathcal{D}) - P_n(A|\mathcal{D})) - \sqrt{n}(\hat{P}_n(A) - P_n(A)) &= \\ &= \sqrt{n}(P_n(A|\mathcal{D}) - \tilde{p}_A) - \sqrt{n}(P_n(A) - \tilde{p}_A) + \sqrt{n} \cdot \hat{P}_n(A) \frac{\hat{P}_n(\mathcal{D}^c) - \frac{1}{n}}{\hat{P}_n(\mathcal{D}) + \frac{1}{n}} \longrightarrow 0 \quad \text{a.s.}[\mathbb{P}], \end{aligned}$$

where $\hat{P}_n(\mathcal{D}) \xrightarrow{a.s.} 1$ and $\hat{P}_n(A) \xrightarrow{a.s.} \tilde{p}_A$ from Proposition 4.2.5 and Lemma A.7 in the Appendix, and $\sqrt{n} \cdot \hat{P}_n(\mathcal{D}^c) \xrightarrow{a.s.} 0$ from Proposition 4.2.3, the discussion after Proposition 4.2.5 and as $\bar{w}^c/\bar{w} < 1/2$ by hypothesis. As a consequence,

$$\sqrt{n}(\hat{P}_n(A) - P_n(A)) \xrightarrow{stably} \mathcal{N}(0, U(A)).$$

□

Curiously, one has from Proposition 4.2.3 that

$$\begin{aligned} \sqrt{n}(P_n(A|\mathcal{D}) - \tilde{p}_A) - \sqrt{n}(P_n(A) - \tilde{p}_A) &= \sqrt{n}P_n(A|\mathcal{D})(1 - P_n(\mathcal{D})) = \\ &= P_n(A|\mathcal{D}) \frac{n}{N_n} \frac{N_n(\mathcal{D}^c)}{M_n(\mathcal{D}^c)} \frac{M_n(\mathcal{D}^c)}{\sqrt{n}} \longrightarrow \infty \quad \text{a.s.}[\mathbb{P}], \end{aligned}$$

whenever $\bar{w} < 2 \cdot \bar{w}^c$ and $\liminf_{n \rightarrow \infty} \mathbb{P}(\tilde{X}_{n+1} \in A | \tilde{\mathcal{F}}_n) > 0$. In those cases $\sqrt{n}(P_n(A) - \tilde{p}_A)$ fails to converge as it is, in fact, $\sqrt{n}(P_n(A|\mathcal{D}) - \tilde{p}_A)$ that is convergent.

4.4 DPS with finite number of colors

The finite-color DPS, also known as the randomly reinforced urn model (RRU), has long been a subject of interest, with some of the most important studies being Durham et al. (1998), Muliere et al. (2006), and Berti et al. (2010). In fact, the name "RRU" has been coined by Muliere et al. (2006), although Durham et al. (1998) are among the earliest to study particular cases of the basic model. On the other hand, Berti et al. (2010) are, as far as we know, first to investigate RRU's with more than one dominant color in the context of continuous responses. The goal of this section is to describe how the aforementioned papers fit within the DPS framework.

Example 4.4.1 (Durham et al., 1998). Let $\alpha_0, \alpha_1 \in (0, 1)$ be such that $\alpha_0 < \alpha_1$. The randomized Pólya urn model of Durham et al. (1998) (introduced in Durham and Yu (1990) and studied in Li et al. (1996)) is the $\{0, 1\}$ -valued DPS with weights of the form

$$W_n = \mathbb{1}_{\{U_n \leq \alpha_{X_n}\}},$$

where $(U_n)_{n \geq 1}$ is a sequence of i.i.d. $\text{Unif}[0, 1]$ random variables that is independent of $(X_n)_{n \geq 1}$. In the context of balls as treatments, α_x can be interpreted as the success probability of treatment x ; hence, according to the above structure, we reinforce the urn with one additional ball of the same color, whenever the corresponding treatment has been successful. It follows that

$$\mathbb{E}[W_n | X_n] = \alpha_{X_n};$$

thus, $\mathcal{D} = \{1\}$, so dominance is interpreted as a treatment having higher success probability. If the urn contains initially balls of color 1, then $\nu(\{1\}) > 0$, so from Proposition 4.2.5 and Lemma A.7 in the Appendix,

$$\frac{1}{n} \sum_{i=1}^n X_i = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\{1\}) \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}],$$

and

$$\mathbb{P}(X_{n+1} = 1 | \mathcal{F}_n) \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}].$$

On the other hand, Proposition 4.2.3 implies $\sum_{i=1}^n (1 - X_i)/n^\gamma \xrightarrow{\text{a.s.}} \infty$, for any $\gamma < \alpha_0/\alpha_1$. Results about k -color extensions with a unique optimal treatment can be similarly deduced.

In the case that $\alpha_0 = \alpha_1 = \alpha$, the sequence X becomes a c.i.d. RRPS, so it follows from Chapter III that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n X_i = \lim_{n \rightarrow \infty} \mathbb{P}(X_{n+1} = 1 | \mathcal{F}_n) = p_1^* \quad \text{a.s.}[\mathbb{P}],$$

and

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (1 - X_i) = \lim_{n \rightarrow \infty} \mathbb{P}(X_{n+1} = 0 | \mathcal{F}_n) = p_0^* \quad \text{a.s.}[\mathbb{P}],$$

where $p_0^* = 1 - p_1^*$, for some $[0, 1]$ -valued measurable function p_1^* such that $\mathbb{P}(p_1^* = p) = 0$, for all $p \in [0, 1]$. In other words, X has directing measure $\tilde{P} = p_0^* \delta_0 + p_1^* \delta_1$. What is more, since the weights are i.i.d., binary, and independent of $(X_n)_{n \geq 1}$, then (see Example 3.2.4)

$$\mathbb{P}(W_n = 1 \text{ i.o.}) = 1;$$

thus, Proposition 3.2.3 applies and

$$p_1^* \sim \text{Beta}(x_1, x_0 + x_1),$$

where $(x_0, x_1) \in \mathbb{R}_+^2$ is the vector indicating the initial composition of the urn. By Theorem 3.3.1, one has

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n X_i - \mathbb{P}(X_{n+1} = 1 | \mathcal{F}_n) \right) \xrightarrow{\text{stably}} \mathcal{N} \left(0, \frac{\alpha(1-\alpha)}{\alpha^2} p_1^*(1-p_1^*) \right),$$

and

$$\sqrt{n} (\mathbb{P}(X_{n+1} = 1 | \mathcal{F}_n) - p_1^*) \xrightarrow{\text{a.s. cond}} \mathcal{N} \left(0, \frac{\alpha}{\alpha^2} p_1^*(1-p_1^*) \right) \quad \text{w.r.t. } \mathcal{F}.$$

If the urn contains balls of k colors such that $\alpha_i = \alpha_1 > \alpha_j$, for all $i \in \{1, \dots, k_0\}$, $j \in \{k_0 + 1, \dots, k\}$ and some $1 \leq k_0 \leq k$, then Theorem 4.2.4 and Proposition 4.2.5 imply

$$\mathbb{P}(X_{n+1} \in A | \mathcal{F}_n) \longrightarrow \sum_{i=1}^{k_0} p_i^* \delta_i(A) \quad \text{a.s.}[\mathbb{P}], \text{ for } A \subseteq \{1, \dots, k\},$$

where, as anticipated by Li et al. (1996), the random masses are Dirichlet distributed, $(p_1^*, \dots, p_{k_0}^*) \sim \text{Dirichlet} \left(\frac{x_1}{\sum_{i=1}^{k_0} x_i}, \dots, \frac{x_{k_0}}{\sum_{i=1}^{k_0} x_i} \right)$, with $(x_1, \dots, x_k) \in \mathbb{R}_+^k$ indicating the initial composition of the urn. \square

Example 4.4.2 (Muliere et al., 2006). Let $\mu_0, \mu_1 \in \mathbb{M}_P([0, \beta])$, for some $\beta < \infty$, and denote by F_0, F_1 their respective cumulative distribution functions. The RRU of Muliere et al. (2006), which extends the scheme of Durham et al. (1998) along the lines of continuous responses, is the $\{0, 1\}$ -valued DPS with weights, given by

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

where $(U_n)_{n \geq 1}$ is a sequence of i.i.d. $\text{Unif}[0, 1]$ random variables that is independent of $(X_n)_{n \geq 1}$. Then $W_n \sim \mu_{X_n}$ and

$$\mathbb{E}[W_n | X_n] = \int_0^\beta t \mu_{X_n}(dt) =: m_{X_n}.$$

Basic convergence results for RRU's can be found in Flournoy et al. (2012). In particular, if $m_0 < m_1$, one has $\mathcal{D} = \{1\}$, and thus Proposition 4.2.5 and Proposition 4.2.3 imply

$$\mathbb{P}(X_{n+1} = 1 | \mathcal{F}_n) \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}], \quad \frac{1}{n} \sum_{i=1}^n X_i \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}] \quad \text{and} \quad \frac{1}{n^\gamma} \sum_{i=1}^n (1 - X_i) \longrightarrow \infty \quad \text{a.s.}[\mathbb{P}],$$

for all $\gamma < m_0/m_1$. It follows that the probability of administering the superior treatment, as determined by the mean response, and the number of patients treated with it both converge to one. The situation remains the same if there are $k > 2$ treatments with corresponding weight distributions μ_1, \dots, μ_k such that $m_1 > m_i$, for $i = 2, \dots, k$.

In the case $m_0 = m_1 = m$, one has $\mathcal{D} = \{0, 1\}$; thus Theorem 4.2.4 applies with $\mathbb{X} = \mathcal{D}$ and it holds

$$\mathbb{P}(X_{n+1} = 1 | \mathcal{F}_n) \longrightarrow p_1^* \quad \text{a.s.}[\mathbb{P}],$$

for some $[0, 1]$ -valued measurable function p_1^* . Moreover, $\sum_{n=1}^\infty \mathbb{E}[X_n^2]/n^2 \leq \sum_{n=1}^\infty 1/n^2 < \infty$, so by Lemma A.7 in the Appendix, one has

$$\frac{1}{n} \sum_{i=1}^n X_i \longrightarrow p_1^* \quad \text{a.s.}[\mathbb{P}].$$

Regarding a central limit result (which exists in the literature as Corollary 2 of Berti et al. (2011)), note that

$$\mathbb{E}[W_{n+1}^2 \delta_{X_{n+1}}(\{i\}) | \mathcal{F}_n] = \mathbb{E}[\mathbb{E}[W_{n+1}^2 | X_{n+1}] \delta_{X_{n+1}}(\{i\}) | \mathcal{F}_n] = \mathbb{E}[W_{n+1}^2 | X_{n+1} = i] \mathbb{P}(X_{n+1} = i | \mathcal{F}_n),$$

for $i = 0, 1$. Therefore, provided both $q_1^* = \lim_{n \rightarrow \infty} \mathbb{E}[W_n^2 | X_n = 1]$ and $q_0^* = \lim_{n \rightarrow \infty} \mathbb{E}[W_n^2 | X_n = 0]$ exist, Theorem 4.3.1 implies

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n X_i - \mathbb{P}(X_{n+1} = 1 | \mathcal{F}_n) \right) \xrightarrow{\text{stably}} \mathcal{N} \left(0, p_1^*(1 - p_1^*) \left(\frac{(1 - p_1^*)q_1^* + p_1^*q_0^*}{m^2} - 1 \right) \right),$$

and

$$\sqrt{n} (\mathbb{P}(X_{n+1} = 1 | \mathcal{F}_n) - p_1^*) \xrightarrow{\text{a.s. cond.}} \mathcal{N} \left(0, p_1^*(1 - p_1^*) \frac{(1 - p_1^*)q_1^* + p_1^*q_0^*}{m^2} \right) \quad \text{w.r.t. } \mathcal{F}.$$

It follows from Proposition 4.3.2 that $\mathbb{P}(p_1^* = p) = 0$, for all $p \in [0, 1]$, which is the original result of Aletti et al. (2009) that motivated the proof of Proposition 4.3.2.

If $\mu_0 = \mu_1 = \mu$, the process X becomes a c.i.d. RRPS that satisfies assumption (A.1) and has directing measure $\tilde{P} = (1 - p_1^*)\delta_0 + p_1^*\delta_1$. As a consequence of Theorem 3.3.1 (or from above with $q_1^* = q_0^*$), one has

$$\sqrt{n}(\mathbb{P}(X_{n+1} = 1|\mathcal{F}_n) - p_1^*) \xrightarrow{a.s.connd} \mathcal{N}\left(0, \frac{p_1^*(1-p_1^*)}{m^2}\sigma^2\right) \quad \text{w.r.t. } \mathcal{F},$$

where $\sigma^2 := \int_0^\beta t^2 \mu(dt)$. □

The paper by Berti et al. (2010) studies the k -color RRU, whose set of dominant colors is given by $\mathcal{D} = \{1, \dots, k_0\}$, for $1 < k_0 < k$. In their model specification, dominance is assumed to occur only in the limit, which is a weaker (and arguably less realistic) assumption than the one in equation (IV.2) of a DPS, which requires the elements in \mathcal{D} to dominate the other colors at each stage n of the experiment. Even though the results in Berti et al. (2010) can be recovered through minor changes in the proofs of our theorems, we present here a version of their model that is within the class of DPSs.

Example 4.4.3 (Berti et al., 2010). Let $\mu_1, \dots, \mu_k \in \mathbb{M}_P([0, \beta])$, for $k \geq 2$ and some $\beta < \infty$, be such that

$$m_1 = \dots = m_{k_0} > \max_{k_0 < j \leq k} m_j,$$

where $m_i := \int_0^\beta t \mu_i(dt)$, for $i = 1, \dots, k$, and $1 \leq k_0 \leq k$. Define $F_i([0, t]) := \mu_i([0, t])$, for $t \in [0, \beta]$ and $i = 1, \dots, k$. We consider the DPS having the weights

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

where $(U_n)_{n \geq 1}$ is a sequence of i.i.d. $\text{Unif}[0, 1]$ random variables that is independent of $(X_n)_{n \geq 1}$. Then $\mathcal{D} = \{1, \dots, k_0\}$. Example 4.4.2 already covered the situations (i) $k = 2$ with $k_0 = 1, 2$; and (ii) $k > 2$ with $k_0 = 1$. In the general case, Proposition 4.2.5 implies

$$\mathbb{P}(X_{n+1} \in \mathcal{D}|\mathcal{F}_n) \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}],$$

and as a consequence,

$$\frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\mathcal{D}) \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}], \quad \text{and} \quad \frac{1}{n^\gamma} \sum_{i=1}^n \delta_{X_i}(\mathcal{D}^c) \longrightarrow \infty \quad \text{a.s.}[\mathbb{P}],$$

for all $\gamma < \max_{k_0 < j \leq k} m_j / m_1$. Moreover,

$$\frac{1}{n} \sum_{i=1}^n W_i \longrightarrow m_1 \quad \text{a.s.}[\mathbb{P}].$$

Denote by $P_n(\cdot|A) := P_n(\cdot \cap A) / P_n(A)$, for each $A \in \mathcal{X}$ such that $P_n(A) > 0$. Fix $j \in \{1, \dots, k_0\}$. It follows from Theorem 4.2.4 that $P_n(\{j\}|\mathcal{D}) \xrightarrow{a.s.} p_j^*$, for some $[0, 1]$ -valued random variable p_j^* , and thus

$$\mathbb{P}(X_{n+1} = j|\mathcal{F}_n) = P_n(\mathcal{D})P_n(\{j\}|\mathcal{D}) + P_n(\mathcal{D}^c)P_n(\{j\}|\mathcal{D}^c) \longrightarrow p_j^* \quad \text{a.s.}[\mathbb{P}].$$

Next, note that

$$\mathbb{E}[W_{n+1}^2 \delta_{X_{n+1}}(\{j\}) | \mathcal{F}_n] = \mathbb{E}[W_{n+1}^2 | X_{n+1} = j] \mathbb{P}(X_{n+1} = j | \mathcal{F}_n).$$

Then, provided $q_j^* = \lim_{n \rightarrow \infty} \mathbb{E}[W_n^2 | X_n = j]$ exists for each $j = 1, \dots, k_0$, we can show as in the proof of Corollary 4.3.3 that

$$\sqrt{n}(\hat{P}_n(\{j\} | \mathcal{D}) - P_n(\{j\} | \mathcal{D})) \xrightarrow{\text{stably}} \mathcal{N}\left(0, \frac{p_j^*}{\bar{w}^2} \left\{ (1 - p_j^*)^2 q_j^* + p_j^* \sum_{i \leq k_0, i \neq j} q_i^* p_i^* \right\} - p_j^* (1 - p_j^*) \right),$$

and

$$\sqrt{n}(P_n(\{j\} | \mathcal{D}) - p_j^*) \xrightarrow{\text{a.s. cond.}} \mathcal{N}\left(0, \frac{p_j^*}{\bar{w}^2} \left\{ (1 - p_j^*)^2 q_j^* + p_j^* \sum_{i \leq k_0, i \neq j} q_i^* p_i^* \right\}\right) \quad \text{w.r.t. } \mathcal{F},$$

where $\hat{P}_n(\{j\} | \mathcal{D}) := \sum_{i=1}^n \delta_{X_i}(\{j\}) / (1 + \sum_{i=1}^n \delta_{X_i}(\mathcal{D}))$ is the relative frequency of j within \mathcal{D} . The latter result strengthens the kind of convergence found in Berti et al. (2010) and implies, together with Proposition 4.3.2, that $\mathbb{P}(p_j^* = p) = 0$, for all $p \in [0, 1]$. Although it holds $P_n(\{j\}) \xrightarrow{\text{a.s.}} p_j^*$, Berti et al. (2010) show that a central limit result for $\sqrt{n}(P_n(\{j\}) - p_j^*)$ is impossible. \square

4.5 Inference

In this section we demonstrate how to use the results in Theorem 4.3.1 and Corollary 4.3.3 for the inference on the random limit of the predictive and empirical distributions. In particular, we suggest procedures for the approximation of the prior and posterior distributions of \tilde{p}_A , depending on the available information.

4.5.1 Confidence intervals for the random limit

Let $A \in \mathcal{X} \cap \mathcal{D}$ be as in Theorem 4.3.1. Using the notation of Theorem 4.3.1, define

$$V_n(A) := \frac{1}{m_n^2} \left\{ (\tilde{P}_n(A^c))^2 s_n(A) + (\tilde{P}_n(A))^2 s_n(A^c) \right\},$$

where A^c is the complement of A in \mathcal{D} , and

$$\tilde{P}_n(\cdot) := \mathbb{P}(\tilde{X}_{n+1} \in \cdot | \tilde{\mathcal{F}}_n), \quad m_n := \frac{1}{n} \sum_{i=1}^n \tilde{W}_i, \quad s_n(A_j) := \frac{1}{n} \sum_{i=1}^n \tilde{W}_i^2 \delta_{\tilde{X}_i}(A_j),$$

for $j = 1, 2$, with $A_1 = A$ and $A_2 = A^c$. It follows from Theorem 4.2.4 that $m_n \xrightarrow{\text{a.s.}} \bar{w}$ and $\tilde{P}_n(A_j) \xrightarrow{\text{a.s.}} \tilde{p}_{A_j}$, for some $\tilde{p}_A, \tilde{p}_{A^c} \in M_+(\mathcal{H})$. On the other hand, we have $\mathbb{E}[\tilde{W}_{n+1}^2 \delta_{\tilde{X}_{n+1}}(A_i) | \tilde{\mathcal{F}}_n] \xrightarrow{\text{a.s.}} q_{A_j}$ by hypothesis, and thus $s_n(A_j) \xrightarrow{\text{a.s.}} q_{A_j}$ from Lemma A.7 in the Appendix. As a consequence,

$$V_n(A) \longrightarrow V(A) \quad \text{a.s.}[\mathbb{P}].$$

Given that a.s. conditional convergence implies stable convergence, it follows from Theorem 4.3.1 and the extended Slutsky's theorem (see Häusler and Luschgy, 2015, Theorem 3.7) that

$$(D_n(A), V_n(A)) \xrightarrow{\tilde{\mathcal{F}}_\infty \text{-stably}} \mathcal{N}(0, V(A)) \times \delta_{V(A)},$$

where $D_n(A) = \sqrt{n}(\tilde{P}_n(A) - \tilde{p}(A))$ and $\tilde{\mathcal{F}}_\infty := \bigvee_{n \in \mathbb{N}} \tilde{\mathcal{F}}_n$. Since $V(A) > 0$ a.s.[\mathbb{P}] from Proposition 4.3.2, it holds

$$\frac{D_n(A)}{\sqrt{V_n(A)}} \xrightarrow{d} N(0, 1).$$

One can then build asymptotic confidence intervals for \tilde{p}_A as

$$\left(\tilde{P}_n(A) - z_\alpha \sqrt{V_n(A)/n}, \tilde{P}_n(A) + z_\alpha \sqrt{V_n(A)/n} \right),$$

where z_α is the appropriate critical value given $100(1 - \alpha)\%$ confidence, with $\alpha \in (0, 1)$. Analogous computations using $C_n(A) + D_n(A)$ lead to asymptotic confidence intervals for \tilde{p}_A , which do not depend on θ and ν , namely,

$$\left(\hat{P}_n(A) - z_\alpha \sqrt{G_n(A)/n}, \hat{P}_n(A) + z_\alpha \sqrt{G_n(A)/n} \right),$$

where $\hat{P}_n(\cdot) := \frac{1}{n} \sum_{i=1}^n \delta_{\tilde{X}_i}(\cdot)$ and

$$G_n(A) = \frac{2}{m_n^2} \left\{ (\hat{P}_n(A^c))^2 s_n(A) + (\hat{P}_n(A))^2 s_n(A^c) \right\} - \hat{P}_n(A) \hat{P}_n(A^c)$$

is a consistent estimator of $V(A) + U(A)$.

4.5.2 Credible intervals for the random limit

Similarly to Chapter III, Section 3.4.2, it is possible to utilize the a.s. conditional convergence in Theorem 4.3.1 for the approximation of the posterior distribution of \tilde{p}_A given $(\tilde{X}_1, \tilde{W}_1, \dots, \tilde{X}_n, \tilde{W}_n)$. It follows from the discussion in Chapter I, Section 1.3 that

$$\mathbb{P}(D_n(A) \in \cdot | \tilde{\mathcal{F}}_n) \xrightarrow{w} \mathcal{N}(0, V(A))(\cdot) \quad \text{a.s.}[\mathbb{P}].$$

As it holds

$$V_n(A) \longrightarrow V(A) \quad \text{a.s.}[\mathbb{P}],$$

then Theorem 4.2 in Fortini and Petrone (2019) implies

$$\mathbb{P}((D_n(A), V_n(A)) \in \cdot | \tilde{\mathcal{F}}_n) \xrightarrow{w} \mathcal{N}(0, V(A)) \times \delta_{V(A)}(\cdot) \quad \text{a.s.}[\mathbb{P}].$$

Note that $V(A) > 0$ a.s.[\mathbb{P}] from Proposition 4.3.2, so $V_n(A) > 0$ a.s.[\mathbb{P}], for all but a finite number of n . As the mapping $(t, u) \mapsto tu$ from \mathbb{R}_+^2 to \mathbb{R}_+ is continuous, we have that

$$\begin{aligned} \mathbb{E} \left[f \left(\frac{D_n(A)}{\sqrt{V_n(A)}} \right) \middle| \tilde{\mathcal{F}}_n \right] &\longrightarrow \int_{\mathbb{R}_+^2} f(tu) \mathcal{N}(0, V(A)) \times \delta_{V(A)}(dt, du) = \\ &= \int_{\mathbb{R}_+} f(t \cdot V(A)^{-1}) \mathcal{N}(0, V(A))(dt) = \\ &= \int_{\mathbb{R}_+} f(s) \mathcal{N}(0, 1)(ds), \end{aligned}$$

for each $f \in C_b(\mathbb{R}_+)$. Since the cumulative distribution function of the Normal distribution is continuous, it follows \mathbb{P} -a.s. that

$$\mathbb{P}(\sqrt{n}(\tilde{P}_n(A) - \tilde{p}_A) \leq t \cdot \sqrt{V_n(A)} | \tilde{\mathcal{F}}_n) \longrightarrow \mathcal{N}(0, 1)((-\infty, t]), \quad \text{for } t \in \mathbb{R}.$$

This result allows us to obtain asymptotic credible intervals around \tilde{p}_A in the sense that

$$\mathbb{P}(\tilde{P}_n(A) - z_\alpha \sqrt{V_n(A)/n} < \tilde{p}_A < \tilde{P}_n(A) + z_\alpha \sqrt{V_n(A)/n} | \tilde{\mathcal{F}}_n) \approx 1 - \alpha,$$

for n large enough, where z_α is the appropriate critical value from the standard Normal distribution given $100(1 - \alpha)\%$ confidence, with $\alpha \in (0, 1)$.

It should be noted that the approximations both here and in the previous subsection can be extended, under the conditions of Corollary 4.3.3, to hold with the original predictive and empirical distributions. In that case, the estimator of the variance $V_n(A)$ would have the same form, but will be composed of

$$m_n = \frac{1}{\sum_{i=1}^n \delta_{X_i}(\mathcal{D})} \sum_{i=1}^n W_i \delta_{X_i}(\mathcal{D}), \quad \text{and} \quad s_n(A_j) = \frac{1}{\sum_{i=1}^n \delta_{X_i}(\mathcal{D})} \sum_{i=1}^n W_i^2 \delta_{X_i}(A_j \cap \mathcal{D}).$$

Chapter V

Further discussion

This chapter elaborates on the results obtained so far by looking at some applications of the randomly reinforced Pólya sequence (RRPS), as well as suggesting ways to extend the basic model. At first, note that the sampling mechanism of a RRPS allows for ties in the observations. This motivates the study of its clustering behavior in Section 5.1, which can be described in short as a weighted version of the Chinese Restaurant Process (see Aldous, 1985). On the other hand, the RRPS of Chapter IV is a generalization of several models that have been developed as particular proposals for randomized, response-adaptive designs for clinical trials; thus, in Section 5.2 we investigate how our infinite-color specification fits into that framework. In Section 5.3 we suggest a way to expand the predictive distributions of the RRPS by incorporating additional information from the partition. Finally, Section 5.4 considers the case of a family of dependent RRPSs and, in particular, discusses interacting RRPSs that are partially conditionally identically distributed (p.c.i.d.).

5.1 Weighted Chinese Restaurant Process

The progressive partitioning of the observations of a RRPS into clusters according to color/species/dose can be interpreted through the Chinese Restaurant Process metaphor (see Pitman, 2010, Section 3.1), which in its simplest form describes the clustering behavior of an exchangeable sequence with a Dirichlet process prior. In the latter case, the sampling procedure can actually be derived from the predictive rules of the Pólya sequence with parameters θ and ν (see Chapter II, Section 2.1), where each observation X_n signals the arrival of a new customer at the (Chinese) restaurant, who is waiting to be seated. It follows for any given time $n + 1$ that there are L_n already occupied tables, so the $n + 1$ th customer is seated randomly at one of the existing tables with a probability proportional to the number of people eating at it, or is prepared a new table with probability $\theta/(\theta + n)$. In addition, each table comes with a numbered placard that reflects its order of appearance. Exchangeability in this case implies that the partitions, resulting from any shuffling of the placards, are equally likely.

Under the RRPS framework, we make two major changes to the basic restaurant idea. First of all, implementing a weighting process changes the probabilities that assign customers to tables. In particular, the n th

customer is given a "weight" W_n (think of it as an attractivity index), so that people are drawn to a table according to its relative total attractiveness. As a result, the placards are no longer exchangeable between tables. On the other hand, we serve food at each table that all of the customers seating at it share, with no two tables serving the same dish. In that situation the sequence X_n will denote the dish served to the n th customer.

Regarding the Pólya sequence with $W_n = 1$, Proposition 11 of Pitman (1996) shows that one can proceed with the seating of customers, independently of the dish selection, and in a way choose the dishes post-factum by drawing them independently from ν for each table. In the general case, we have Theorem 3.2.1 of Chapter II, which implies that whenever the weights assigned to customers $n, n + 1, n + 2, \dots$ are independent of the dish served to the n th customer, given all past information, the food selection can happen again at the end, independent of the seating process. Moreover, the proportion of people in the population that will seat at each table is equal to the weighted average. If the conditions of Proposition 3.3.3 are further satisfied, then the number of people at each table grows at a rate n , so no table/dish becomes unpopular.

The model specification in Chapter IV assumes weights that are generated from distributions, which depend on the particular dish served at the table. In that case the partitioning of the observations into clusters is not independent of the food selection process. What results like Theorem 4.2.5 suggest though is that only a part of the tables will continue to accomodate customers, whereas the rest of them will become neglected with time. In this way the restaurant will have a sparse structure in the limit as only tables with "dominant" dishes will tend to be prepared.

The Indian buffet process (see Griffiths and Ghahramani, 2011), which is out of the scope of the thesis, characterizes the situation, in which each customer gets to sample more than one dish, with the possibility that people seating at the same table may get different dishes and that dishes may be shared across tables. In this analogy, tables become irrelevant and arriving customers select dishes according to their popularity or may choose to try untested ones. As a result, customers can be partitioned with respect to the dishes that they try and at the end may become associated with more than one cluster. Similarly to the work so far, Berti et al. (2015) implement a weighting scheme into the Indian buffet process and study its effects on the clustering behavior.

5.2 Clinical trials

Design of clinical trials has been an area of ongoing research for some time as researchers have tried to balance between the competing targets of gathering enough evidence for proper inference and of administering the superior treatment to the greatest number of patients. Within this field randomized, response-adaptive designs form an important class of protocols since they have some desirable characteristics (see, e.g., Rosenberger, 2002; Hu and Rosenberger, 2006; Zhang, 2015, for a discussion on the topic). For example, in contrast to standard designs with a one-time randomization of patients into groups, response-adaptive protocols enrol patients one at a time, so that when a new patient arrives, the clinician would select a treatment randomly according to predictive rules, which take into account all past information on the treatments' efficacy and reinforce the treatments having the better responses.

The dominant Pólya sequence (DPS) of Chapter IV encodes within its sampling mechanism such a response-driven design. In fact, the models of Section 4.4 have been developed with such an application in mind. In order to see this more clearly, suppose that each observation X_n indicates the treatment given to patient n and let W_n denote his or her response to it. The sampling scheme in Example 4.4.1 studies the simplest case of two competing treatments, say $\{0, 1\}$, with binary weights, indicating success or failure for the corresponding treatment. Each treatment has a different intrinsic probability of success, α_0, α_1 , so we may consider the treatment with the higher success probability as the better one. It follows from the results in Section 4.4 that the probability of assigning the better treatment to the next patient, given all past information on successes and failures, converges to 1.

In contrast, Example 4.4.2 considers the same case, but with continuous responses, which may, for example, represent a measurement on a favorable covariate such as the patient's survival time after treatment (as per Muliere et al., 2006). It follows again that the probability of assigning the better treatment goes to 1, where better is a function of the expected value of W_n given X_n . The case $m_1 := \mathbb{E}[W_n|X_n = 1] = \mathbb{E}[W_n|X_n = 0] =: m_0$ can be regarded as the null hypothesis for a statistical test on equivalence, and thus one can use the central limit theorems from Section 4.3 to derive p -values.

Example 4.4.3 is concerned with the situation of more than one "dominant" treatment from a set of k treatments, where dominance is again given in terms of the conditional expectations m_1, \dots, m_k . Under the assumption that $m_1 = \dots = m_{k_0}$, for some $1 \leq k_0 \leq k$, one has that both

$$\mathbb{P}(X_{n+1} \in \{1, \dots, k_0\} | X_1, \dots, X_n) \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}], \quad \text{and} \quad \frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\{1, \dots, k_0\}) \longrightarrow 1 \quad \text{a.s.}[\mathbb{P}].$$

It is hard to imagine a practical situation, in which two or more treatments have the same overall effect on people and for which the design of Example 4.4.3 would be applicable. In the clinical trials setting, however, one might appreciate the theoretical extension of the models in Section 4.4 to the infinite case. At first, suppose a different framework of interpretation – one of competing doses of a single treatment. As there is a continuum of doses, the fact that some of them are equally effective is highly probable. Theorem 4.2.2 shows that, under the assumption of positive probability of administering *near*-dominant doses, both the probability of assigning and the number of people assigned to a dose, which is at some positive distance from the dominant set, converge to 0. The stronger condition of non-zero probability of administering a dominant dose, although not needed for optimality, is less realistic in nature; yet, whenever satisfied, Theorem 4.3.1 provides tools to analyze the dynamics within the dominant set itself.

All in all, the protocol behind a DPS can be considered optimal in that the probability of administering the superior treatments/doses converges to 1 as the experiment proceeds. However, as argued in Aletti et al. (2018b), this need not be the preferred design in a clinical trial as it has suboptimal statistical properties. Instead, Aletti et al. (2018b) provide a modification of the model in Example 4.4.2 in order to have only a proportion of the patients assigned to the superior treatment, which ultimately resolves some of the inferential issues.

5.3 Randomly reinforced Poisson-Dirichlet sequence

Modeling after the exchangeable case, we get an immediate extension of RRPS by considering an \mathbb{X} -valued sequence of random variables $X = (X_n)_{n \geq 1}$ such that

$$\mathbb{P}(X_{n+1} \in \cdot | \mathcal{F}_n) = \sum_{i=1}^n \frac{W_i - \alpha / C_i(\Pi_n)}{\theta + \sum_{j=1}^n W_j} \delta_{X_k^*}(\cdot) + \frac{\theta + \alpha L_n}{\theta + \sum_{j=1}^n W_j} \nu(\cdot), \quad (\text{V.1})$$

for some sequence of non-negative random weights $W = (W_n)_{n \geq 1}$, $\nu \in \mathbb{M}_P(\mathbb{X})$, $\alpha \in \mathbb{R}_+$ and $\theta > -\alpha$, where $\mathcal{F}_n = \mathcal{F}_n^X \vee \mathcal{F}_n^W$, for $n \geq 0$, $\Pi_n = \{\Pi_{n,1}, \dots\}$ is the random partition of $\{1, \dots, n\}$ that is generated by (X_1, \dots, X_n) and has length L_n , and $C_i(\Pi_n) = k$ is the cardinality of the block in Π_n , which contains i . Denote by $(X_k^*)_{k=1}^{L_n}$ the distinct values in (X_1, \dots, X_n) . Then

$$\mathbb{P}(X_{n+1} \in \cdot | \mathcal{F}_n) = \sum_{k=1}^{L_n} \frac{(\sum_{i \in \Pi_{n,k}} W_i) - \alpha}{\theta + \sum_{j=1}^n W_j} \delta_{X_k^*}(\cdot) + \frac{\theta + \alpha L_n}{\theta + \sum_{j=1}^n W_j} \nu(\cdot).$$

We call the model with predictive distributions as in (V.1) a *randomly reinforced Poisson-Dirichlet sequence* (RRPD) with parameters α , θ , ν and W . If it holds $W_n = 1$, then we get the exchangeable two-parameter Poisson-Dirichlet process (see Pitman and Yor, 1997), given by

$$\mathbb{P}(X_{n+1} \in \cdot | X_1, \dots, X_n) = \sum_{k=1}^{L_n} \frac{\text{card}(\Pi_{n,k}) - \alpha}{\theta + \sum_{j=1}^n W_j} \delta_{X_k^*}(\cdot) + \frac{\theta + \alpha L_n}{\theta + \sum_{j=1}^n W_j} \nu(\cdot),$$

whereas for $\alpha = 0$, we are back to the RRPS case. The generalized Poisson-Dirichlet process (GPD) of Bassetti et al. (2010) is the RRPD with independent weights such that W_n is independent of (X_1, \dots, X_n) . In particular, the GPD satisfies assumption (A.1) of Chapter III and is shown by Bassetti et al. (2010, Theorem 2.1) to be conditionally identically distributed with respect to the filtration $\mathcal{F}^* = (\mathcal{F}_n^*)_{n \geq 0}$, given by $\mathcal{F}_n^* := \mathcal{F}_n^X \vee \mathcal{F}_\infty^W$, for $n \geq 0$. As a result,

$$P_n(\cdot) := \mathbb{P}(X_{n+1} \in \cdot | \mathcal{F}_n) \xrightarrow{w} \tilde{P}(\cdot) \quad \text{a.s.}[\mathbb{P}], \quad \text{and} \quad \hat{P}_n(\cdot) := \frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\cdot) \xrightarrow{w} \tilde{P}(\cdot) \quad \text{a.s.}[\mathbb{P}],$$

for some $\tilde{P} \in \mathbb{K}_P(\Omega, \mathbb{X})$. Suppose ν is diffuse. It follows from the discussion after Theorem 3.2.1 that

$$\tilde{P} = \sum_k p_k^* \delta_{X_k^*} + \left(1 - \sum_k p_k^*\right) \nu,$$

where

$$p_k^* = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\{X_k^*\}) \quad \text{a.s.}[\mathbb{P}],$$

and $(X_k^*)_{k \geq 1}$ are i.i.d. (ν) conditionally given $(p_k^*)_{k \geq 1}$. If $\sup_{n \in \mathbb{N}} \mathbb{E}[W_n^2] < \infty$ and $\lim_{n \rightarrow \infty} \mathbb{E}[W_n] = m$, then Bassetti et al. (2010, Proposition 4.1) prove that

$$\frac{\theta + \alpha L_n}{\theta + \sum_{j=1}^n W_j} \longrightarrow R \quad \text{a.s.}[\mathbb{P}], \quad \text{and} \quad \frac{L_n}{n} \longrightarrow R \quad \text{a.s.}[\mathbb{P}],$$

for some $R \in M_+(\mathcal{H})$. In particular, if $m > \alpha$, they show that $R = 0$ a.s.[\mathbb{P}], in which case \tilde{P} is proper. If it actually holds that $\sup_{n \in \mathbb{N}} \mathbb{E}[W_n^u] < \infty$, for some $u > 2$, and $q = \lim_{n \rightarrow \infty} \mathbb{E}[W_n^2]$, then Theorem 4.1 in Bassetti et al. (2010) implies

$$\sqrt{n}(\hat{P}_n(A) - P_n(A)) \xrightarrow{\text{stably}} \mathcal{N}(0, U(A)), \quad \text{and} \quad \sqrt{n}(P_n(A) - \tilde{P}(A)) \xrightarrow{\text{a.s. cond.}} \mathcal{N}(0, V(A)) \quad \text{w.r.t. } \mathcal{F},$$

for $A \in \mathcal{X}$, where

$$U(A) = \left(\frac{q}{m^2} - 1 \right) \tilde{P}(A)(1 - \tilde{P}(A)) + \frac{\alpha}{m} \left(\frac{\alpha}{m} - 2 \right) R \cdot \nu(A)(1 - \nu(A)),$$

and

$$V(A) = \frac{q}{m^2} \tilde{P}(A)(1 - \tilde{P}(A)) + \frac{\alpha^2}{m^2} R \cdot \nu(A)(1 - \nu(A)).$$

In the case $m > \alpha$, then $V(A) = \frac{q}{m^2} \tilde{P}(A)(1 - \tilde{P}(A))$ and $U(A) = V(A) - \tilde{P}(A)(1 - \tilde{P}(A))$, which is very reminiscent of the results in Theorem 3.3.1 of Chapter III. In fact, we can relax the conditions in Theorem 3.3.1 to match those for a GDP.

We leave for future research the extension of the results in Chapter IV to RRPDs having weights of the form

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

where $(U_n)_{n \geq 1}$ is a sequence of independent random variables such that U_n is independent of (X_1, \dots, X_n) .

5.4 Multi-experiment data

The current section should serve as the starting point for future research on groups of interacting RRPSs. As a matter of fact, such processes are already known in the literature through the works of Paganoni and Secchi (2004) and Fortini et al. (2018). The following definition, which is derived from Example 3.8 of Fortini et al. (2018), formalizes the notion of a family of interacting RRPSs.

Definition 5.4.1. *An array of \mathbb{X} -valued random variables $X = [X_{n,i}]_{n \geq 1, i \in I}$ represents a family of interacting randomly reinforced Pólya sequences (IRRRPS) if there exist $(\nu_i)_{i \in I} \subseteq \mathbb{M}_P(\mathbb{X})$ such that $X_{1,i} \sim \nu_i$, for $i \in I$, constants $(\theta_i)_{i \in I} \subseteq \mathbb{R}_+$, and an array of non-negative random variables $W = [W_{n,i}]_{n \geq 1, i \in I}$ such that a version of the conditional distribution of $X_{n+1,i}$ given $\mathcal{F}_n := \sigma(\mathbf{X}_1^I, \mathbf{W}_1^I, \dots, \mathbf{X}_n^I, \mathbf{W}_n^I)$ is the transition probability kernel*

$$\mathbb{P}(X_{n+1,i} \in \cdot | \mathcal{F}_n) = \sum_{k=1}^n \frac{W_{k,i}}{\theta_i + \sum_{j=1}^n W_{j,i}} \delta_{X_{k,i}}(\cdot) + \frac{\theta_i}{\theta_i + \sum_{j=1}^n W_{j,i}} \nu_i(\cdot), \quad \text{for } n \geq 1 \text{ and } i \in I,$$

where $\mathbf{X}_n^I = (X_{n,i}, i \in I)$ and $\mathbf{W}_n^I = (W_{n,i}, i \in I)$ represent the n th rows of X and W , respectively.

If $I = \{1\}$, then we are back to the univariate case. Otherwise, from the towering property of conditional expectations and from conditional determinism it holds

$$\mathbb{P}(X_{n+1,i} \in \cdot | X_{1,i}, W_{1,i}, \dots, X_{n,i}, W_{n,i}) = \mathbb{P}(X_{n+1,i} \in \cdot | \mathcal{F}_n),$$

so that each sequence $(X_{n,i})_{n \geq 1}$, for $i \in I$, constitutes a RRPS in itself. As the filtration system $(\mathcal{F}_n)_{n \geq 1}$ contains the *combined* history of the individual RRPSs, one way to read the above equality is to assume, for any given $i \in I$, that the weight $W_{n,i}$ depends on the joint set of observations $\mathbf{X}_1^I, \dots, \mathbf{X}_n^I$. In fact, the *interacting reinforced-urn system* of Paganoni and Secchi (2004) is a countable collection of two-color randomly reinforced urns (refer to Example 4.4.2 from Chapter IV) such that the reinforcement of each urn is a function of the observations from the other urns plus an independent component. In particular, their Example 3.1 considers a system of two urns, $I = \{1, 2\}$, whose weights are given by $W_{n,1} = X_{n,2}$ and $W_{n,2} = X_{n,1}$. Such models are particularly important in that they represent intuitive constructions of dependent stochastic processes with reinforcement, and yet are such that if the individual sequences were exchangeable, the joint process fails to be partially exchangeable in the sense of de Finetti. The next subsection elaborates more on that.

5.4.1 Partially c.i.d. interacting RRPS

In the context of multiple experiments, partial exchangeability in the sense of de Finetti (see Fortini et al., 2018, for historical references) is the standard dependence structure for inference in Bayesian statistics. Roughly speaking, an array of random variables $[X_{n,i}]_{n \geq 1, i \in I}$ is said to be *partially exchangeable* if the joint probability law is invariant under permutations within, but not across the different groups. Historically, partially exchangeable laws have been specified through a system of dependent random probability measures (see, e.g. Bassetti et al., 2018; Camerlenghi et al., 2019, and references therein). In contrast, predictive construction of partially exchangeable processes remains an open question as the introduction of interactions in the system leads to complicated predictive rules. In fact, what can be considered natural predictive constructions could easily break the symmetry required by partial exchangeability. For example, let $I = \{1, 2\}$ and suppose $X_{n+1,1}$ and $X_{n+1,2}$ are conditionally independent given the past observations $\mathbf{X}_1^I, \dots, \mathbf{X}_n^I$. This scheme could represent a game of sequential decisions between two players, in which the players choose their next moves independently from one another, but have access to the whole decision history of their opponent. Under this assumption, the joint process is not partially exchangeable (except trivially, when the sequences are unconditionally independent) as in that case $X_{n+1,1}$ and $X_{n+1,2}$ would be independent, but conditionally on the tail information.

Addressing these issues requires the relaxation of some of the structure imposed by partial exchangeability, yet much as in the univariate case, one would like to simultaneously preserve most of its important properties, at least in the limit. To that end, Fortini et al. (2018) have proposed an extension of the conditional identity in distribution property, which they call partial conditional identity in distribution. An array of random variables $[X_{n,i}]_{n \geq 1, i \in I}$ is said to be *partially conditionally identically distributed* (p.c.i.d.) if the future observations of each individual sequence are conditionally identically distributed given all past observations *and* the concomitant values of the other variables; in other words, it holds, for each $k, n \geq 1$, $i \in I$ and any $f \in M_b(\mathcal{X})$, that

$$\mathbb{E}[f(X_{n+k,i}) | \mathcal{F}_n^i] = \mathbb{E}[f(X_{n+1,i}) | \mathcal{F}_n^i],$$

where $\mathcal{F}_n^i := \sigma(\mathbf{X}_1^I, \mathbf{W}_1^I, \dots, \mathbf{X}_n^I, \mathbf{W}_n^I) \vee \sigma(X_{n+1,j} : j \neq i)$. Fortini et al. (2018) prove among all that p.c.i.d. arrays are asymptotically partially exchangeable. Now, contrary to the partial exchangeability case, p.c.i.d.

probability laws can be readily constructed by interweaving interactions into a set of predictive rules, as long as the individual sequences are updated independently given the past.

Regarding IRRPS, the latter would amount to the following assumption,

Assumption B.1. $(X_{n+1,i})_{i \in I}$ are mutually independent given \mathcal{F}_n , for each $n \geq 0$. (B.1)

As the individual sequences of a p.c.i.d. process are c.i.d. themselves, we may also require as in Chapter III that

Assumption B.2. $X_{n+1,i}$ is independent of $W_{n+1,i}$ given \mathcal{F}_n , for each $n \geq 0$ and $i \in I$. (B.2)

Under those two conditions, Fortini et al. (2018, Example 3.8) prove that a family of IRRPS is p.c.i.d. An interesting example of a p.c.i.d. IRRPS is due to (Paganoni and Secchi, 2004, Example 3.1), which we outline next. In fact, by applying Theorem 3.2.3 from Chapter III we are able to prove a conjecture of theirs regarding the univariate distributions of the random limits of the predictive distributions.

Example 5.4.2 (Coupled Pólya urns). Suppose $\mathbb{X} = \{0, 1\}$ and $I = \{1, 2\}$. Let $(b_i, w_i)_{i=1,2} \subseteq \mathbb{R}_+^2$ be such that $b_i + w_i > 0$. Denote by $\theta_i := b_i + w_i$ and $\nu_i := b_i \delta_1 + w_i \delta_0$, for $i = 1, 2$. Assume that the predictive distributions of the two sequences of IRRPSs $[X_{n,i}]_{n \geq 1, i=1,2}$ are given by

$$\mathbb{P}(X_{n+1,i} = 1 \mid \mathcal{F}_n) = \frac{b_i + \sum_{k=1}^n X_{k,i} X_{k,j}}{b_i + w_i + \sum_{k=1}^n X_{k,j}}, \quad \text{for } i = 1, 2,$$

where $j \in I \setminus \{i\}$ and $\mathcal{F}_n := \sigma(X_{1,1}, X_{1,2}, \dots, X_{n,1}, X_{n,2})$. The model is comprised of two interacting two-color randomly reinforced urns such that each urn is reinforced only when the ball extracted from the other urn has color 1. Paganoni and Secchi (2004) demonstrate that the predictive distributions of both sequences converge to a random vector (Z_1, Z_2) such that $\text{Cov}(Z_1, Z_2) > 0$. They further conjecture that Z_1 and Z_2 are individually distributed according to $\text{Beta}(b_1, w_1)$ and $\text{Beta}(b_2, w_2)$ distributions, respectively. As each of the individual sequences forms a RRPS with binary weights such that

$$\sum_{n=1}^{\infty} \mathbb{E}[X_{n,i} \mid \mathcal{F}_{n-1}] \geq \sum_{n=1}^{\infty} \frac{b_i}{b_i + w_i + n} = +\infty \quad \text{a.s.}[\mathbb{P}], \text{ for } i = 1, 2,$$

then $\mathbb{P}(X_{n,i} = 1 \text{ i.o.}) = 1$, for each $i = 1, 2$, so it follows from Theorem 3.2.3 of Chapter III that the proportion of balls of color 1 in each urn converges to a $\text{Beta}(b_i, w_i)$ random variable.

5.4.2 Related models

A different way to elicit dependence between a group of RRPSs is through shared atoms. In order to do so, one has to allow the base measure ν to be a random discrete probability measure and/or to depend on the time n .

Let G_0 be a random discrete probability measure on \mathbb{X} . Consider a pair of IRRPS $[X_{n,i}]_{n \geq 1, i=1,2}$ that have independent weights and are characterized by the predictive rules

$$\mathbb{P}(X_{n+1,i} \in \cdot | \mathcal{F}_n) = \sum_{k=1}^n \frac{W_{k,i}}{\theta_i + \sum_{j=1}^n W_{j,i}} \delta_{X_{k,i}}(\cdot) + \frac{\theta_i}{\theta_i + \sum_{j=1}^n W_{j,i}} \left(\frac{G_0(\cdot) + \nu(\cdot)}{G_0(\mathbb{X}) + \nu(\mathbb{X})} \right), \quad \text{for } i = 1, 2,$$

where $\mathcal{F}_n := \sigma(\mathbf{X}_1^I, \mathbf{W}_1^I, \dots, \mathbf{X}_n^I, \mathbf{W}_n^I) \vee \sigma(G_0)$, for $n \geq 0$. This process can be seen as a weighted version of the bivariate Dirichlet process of Walker and Muliere (2003), which can be recovered in the case $\theta_1 = \theta_2 = \theta$, $W_{n,i} = 1$, for $n \geq 1$ and $i = 1, 2$, and $G_0 = \sum_{j=1}^k \delta_{Z_j}$, where $(Z_j)_{j=1}^k$ is a sample from a $\text{DP}(\theta, \nu)$ process. It follows that $(X_{n,1})_{n \geq 1}$ and $(X_{n,2})_{n \geq 1}$ are conditionally independent RRPSs given G_0 .

Take instead the predictive system

$$\mathbb{P}(X_{n+1,i} \in \cdot | \mathcal{F}_n) = \sum_{k=1}^n \frac{W_{k,i}}{\theta_i + \sum_{j=1}^n W_{j,i}} \delta_{X_{k,i}}(\cdot) + \frac{\theta_i}{\theta_i + \sum_{j=1}^n W_{j,i}} G_n(\cdot), \quad \text{for } i = 1, 2,$$

where, for some RRPS $(Z_n)_{n \geq 1}$ with parameters θ, ν and $(Y_n)_{n \geq 1}$, it holds

$$G_n(\cdot) := \mathbb{P}(Z_{n+1} \in \cdot | \mathcal{F}_n) = \sum_{i=1}^n \frac{Y_i}{\theta_0 + \sum_{j=1}^n Y_j} \delta_{Z_i}(\cdot) + \frac{\theta_0}{\theta_0 + \sum_{j=1}^n Y_j} \nu(\cdot),$$

with $\mathcal{F}_n := \mathcal{F}_n^X \vee \mathcal{F}_n^Z$. It is straightforward to show that $[X_{n,i}]_{n \geq 1, i=1,2}$ is p.c.i.d. under (B.2) and the following extended version of (B.1),

$$Z_{n+1}, X_{n+1,1} \text{ and } X_{n+1,2} \text{ are mutually independent given } \mathcal{F}_n, \text{ for each } n \geq 0. \quad (\text{V.1})$$

Moreover, $(Z_n)_{n \geq 1}$ is c.i.d. under its own version of (B.2), in which case $G_n \xrightarrow{w} G_0$ a.s. [P], for some random probability measure G_0 . Provided further $W_{n,i} = 1$, for $n \geq 1$ and $i = 1, 2$, the predictive distributions $\mathbb{P}(X_{n+1,i} \in \cdot | \mathcal{F}_n)$ converge weakly on a set of probability one to a random probability measure $G_i(\cdot)$, which is a $\text{DP}(\theta_i, G_0)$ process. The hierarchical Dirichlet process of Teh et al. (2006) is then the special case with $Y_n = 1$, for $n \geq 1$, so that $G_0 \sim \text{DP}(\theta_0, \nu)$.

5.5 Conclusion

In the final section we provide, in no particular order, comments, as well as suggestions on how to extend the research on RRPSs or any of topics that were touched upon in the thesis. We have said in Chapter I that one of the main reasons for the development of the stochastic processes, which we consider, is their potential role in approximate Bayesian analysis. Before going into that direction though, we need a merging of opinions type of result as in Blackwell and Dubins (1962) that will ensure the consistency of the conclusions reached by the approximation. On the other hand, our basic model is defined with respect to a sequence of weights $(W_n)_{n \geq 1}$, which may be non-observable depending on the context. We have suggested in Section 3.4, Chapter III a way to make draws from the posterior distribution of the weights in the c.i.d. case, yet inference on even the dominance function w of a DPS is more critical because of the particular meaning

attached to the weighting process. In fact, our current research is aimed in this direction. We recall also the comment, made in Section 2.3, Chapter II, regarding the lack of studies on Pólya urn models with absolutely continuous reinforcements, for which we believe the measure-valued Pólya urn process of Janson (2018) to be an excellent starting point. Immediate extensions of the theory in Chapter III and Chapter IV could involve work on uniform limit theorems or further inquiry into the random partition of the RRPS, in particular as regards to the problem of species sampling (see, e.g. Favaro et al., 2012a,b). Lastly, we have outlined in this chapter some possible generalizations of the RRPS in the uni- and multivariate setting.

Appendix A

Miscellaneous results

Results in measure theory and real analysis

Lemma A.1 (Two-Valued Measure is Dirac Measure). *Let $\mu \in \mathbb{M}_P(\mathbb{R}_+)$ be such that $\mu \in \{0, 1\}$. Then there exists $x \in \mathbb{R}$ such that $\mu(\{x\}) = 1$ or, equivalently, $\mu = \delta_x$.*

Proof. Define $I_n := [-n, n]$, for $n \geq 1$. Suppose, by contradiction, $\mu(I_n) = 0$ for all $n \geq 1$. Then $1 = \mu(\mathbb{R}) = \mu(\bigcup_{n=1}^{\infty} I_n) = \lim_{n \rightarrow \infty} \mu(I_n) = 0$, absurd, hence there exists $n \in \mathbb{N}$ such that $\mu(I_n) = 1$. Denote $J_0 = I_n$. Set $J_0^- = [-n, 0]$ and $J_0^+ = [0, n]$. Then either $\mu(J_0^-) = 1$ or $\mu(J_0^+) = 1$. Denote that set by J_1 . Iterating the procedure yields a sequence $(J_k)_{k \geq 0}$ such that $\mu(J_k) = 1$, $\text{diam}(J_k) \leq n/2^{k-1} \rightarrow 0$ and $J_k \downarrow \bigcap_{k=0}^{\infty} J_k =: J$. As \mathbb{R} is complete, one has from Cantor's intersection theorem that $J = \{x\}$, for some $x \in \mathbb{R}$. As a consequence,

$$\mu(\{x\}) = \mu(J) = \lim_{k \rightarrow \infty} \mu(J_k) = 1.$$

□

Lemma A.2 (Approximation of Bounded Functions). *Let (X, \mathcal{F}) be a measurable space, and $f \in M_b(\mathcal{F})$. Then, for each $\epsilon > 0$, there exist $f_1, f_2 \in M_0(\mathcal{F})$ such that $f_1 \leq f \leq f_2$ and*

$$\sup_{x \in X} (f_2(x) - f_1(x)) < \epsilon.$$

Proof. Suppose $|f| \leq M$, for some $M \in \mathbb{R}_+$. From standard approximation results for measurable functions, there exist $(g_n)_{n \geq 1}, (\hat{h}_n)_{n \geq 1} \subseteq M_0(\mathcal{F})$ such that $g_n(x) \leq g_{n+1}(x)$ and $\hat{h}_n(x) \leq \hat{h}_{n+1}(x)$, for $x \in X$ and $n \geq 1$, and it holds

$$\lim_{n \rightarrow \infty} \sup_{x \in X} |f(x) - g_n(x)| = 0, \quad \text{and} \quad \lim_{n \rightarrow \infty} \sup_{x \in X} |(M - f(x)) - \hat{h}_n(x)| = 0.$$

Define $h_n := M - \hat{h}_n$, for $n \geq 1$. Then $h_n(x) \geq h_{n+1}(x)$, for $x \in X$ and $n \geq 1$, and $\lim_{n \rightarrow \infty} \sup_{x \in X} |f(x) - h_n(x)| = 0$. Let $\epsilon > 0$. There exist $\bar{n}_1, \bar{n}_2 \in \mathbb{N}$ such that $\sup_{x \in X} |f(x) - g_n(x)| < \epsilon/2$, for $n \geq \bar{n}_1$, and

$\sup_{x \in X} |h_n(x) - f(x)| < \epsilon/2$, for $n \geq \bar{n}_2$. Therefore,

$$\sup_{x \in X} |h_n(x) - g_n(x)| < \epsilon, \quad \text{for } n \geq \max\{\bar{n}_1, \bar{n}_2\}.$$

□

Lemma A.3 (Weighted Average). *Let $(a_n)_{n \geq 1}, (b_n)_{n \geq 1}, (h_n)_{n \geq 1} \subseteq \mathbb{R}$ be such that $a_n \geq 0$ and $h_n \rightarrow \infty$. If $\frac{1}{h_n} \sum_{i=1}^n a_i \rightarrow 1$ and $b_n \rightarrow b$, for some $b \in \mathbb{R}$, then*

$$\frac{1}{h_n} \sum_{i=1}^n a_i b_i \rightarrow b.$$

Proof. It follows that,

$$\frac{1}{h_n} \sum_{i=1}^n a_i b_i = b \frac{1}{h_n} \sum_{i=1}^n a_i + \frac{1}{h_n} \sum_{i=1}^n a_i (b_i - b).$$

Fix $\epsilon > 0$. Then there exists $\bar{n}_1 \in \mathbb{N}$ such that $|b_n - b| < \epsilon/2$, for $n \geq \bar{n}_1$; thus,

$$\left| \frac{1}{h_n} \sum_{i=1}^n a_i (b_i - b) \right| \leq \frac{1}{h_n} \sum_{i=1}^n a_i |b_i - b| \leq \frac{1}{h_n} \sum_{i=1}^{\bar{n}_1} a_i |b_i - b| + \frac{\epsilon}{2} \frac{1}{h_n} \sum_{i=1}^n a_i.$$

But $\frac{1}{h_n} \sum_{i=1}^{\bar{n}_1} a_i |b_i - b| \rightarrow 0$ as $h_n \rightarrow \infty$, so there is $\bar{n}_2 \in \mathbb{N}$ such that $\frac{1}{h_n} \sum_{i=1}^{\bar{n}_1} a_i |b_i - b| < \epsilon/4$, for $n \geq \bar{n}_2$. Finally, take $\bar{n}_3 \in \mathbb{N}$ such that $\frac{1}{h_n} \sum_{i=1}^n a_i < 3/2$, for $n \geq \bar{n}_3$. As a result, $\left| \frac{1}{h_n} \sum_{i=1}^n a_i (b_i - b) \right| < \epsilon$, for $n \geq \max\{\bar{n}_1, \bar{n}_2, \bar{n}_3\}$, so $\frac{1}{h_n} \sum_{i=1}^n a_i (b_i - b) \rightarrow 0$ and the conclusion follows. □

Lemma A.4 (Weighted Average 2). *Let $(a_n)_{n \geq 1}, (b_n)_{n \geq 1} \subseteq \mathbb{R}$ be such that $a_n \geq 0$. If $\frac{1}{n} \sum_{i=1}^n a_i \rightarrow a$ and $b_n \rightarrow b$, for some $a, b \in \mathbb{R}$, then $\frac{1}{n} \sum_{i=1}^n a_i b_i \rightarrow ab$.*

Proof. The proof is similar to the one for Lemma A.3. □

Lemma A.5 (Abel's Test). *Let $(a_n)_{n \geq 1} \subseteq \mathbb{R}$ be such that $\sum_{n=1}^{\infty} a_n < \infty$, and $(b_n)_{n \geq 1} \subseteq \mathbb{R}$ be a bounded monotone sequence. Then $\sum_{n=1}^{\infty} a_n b_n < \infty$. In particular, $\sum_{n=1}^{\infty} n^{-1} a_n < \infty$, and thus $n \cdot \sum_{m \geq n} m^{-1} a_m \rightarrow 0$.*

Proof. Set $a_0 = 0$. Define $S_n := \sum_{i=0}^n a_i$, for $n \geq 0$. Then $a_n = S_n - S_{n-1}$ and

$$\sum_{i=1}^n a_i b_i = \sum_{i=1}^n (S_i - S_{i-1}) b_i = \sum_{i=1}^n S_i (b_i - b_{i+1}) + S_n b_n.$$

As $(b_n)_{n \geq 1}$ is bounded and monotone, then $b_n \rightarrow b$, for some $b \in \mathbb{R}$, and thus $S_n b_n$ converges. Take $M \in \mathbb{R}$ such that $|\sum_{n=1}^{\infty} a_n| < M$. From the monotonicity of $(b_n)_{n \geq 1}$, one has

$$\sum_{i=1}^n |S_i (b_i - b_{i+1})| \leq M \sum_{i=1}^n |b_i - b_{i+1}| \leq M |b_1 - b_{n+1}| \rightarrow M |b_1 - b|;$$

hence $\sum_{i=1}^n S_i (b_i - b_{i+1})$ converges absolutely. □

Results on conditional expectations

Lemma A.6 (Joint Conditional Expectation). *Let $(\Omega, \mathcal{H}, \mathbb{P})$ be a probability space, $\mathcal{F} \subseteq \mathcal{H}$ be a sub- σ -algebra, $(\mathbb{X}, \mathcal{X})$, $(\mathbb{Y}, \mathcal{Y})$ and $(\mathbb{Z}, \mathcal{Z})$ be measurable spaces, X, Y and Z be random variables on $(\Omega, \mathcal{H}, \mathbb{P})$ with values in \mathbb{X}, \mathbb{Y} and \mathbb{Z} , respectively, and $\kappa \in \mathbb{K}_P(\Omega, \mathbb{X})$ be a version of the conditional distribution of X given \mathcal{F} . If Y is \mathcal{F} -measurable and Z is independent of $\sigma(X) \vee \mathcal{F}$, then the mapping defined by*

$$\kappa f(Y(\omega)) := \int_{\mathbb{X}} \int_{\Omega} f(x, Y(\omega), Z(\omega')) \mathbb{P}(d\omega') \kappa(\omega, dx), \quad \text{for } \omega \in \Omega,$$

is a version of $\mathbb{E}[f(X, Y, Z)|\mathcal{F}]$ for each $f \in M_+(\mathcal{X} \otimes \mathcal{Y} \otimes \mathcal{Z})$.

Proof. Let $f \in M_+(\mathcal{X} \otimes \mathcal{Y} \otimes \mathcal{Z})$. By Theorem 2.19 in Cinlar (2011, Chapter IV), one has that

$$h(X(\omega), Y(\omega)) := \int_{\Omega} f(X(\omega), Y(\omega), Z(\omega')) \mathbb{P}(d\omega'), \quad \text{for } \omega \in \Omega,$$

is a version of $\mathbb{E}[f(X, Y, Z)|\mathcal{F} \vee \sigma(X)]$, for some $h \in M_+(\mathcal{X} \otimes \mathcal{Y})$. From repeated conditioning and the same theorem, the map $\omega \mapsto \int_{\mathbb{X}} h(x, Y(\omega)) \kappa(\omega, dx)$ from Ω to \mathbb{R}_+ is a version of $\mathbb{E}[f(X, Y, Z)|\mathcal{F}]$. \square

Lemma A.7 (Empirical Convergence from Predictive Convergence [Berti et al., 2011, Lemma 3]). *Let $(\Omega, \mathcal{H}, \mathbb{P})$ be a probability space, $\mathcal{F} = (\mathcal{F}_n)_{n \geq 0}$ be a filtration on (Ω, \mathcal{H}) , and $(X_n)_{n \geq 1}$ be an \mathcal{F} -adapted sequence of real-valued random variables on $(\Omega, \mathcal{H}, \mathbb{P})$. If*

$$\sum_{n=1}^{\infty} \frac{\mathbb{E}[X_n^2]}{n^2} < \infty, \quad \text{and} \quad \mathbb{E}[X_{n+1}|\mathcal{F}_n] \longrightarrow X \quad \text{a.s.}[\mathbb{P}],$$

for some $X \in M(\mathcal{H})$, then

$$n \cdot \sum_{k \geq n} \frac{X_k}{k^2} \longrightarrow X \quad \text{a.s.}[\mathbb{P}], \quad \text{and} \quad \frac{1}{n} \sum_{k=1}^n X_k \longrightarrow X \quad \text{a.s.}[\mathbb{P}].$$

Proof. Define $M_n := \sum_{k=1}^n \frac{1}{k} (X_k - \mathbb{E}[X_k|\mathcal{F}_{k-1}])$, for $n \geq 1$. It follows that,

$$\mathbb{E}[M_{n+1}|\mathcal{F}_n] = \mathbb{E}\left[\sum_{k=1}^{n+1} \frac{X_k - \mathbb{E}[X_k|\mathcal{F}_{k-1}]}{k} \middle| \mathcal{F}_n\right] = M_n + \frac{1}{n+1} \mathbb{E}[X_{n+1} - \mathbb{E}[X_{n+1}|\mathcal{F}_n]|\mathcal{F}_n] = M_n,$$

and $\sup_{n \in \mathbb{N}} \mathbb{E}[M_n^2] \leq 4 \sum_{n=1}^{\infty} \mathbb{E}[X_n^2]/n^2 < \infty$; therefore, $(M_n)_{n \geq 1}$ is a uniformly integrable \mathcal{F} -martingale and converges \mathbb{P} -almost surely. From Lemma A.5,

$$n \cdot \sum_{k \geq n} \frac{X_k - \mathbb{E}[X_k|\mathcal{F}_{k-1}]}{k^2} \longrightarrow 0 \quad \text{a.s.}[\mathbb{P}].$$

But $n \sum_{k \geq n} 1/k^2 \rightarrow 1$, therefore,

$$n \cdot \sum_{k \geq n} \frac{X_k}{k^2} = n \cdot \sum_{k \geq n} \frac{X_k - \mathbb{E}[X_k|\mathcal{F}_{k-1}]}{k^2} + n \cdot \sum_{k \geq n} \frac{\mathbb{E}[X_k|\mathcal{F}_{k-1}]}{k^2} \longrightarrow X \quad \text{a.s.}[\mathbb{P}].$$

Similarly, by Kronecker's lemma,

$$\frac{1}{n} \sum_{k=1}^n (X_k - \mathbb{E}[X_k | \mathcal{F}_{k-1}]) \longrightarrow 0 \quad \text{a.s.}[\mathbb{P}].$$

As $\mathbb{E}[X_{n+1} | \mathcal{F}_n] \xrightarrow{\text{a.s.}} X$, then $\frac{1}{n} \sum_{k=1}^n \mathbb{E}[X_k | \mathcal{F}_{k-1}] \xrightarrow{\text{a.s.}} X$, so that

$$\frac{1}{n} \sum_{k=1}^n X_k = \frac{1}{n} \sum_{k=1}^n \mathbb{E}[X_k | \mathcal{F}_{k-1}] + \frac{1}{n} \sum_{k=1}^n (X_k - \mathbb{E}[X_k | \mathcal{F}_{k-1}]) \longrightarrow X \quad \text{a.s.}[\mathbb{P}].$$

□

Lemma A.8 (Conditional Criteria for A.S. Convergence [Pemantle and Volkov, 1999, Lemma 3.2]). *Let $(\Omega, \mathcal{H}, \mathbb{P})$ be a probability space, $\mathcal{F} = (\mathcal{F}_n)_{n \geq 0}$ be a filtration on (Ω, \mathcal{H}) , and $(X_n)_{n \geq 1}$ be an \mathcal{F} -adapted sequence of real-valued random variables on $(\Omega, \mathcal{H}, \mathbb{P})$. If*

$$\sum_{n=1}^{\infty} \mathbb{E}[X_{n+1} - X_n | \mathcal{F}_n] < \infty \quad \text{a.s.}[\mathbb{P}], \quad \text{and} \quad \sum_{n=1}^{\infty} \mathbb{E}[(X_{n+1} - X_n)^2 | \mathcal{F}_n] < \infty \quad \text{a.s.}[\mathbb{P}],$$

then there exists $X \in \mathcal{L}(\mathcal{H})$ such that

$$X_n \longrightarrow X \quad \text{a.s.}[\mathbb{P}].$$

Proof. Let $M \in \mathbb{R}$. Define $T_M := \inf\{n \in \mathbb{N} : \sum_{i=1}^n \mathbb{E}[(X_{i+1} - X_i)^2 | \mathcal{F}_i] > M\}$ and $X_n^{(M)} := X_{n \wedge T_M} - \sum_{i=1}^{n \wedge T_M} \mathbb{E}[X_{i+1} - X_i | \mathcal{F}_i]$, for $n \geq 1$. Then $(X_n^{(M)})_{n \geq 1}$ is an \mathcal{F} -martingale such that

$$\mathbb{E}[(X_{n+1}^{(M)} - X_n^{(M)})^2 | \mathcal{F}_n] \leq \text{Var}(X_{n+1} - X_n | \mathcal{F}_n) \cdot \mathbb{1}_{\{T_M > n\}} \leq \mathbb{E}[(X_{n+1} - X_n)^2 | \mathcal{F}_n] \cdot \mathbb{1}_{\{T_M > n\}};$$

thus, $X_n^{(M)}$ converges a.s.[\mathbb{P}] and in L^2 to a finite limit, say C_M . As $T_M = \infty$ on $\{\sum_{n=1}^{\infty} \mathbb{E}[(X_{n+1} - X_n)^2 | \mathcal{F}_n] < \infty\}$ for a sufficiently large M , then

$$X_n \longrightarrow C_M + \sum_{n=1}^{\infty} \mathbb{E}[X_{n+1} - X_n | \mathcal{F}_n] \quad \text{a.s.}[\mathbb{P}].$$

□

Results on c.i.d. processes

Lemma A.9 (Convergence of Predictive/Empirical Distribution on Random Sets). *Let $(\Omega, \mathcal{H}, \mathbb{P})$ be a probability space, $\mathcal{F} = (\mathcal{F}_n)_{n \geq 0}$ be a filtration on (Ω, \mathcal{H}) , \mathbb{X} and \mathbb{Y} be c.s.m.s. equipped with their associated Borel σ -algebras \mathcal{X} and \mathcal{Y} , $(X_n)_{n \geq 1}$ be an \mathcal{F} -c.i.d. sequence of \mathbb{X} -valued random variables on $(\Omega, \mathcal{H}, \mathbb{P})$ with directing measure $\tilde{P} \in \mathbb{K}_P(\Omega, \mathbb{X})$, T be an a.s.-finite \mathcal{F} -stopping time, and Y be an \mathbb{Y} -valued \mathcal{F}_T -measurable random variable. Then,*

$$\mathbb{E}[f(Y, X_{n+1}) | \mathcal{F}_n] \longrightarrow \int_{\mathbb{X}} f(Y, x) \tilde{P}(dx) \quad \text{a.s.}[\mathbb{P}] \text{ and in } L^1,$$

and

$$\frac{1}{n} \sum_{i=1}^n f(Y, X_i) \longrightarrow \int_{\mathbb{X}} f(Y, x) \tilde{P}(dx) \quad \text{a.s.}[\mathbb{P}],$$

for each $f \in M_b(\mathcal{Y} \otimes \mathcal{X})$. In particular,

$$\mathbb{P}(X_{n+1} = X_T | \mathcal{F}_n) \longrightarrow \tilde{P}(\{X_T\}) \quad \text{a.s.}[\mathbb{P}] \text{ and in } L^1,$$

and

$$\frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\{X_T\}) \longrightarrow \tilde{P}(\{X_T\}) \quad \text{a.s.}[\mathbb{P}].$$

Proof. Let $f \in M_b(\mathcal{X})$. As $(X_n)_{n \geq 0}$ is \mathcal{F} -c.i.d., then $(\mathbb{E}[f(X_{n+1}) | \mathcal{F}_n])_{n \geq 1}$ is an \mathcal{F} -martingale that converges \mathbb{P} -a.s. to $\tilde{P}_f := \int_{\mathbb{X}} f(x) \tilde{P}(dx)$. In fact, $\mathbb{E}[\tilde{P}_f | \mathcal{F}_n] = \mathbb{E}[f(X_{n+1}) | \mathcal{F}_n]$ a.s. $[\mathbb{P}]$, for $n \geq 1$. It follows that

$$\begin{aligned} \mathbb{E}[\tilde{P}_f | \mathcal{F}_{T+n}] &= \mathbb{E}[\tilde{P}_f | \mathcal{F}_{T+n}] \mathbb{1}_{\{T < \infty\}} = \sum_{m=1}^{\infty} \mathbb{E}[\tilde{P}_f | \mathcal{F}_{T+n}] \mathbb{1}_{\{T=m\}} = \\ &= \sum_{m=1}^{\infty} \mathbb{E}[\tilde{P}_f | \mathcal{F}_{m+n}] \mathbb{1}_{\{T=m\}} = \sum_{m=1}^{\infty} \mathbb{E}[f(X_{m+n+1}) | \mathcal{F}_{m+n}] \mathbb{1}_{\{T=m\}} = \mathbb{E}[f(X_{T+n+1}) | \mathcal{F}_{T+n}], \end{aligned}$$

where we have used that¹ $\mathbb{E}[Z | \mathcal{F}_{T+n}] \mathbb{1}_{\{T=m\}} = \mathbb{E}[Z | \mathcal{F}_{m+n}] \mathbb{1}_{\{T=m\}}$ a.s. $[\mathbb{P}]$, for each $m \geq 1$ and any $Z \in \mathcal{L}(\mathcal{H})$. Let $V \in M_b(\mathcal{F}_T)$. As $\mathcal{F}_T \subseteq \mathcal{F}_{T+n}$, one has \mathbb{P} -a.s. that

$$\begin{aligned} \mathbb{E}\left[\int_{\mathbb{X}} V f(x) \tilde{P}(dx) | \mathcal{F}_{T+n}\right] &= V \cdot \mathbb{E}[\tilde{P}_f | \mathcal{F}_{T+n}] = \\ &= V \cdot \mathbb{E}[f(X_{T+n+1}) | \mathcal{F}_{T+n}] = \mathbb{E}[V f(X_{T+n+1}) | \mathcal{F}_{T+n}]. \end{aligned}$$

In particular, the result holds with $V = \mathbb{1}_{\{Y \in A\}}$ and $f = \mathbb{1}_B$, for any $A \in \mathcal{Y}$ and $B \in \mathcal{X}$. Define

$$\mathcal{M} := \left\{ f \in M_b(\mathcal{Y} \otimes \mathcal{X}) : \mathbb{E}\left[\int_{\mathbb{X}} f(Y, x) \tilde{P}(dx) | \mathcal{F}_{T+n}\right] = \mathbb{E}[f(Y, X_{T+n+1}) | \mathcal{F}_{T+n}] \quad \text{a.s.}[\mathbb{P}], \text{ for } n \geq 1 \right\}.$$

Let $a, b \in \mathbb{R}$ and $f, g \in \mathcal{M}$. From the linearity of integrals w.r.t. transitional probability kernels, one has \mathbb{P} -a.s. that

$$\begin{aligned} \mathbb{E}\left[\int_{\mathbb{X}} [af(Y, x) + bg(Y, x)] \tilde{P}(dx) | \mathcal{F}_{T+n}\right] &= \mathbb{E}\left[\int_{\mathbb{X}} af(Y, x) \tilde{P}(dx) + \int_{\mathbb{X}} bg(Y, x) \tilde{P}(dx) | \mathcal{F}_{T+n}\right] = \\ &= a \cdot \mathbb{E}\left[\int_{\mathbb{X}} f(Y, x) \tilde{P}(dx) | \mathcal{F}_{T+n}\right] + b \cdot \mathbb{E}\left[\int_{\mathbb{X}} g(Y, x) \tilde{P}(dx) | \mathcal{F}_{T+n}\right] = \end{aligned}$$

¹Let $A \in \mathcal{F}_{m+n}$ and $D = \{T = m\}$. As both $A \cap D \in \mathcal{F}_{m+n}$ and $A \cap D \in \mathcal{F}_{T+n}$, one has

$$\begin{aligned} \int_A \mathbb{E}[f(X_{m+n+1}) | \mathcal{F}_{m+n}](\omega) \mathbb{1}_D(\omega) \mathbb{P}(d\omega) &= \int_{A \cap D} f(X_{m+n+1}(\omega)) \mathbb{P}(d\omega) = \\ &= \int_{A \cap D} f(X_{T(\omega)+n+1}(\omega)) \mathbb{P}(d\omega) = \int_A \mathbb{E}[f(X_{T+n+1}) | \mathcal{F}_{T+n}](\omega) \mathbb{1}_D(\omega) \mathbb{P}(d\omega). \end{aligned}$$

But $\mathbb{E}[f(X_{T+n+1}) | \mathcal{F}_{T+n}] \in M(\mathcal{F}_{T+n})$, so $\mathbb{E}[f(X_{T+n+1}) | \mathcal{F}_{T+n}] \cdot \mathbb{1}_D \in M(\mathcal{F}_{m+n})$ from the properties of stopping times. As a consequence,

$$\mathbb{E}[f(X_{m+n+1}) | \mathcal{F}_{m+n}] \mathbb{1}_D = \mathbb{E}[f(X_{T+n+1}) | \mathcal{F}_{T+n}] \mathbb{1}_D \quad \text{a.s.}[\mathbb{P}].$$

The case of an \mathcal{H} -measurable integrand follows in an analogous way.

$$\begin{aligned}
&= a \cdot \mathbb{E}[f(Y, X_{T+n+1})|\mathcal{F}_{T+n}] + b \cdot \mathbb{E}[g(Y, X_{T+n+1})|\mathcal{F}_{T+n}] = \\
&= \mathbb{E}[af(Y, X_{T+n+1}) + bg(Y, X_{T+n+1})|\mathcal{F}_{T+n}];
\end{aligned}$$

thus, $af + bg \in \mathcal{M}$. Let $(f_m)_{m \geq 1} \subseteq \mathcal{M}$ be such that $f_m \geq 0$ and $f_m \uparrow f$, for some $f \in M_{b,+}(\mathcal{Y} \otimes \mathcal{X})$. From the monotone convergence of integrals w.r.t. transitional probability kernels, one has \mathbb{P} -a.s. that

$$\begin{aligned}
\mathbb{E}\left[\int_{\mathbb{X}} f(Y, x)\tilde{P}(dx)|\mathcal{F}_{T+n}\right] &= \lim_{m \rightarrow \infty} \mathbb{E}\left[\int_{\mathbb{X}} f_m(Y, x)\tilde{P}(dx)|\mathcal{F}_{T+n}\right] = \\
&= \lim_{m \rightarrow \infty} \mathbb{E}[f_m(Y, X_{T+n+1})|\mathcal{F}_{T+n}] = \mathbb{E}[f(Y, X_{T+n+1})|\mathcal{F}_{T+n}],
\end{aligned}$$

where we have used that $\int_{\mathbb{X}} f_m(Y, x)\tilde{P}(dx) \uparrow \int_{\mathbb{X}} f(Y, x)\tilde{P}(dx)$. As a result $f \in \mathcal{M}$, so $\mathcal{M} \supseteq M_b(\mathcal{Y} \otimes \mathcal{X})$ from a monotone class argument; therefore, for each $f \in M_b(\mathcal{Y} \otimes \mathcal{X})$,

$$\mathbb{E}\left[\int_{\mathbb{X}} f(Y, x)\tilde{P}(dx)|\mathcal{F}_{T+n}\right] = \mathbb{E}[f(Y, X_{T+n+1})|\mathcal{F}_{T+n}] \quad \text{a.s.}[\mathbb{P}], \text{ for } n \geq 1.$$

Let $f \in M_b(\mathcal{Y} \otimes \mathcal{X})$. It follows that $(\mathbb{E}[f(Y, X_{T+n+1})|\mathcal{F}_{T+n}])_{n \geq 0}$ is a uniformly integrable martingale w.r.t. $(\mathcal{F}_{T+n})_{n \geq 0}$, and hence it converges a.s. $[\mathbb{P}]$ and in L^1 . In fact,

$$\mathbb{E}[f(Y, X_{T+n+1})|\mathcal{F}_{T+n}] \longrightarrow \mathbb{E}\left[\int_{\mathbb{X}} f(Y, x)\tilde{P}(dx)|\mathcal{F}_{\infty}\right] = \int_{\mathbb{X}} f(Y, x)\tilde{P}(dx) \quad \text{a.s.}[\mathbb{P}],$$

by noting that $\mathcal{F}_T \subseteq \mathcal{F}_{\infty}$. Now,

$$\lim_{n \rightarrow \infty} \mathbb{E}[f(Y, X_{n+1})|\mathcal{F}_n] \cdot \mathbb{1}_{\{T=m\}} = \lim_{n \rightarrow \infty} \mathbb{E}[f(Y, X_{T+n+1})|\mathcal{F}_{T+n}] \cdot \mathbb{1}_{\{T=m\}} = \int_{\mathbb{X}} f(Y, x)\tilde{P}(dx) \cdot \mathbb{1}_{\{T=m\}},$$

exists \mathbb{P} -a.s. for every $m \geq 1$. From the dominated convergence theorem w.r.t. the counting measure², one has that the following limit exists \mathbb{P} -a.s. as well,

$$\begin{aligned}
\lim_{n \rightarrow \infty} \mathbb{E}[f(Y, X_{n+1})|\mathcal{F}_n] &= \lim_{n \rightarrow \infty} \sum_{m=1}^{\infty} \mathbb{E}[f(Y, X_{n+1})|\mathcal{F}_n] \cdot \mathbb{1}_{\{T=m\}} = \\
&= \sum_{m=1}^{\infty} \lim_{n \rightarrow \infty} \mathbb{E}[f(Y, X_{n+1})|\mathcal{F}_n] \cdot \mathbb{1}_{\{T=m\}} = \int_{\mathbb{X}} f(Y, x)\tilde{P}(dx).
\end{aligned}$$

As $(\mathbb{E}[f(Y, X_{n+1})|\mathcal{F}_n])_{n \geq 0}$ is uniformly integrable, then

$$\mathbb{E}[f(Y, X_{n+1})|\mathcal{F}_n] \longrightarrow \int_{\mathbb{X}} f(Y, x)\tilde{P}(dx) \quad \text{a.s.}[\mathbb{P}] \text{ and in } L^1.$$

On the other hand, by Lemma A.7,

$$\frac{1}{n} \sum_{i=1}^n f(Y, X_i) \longrightarrow \int_{\mathbb{X}} f(Y, x)\tilde{P}(dx) \quad \text{a.s.}[\mathbb{P}].$$

□

²Indeed, as f is bounded, then $|\mathbb{E}[f(Y, X_{n+1})|\mathcal{F}_n] \cdot \mathbb{1}_{\{T=m\}}| \leq M \mathbb{1}_{\{T=m\}}$ a.s. $[\mathbb{P}]$ for some $M \in \mathbb{R}$ and each $m \geq 1$, and it holds that

$$\sum_{m=1}^{\infty} M \mathbb{1}_{\{T=m\}} = M \mathbb{1}_{\{T < \infty\}} = M < \infty \quad \text{a.s.}[\mathbb{P}].$$

Central limit results

Theorem A.1 (A.S. Conditional Convergence of Martingales). *Let $(\Omega, \mathcal{H}, \mathbb{P})$ be a probability space, $\mathcal{G} = (\mathcal{G}_n)_{n \geq 0}$ be a filtration on (Ω, \mathcal{H}) , $(M_n)_{n \geq 1}$ be a real-valued \mathcal{G} -martingale, and $H_n \in \mathcal{G}_n$ be such that $\mathbb{P}(H_n^c \text{ i.o.}) = 0$. Suppose that $M_n \rightarrow M$ in L^1 , for some $M \in M(\mathcal{H})$. If it holds*

- (i) $\mathbb{E}[\sup_{n \in \mathbb{N}} \sqrt{n} \cdot \mathbb{1}_{H_n} |M_n - M_{n+1}|] < \infty$;
- (ii) $n \cdot \sum_{m \geq n} (M_m - M_{m+1})^2 \xrightarrow{\text{a.s.}} U$, for some $U \in M_+(\mathcal{H})$;

then

$$\sqrt{n}(M_n - M) \xrightarrow{\text{a.s. cond.}} \mathcal{N}(0, U) \quad \text{w.r.t. } \mathcal{G}.$$

Proof. The proof itself uses Theorem A.1 in Crimaldi (2009), which we state here in a shortened version.

Theorem: *For each $n \geq 1$, let $(M_{n,h})_{h \geq 0}$ be a real-valued martingale w.r.t. a filtration $(\mathcal{F}_{n,h})_{h \geq 0}$ with $M_{n,0} = 0$ and such that $M_{n,h} \rightarrow M_{n,\infty}$ in L^1 , for some $M_{n,\infty} \in M(\mathcal{H})$. Suppose $(\mathcal{F}_{n,1})_{n \geq 1}$ is a filtration on (Ω, \mathcal{H}) . Denote by \mathcal{U} the completion of $\bigvee_{n \in \mathbb{N}} \mathcal{F}_{n,1}$ in \mathcal{H} ,*

$$X_{n,h} := M_{n,h} - M_{n,h-1}, \quad U_n := \sum_{h=1}^{\infty} X_{n,h}^2, \quad X_n^* := \sup_{h \in \mathbb{N}} |X_{n,h}|, \quad \text{for } h, n \geq 1.$$

If it holds

- (a) $X_n^* \xrightarrow{\text{a.s.}} 0$;
- (b) $(X_n^*)_{n \geq 1}$ is dominated in L^1 ;
- (c) $U_n \xrightarrow{\text{a.s.}} U$, for some $U \in M_+(\mathcal{U})$;

then

$$M_{n,\infty} \xrightarrow{\text{a.s. cond.}} \mathcal{N}(0, U) \quad \text{w.r.t. } \mathcal{G}.$$

where in our case $M_{n,0} = M_{n,1} := 0$, $\mathcal{F}_{n,0} = \mathcal{F}_{n,1} := \mathcal{G}_n$ and, for $h \geq 2$,

$$\mathcal{F}_{n,h} := \mathcal{G}_{n+h-1}, \quad M_{n,h} := \sqrt{n} \cdot \mathbb{1}_{H_n} (M_n - M_{n+h-1}).$$

It follows that

$$\mathbb{E}[M_{n,h+1} | \mathcal{F}_{n,h}] = \sqrt{n} \cdot \mathbb{1}_{H_n} (M_n - \mathbb{E}[M_{n+h} | \mathcal{G}_{n+h-1}]) = M_{n,h};$$

thus, $(M_{n,h})_{h \geq 1}$ is a martingale w.r.t. $(\mathcal{F}_{n,h})_{h \geq 0}$ such that

$$M_{n,h} \longrightarrow \sqrt{n} \cdot \mathbb{1}_{H_n} (M_n - M) =: M_{n,\infty} \quad \text{in } L^1, \text{ as } h \rightarrow \infty.$$

On the other hand, $X_{n,h} = \sqrt{n} \cdot \mathbb{1}_{H_n} (M_{n+h-2} - M_{n+h-1})$, so

$$\sum_{h=1}^{\infty} X_{n,h}^2 = n \cdot \mathbb{1}_{H_n} \sum_{m \geq n} (M_m - M_{m+1})^2 \longrightarrow U \quad \text{a.s.}[\mathbb{P}],$$

and

$$X_n^* = \sqrt{n} \cdot \mathbb{1}_{H_n} \sup_{m \geq n} |M_m - M_{m+1}| \leq \mathbb{1}_{H_n} \sup_{m \geq n} \sqrt{m} |M_m - M_{m+1}| \leq \sup_{n \in \mathbb{N}} \sqrt{n} \cdot \mathbb{1}_{H_n} |M_n - M_{n+1}| \quad \text{a.s.}[\mathbb{P}],$$

where we have used $\mathbb{1}_{H_n} \xrightarrow{a.s.} 1$ for both results. Moreover,

$$n(M_n - M_{n+1})^2 = n \cdot \sum_{m \geq n} (M_m - M_{m+1})^2 - \frac{n}{n+1} \sum_{m \geq n+1} (M_m - M_{m+1})^2 \longrightarrow 0 \quad \text{a.s.}[\mathbb{P}],$$

which implies $\sup_{m \geq n} \sqrt{m} |M_m - M_{m+1}| \xrightarrow{a.s.} 0$, and thus $X_n^* \xrightarrow{a.s.} 0$. As a result of the above theorem,

$$\sqrt{n} \cdot \mathbb{1}_{H_n} (M_n - M) \xrightarrow{a.s. \text{ cond.}} \mathcal{N}(0, U) \quad \text{w.r.t. } \mathcal{G}.$$

But $\mathbb{1}_{H_n} \xrightarrow{a.s.} 1$ and $H_n \in \mathcal{G}_n$; therefore,

$$\sqrt{n} (M_n - M) \xrightarrow{a.s. \text{ cond.}} \mathcal{N}(0, U) \quad \text{w.r.t. } \mathcal{G}.$$

□

Bibliography

- Airoldi, E. M., Costa, T., Bassetti, F., Leisen, F., and Guindani, M. (2014). Generalized species sampling priors with latent Beta reinforcements. *Journal of the American Statistical Association*, 109(508):1466–1480.
- Aldous, D. J. (1985). Exchangeability and related topics. *École d’Été de Probabilités de Saint-Flour XIII 1983*, 1117:1–198.
- Aletti, G., Ghiglietti, A., and Rosenberger, W. F. (2018a). Nonparametric covariate-adjusted response-adaptive design based on a functional urn model. *Annals of Statistics*, 46(6B):3838–3866.
- Aletti, G., Ghiglietti, A., and Vidyashankar, A. N. (2018b). Dynamics of an adaptive randomly reinforced urn. *Bernoulli*, 24(3):2204–2255.
- Aletti, G., May, C., and Secchi, P. (2007). On the distribution of the limit proportion for a two-color, randomly reinforced urn with equal reinforcement distributions. *Advances in Applied Probability*, 39(3):690–707.
- Aletti, G., May, C., and Secchi, P. (2009). A central limit theorem, and related results, for a two-color randomly reinforced urn. *Advances in Applied Probability*, 41(3):829–844.
- Athreya, K. B. and Ney, P. (1972). *Branching Processes*. Springer, Berlin, Heidelberg.
- Bandyopadhyay, A. and Thacker, D. (2014). Rate of convergence and large deviation for the infinite color Pólya urn schemes. *Statistics and Probability Letters*, 92:232–240.
- Bandyopadhyay, A. and Thacker, D. (2016). A new approach to Pólya urn schemes and its infinite color generalization. arXiv:1606.05317.
- Bandyopadhyay, A. and Thacker, D. (2017). Pólya urn schemes with infinitely many colors. *Bernoulli*, 23(4B):3243–3267.
- Bassetti, F., Casarin, R., and Rossini, L. (2018). Hierarchical species sampling models. arXiv:1803.05793.
- Bassetti, F., Crimaldi, I., and Leisen, F. (2010). Conditionally identically distributed species sampling sequences. *Advances in Applied Probability*, 42(2):433–459.

- Beggs, A. W. (2005). On the convergence of reinforcement learning. *Journal of Economic Theory*, 122(1):1–36.
- Berti, P., Crimaldi, I., Pratelli, L., and Rigo, P. (2009). Rate of convergence of predictive distributions for dependent data. *Bernoulli*, 15(4):1351–1367.
- Berti, P., Crimaldi, I., Pratelli, L., and Rigo, P. (2010). Central limit theorems for multicolor urns with dominated colors. *Stochastic Processes and their Applications*, 120(8):1473–1491.
- Berti, P., Crimaldi, I., Pratelli, L., and Rigo, P. (2011). A central limit theorem and its applications to multicolor randomly reinforced urns. *Journal of Applied Probability*, 48(2):527–546.
- Berti, P., Crimaldi, I., Pratelli, L., and Rigo, P. (2015). Central limit theorems for an Indian buffet model with random weights. *Annals of Applied Probability*, 25(2):523–547.
- Berti, P., Dreassi, E., Pratelli, L., and Rigo, P. (2019). A predictive approach to Bayesian nonparametrics. <http://www-dimat.unipv.it/~rigo/mg.pdf>. Accessed: 2019-11-08.
- Berti, P., Pratelli, L., and Rigo, P. (2004). Limit theorems for a class of identically distributed random variables. *Annals of Probability*, 32(3):2029–2052.
- Berti, P., Pratelli, L., and Rigo, P. (2006). Almost sure weak convergence of random probability measures. *Stochastics*, 78(2):91–97.
- Berti, P., Pratelli, L., and Rigo, P. (2013). Exchangeable sequences driven by an absolutely continuous random measure. *Annals of Probability*, 41(3B):2090–2102.
- Blackwell, D. and Dubins, L. (1962). Merging of opinions with increasing information. *Annals of Mathematical Statistics*, 33(3):882–886.
- Blackwell, D. and MacQueen, J. B. (1973). Ferguson distributions via Pólya urn schemes. *Annals of Statistics*, 1(2):353–355.
- Camerlenghi, F., Lijoi, A., Orbanz, P., and Prünster, I. (2019). Distribution theory for hierarchical processes. *Annals of Statistics*, 47(1):67–92.
- Cassese, A., Zhu, W., Guindani, M., and Vannucci, M. (2019). A Bayesian nonparametric spiked prior for dynamic model selection. *Bayesian Analysis*, 14(2):553–572.
- Cinlar, E. (2011). *Probability and Stochastics*. Springer, New York.
- Crimaldi, I. (2009). An almost sure conditional convergence result and an application to a generalized Pólya urn. *International Mathematical Forum*, 4(23):1139–1156.
- Crimaldi, I., Letta, G., and Pratelli, L. (2007). A strong form of stable convergence. In Donati-Martin, C., Émery, M., Rouault, A., and Stricker, C., editors, *Séminaire de Probabilités XL. Lecture Notes in Mathematics*, volume 1899, pages 203–225. Springer, Berlin, Heidelberg.

- de Finetti, B. (1931). Funzione caratteristica di un fenomeno aleatorio. *Atti della R. Accademia Nazionale dei Lincei, Serie 6. Memorie, Classe di Scienze Fisiche, Matematiche e Naturale*, 4:251–299.
- de Finetti, B. (1937). La Pr evision: ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincar e*, 7(1):1–68.
- Durham, S. C., Flournoy, N., and Li, W. (1998). A sequential design for maximizing the probability of a response. *Canadian Journal of Statistics*, 26(3):479–495.
- Durham, S. C. and Yu, K. F. (1990). Randomized play-the-leader rules for sequential sampling from two populations. *Probability in the Engineering and Information Sciences*, 4(3):355–367.
- Durrett, R. (2010). *Probability: Theory and Examples*. Cambridge University Press, Cambridge, UK, 4 edition.
- Favaro, S., Lijoi, A., and Pr unster, I. (2012a). Asymptotics for a Bayesian nonparametric estimator of species variety. *Bernoulli*, 18(4):1267–1283.
- Favaro, S., Lijoi, A., and Pr unster, I. (2012b). A new estimator of the discovery probability. *Biometrics*, 68(4):1188–1196.
- Flournoy, N., May, C., and Secchi, P. (2012). Asymptotically optimal response-adaptive designs for allocating the best treatment: an overview. *International Statistical Review*, 80(2):293–305.
- Fortini, S. and Petrone, S. (2012a). Hierarchical reinforced urn processes. *Statistical & Probability Letters*, 82(8):1521–1529.
- Fortini, S. and Petrone, S. (2012b). Predictive construction of priors in Bayesian nonparametrics. *Brazilian Journal of Probability and Statistics*, 26(4):423–449.
- Fortini, S. and Petrone, S. (2019). Quasi-Bayes properties of a recursive procedure for mixtures. arXiv:1902.10708.
- Fortini, S., Petrone, S., and Sporysheva, P. (2018). On a notion of partially conditionally identically distributed sequences. *Stochastic Processes and their Applications*, 128(3):819–846.
- Freedman, D. (1965). Bernard Friedman’s urn. *Annals of Mathematical Statistics*, 36(3):956–970.
- Ghosal, S. and van der Vaart, A. (2017). *Fundamentals of Nonparametric Bayesian Inference*. Cambridge University Press, Cambridge, UK.
- Gibbs, A. D. and Su, F. E. (2002). On choosing and bounding probability metrics. *International Statistical Review*, 70(3):419–435.
- Griffiths, T. L. and Ghahramani, Z. (2011). The Indian buffet process: an introduction and review. *Journal of Machine Learning Research*, 12:1185–1224.
- Hansen, B. and Pitman, J. (2000). Prediction rules for exchangeable sequences related to species sampling. *Statistics & Probability Letters*, 46(3):251–256.

- Häusler, E. and Luschgy, H. (2015). *Stable Convergence and Stable Limit Theorems*. Springer, Cham, Switzerland.
- Hewitt, E. and Savage, L. J. (1955). Symmetric measures on Cartesian products. *Transactions of the American Mathematical Society*, 80:470–501.
- Hoppe, F. M. (1984). Pólya-like urns and the Ewens' sampling formula. *Journal of Mathematical Biology*, 20(1):91–94.
- Hu, F. and Rosenberger, W. F. (2006). *The Theory of Response-Adaptive Randomization in Clinical Trials*. Wiley, Hoboken, NJ.
- Janson, S. (2018). A.s. convergence for infinite colour Pólya urns associated with random walks. arXiv:1803.04207.
- Janson, S. (2019). Random replacements in Pólya urns with infinitely many colours. *Electronic Communications in Probability*, 24. paper no. 23, 11 pp.
- Kallenberg, O. (1988). Spreading and predictable sampling in exchangeable sequences and processes. *Annals of Probability*, 16(2):508–534.
- Kallenberg, O. (2002). *Foundations of Modern Probability*. Springer, New York, 2 edition.
- Kallenberg, O. (2017). *Random Measures, Theory and Applications*. Springer, Cham, Switzerland.
- Krasker, W. S. and Pratt, J. W. (1986). Discussion: On the consistency of Bayes estimates. *Annals of Statistics*, 14(1):55–58.
- Li, W., Durham, S. D., and Flournoy, N. (1996). Randomized Pólya urn designs. In *Proceedings of the Biometric Section of the American Statistical Association*, pages 166–170.
- Mahmoud, H. (2008). *Pólya Urn Models*. CRC Press, Boca Raton, FL.
- Mailler, C. and Marckert, J.-F. (2017). Measure-valued Pólya urn processes. *Electronic Communications in Probability*, 22. paper no. 26, 33 pp.
- Martin, C. F. and Ho, Y. C. (2002). Value of information in the Pólya urn process. *Information Sciences*, 147:65–90.
- May, C. and Flournoy, N. (2009). Asymptotics in response-adaptive designs generated by a two-color, randomly reinforced urn. *Annals of Statistics*, 37(2):1058–1078.
- May, C., Paganoni, A. M., and Secchi, P. (2005). On a two-color generalized Pólya urn. *Metron*, 63(1):115–134.
- Muliere, P., Paganoni, A. M., and Secchi, P. (2006). A randomly reinforced urn. *Journal of Statistical Planning and Inference*, 136(6):1853–1874.
- Paganoni, A. M. and Secchi, P. (2004). Interacting reinforced-urn systems. *Advances in Applied Probability*, 36(3):791–804.

- Pemantle, R. (1990). A time-dependent version of Pólya's urn. *Journal of Theoretical Probability*, 3(4):627–637.
- Pemantle, R. (2007). A survey on random processes with reinforcement. *Probability Surveys*, 4:1–79.
- Pemantle, R. and Volkov, S. (1999). Vertex-reinforced random walk on \mathbb{Z} has finite range. *Annals of Probability*, 27(3):1368–1388.
- Petrone, S. and Raftery, A. E. (1997). A note on the Dirichlet process prior in Bayesian nonparametric inference with partial exchangeability. *Statistical & Probability Letters*, 36(1):69–83.
- Pitman, J. (1995). Exchangeable and partially exchangeable random partitions. *Probability Theory and Related Fields*, 102(2):145–158.
- Pitman, J. (1996). Some developments of the Blackwell-MacQueen urn scheme. In *Lecture Notes-Monograph Series*, volume 30, pages 245–267. Institute of Mathematical Statistics, Hayward, CA.
- Pitman, J. (2010). Combinatorial Stochastic Processes. In *Lecture Notes in Mathematics*, volume 1875. Springer, Berlin, Heidelberg.
- Pitman, J. and Yor, M. (1997). The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator. *Annals of Probability*, 25(2):855–900.
- Pólya, G. (1930). Sur quelques points de la théorie des probabilités. *Annales de l'Institut Henri Poincaré*, 1(2):117–161.
- Rosenberger, W. F. (2002). Randomized urn models and sequential design. *Sequential Analysis*, 21(1):1–28.
- Rosenberger, W. F. and Lachin, J. M. (2002). *Randomization in Clinical Trials: Theory and Practice*. Wiley, New York.
- Teh, Y. W., Jordan, M., Beal, M., and Blei, D. (2006). Hierarchical Dirichlet processes. *Journal of the American Statistical Association*, 101(476):1566–1581.
- Walker, S. and Muliere, P. (2003). A bivariate Dirichlet process. *Statistics & Probability Letters*, 64(1):1–7.
- Williams, D. (1991). *Probability with Martingales*. Cambridge University Press, Cambridge, UK.
- Zhang, L.-X. (2015). Response-adaptive randomization: an overview of designs and asymptotic theory. In Sverdlov, O., editor, *Modern Adaptive Randomized Clinical Trials-Statistical and Practical Aspects*, pages 221–250. CRC Press, Boca Raton, FL.
- Zhang, L.-X., Hu, F., Cheung, S. H., and Chan, W. S. (2011). Immigrated urn models - theoretical properties and applications. *Annals of Statistics*, 39(1):643–671.
- Zhang, L.-X., Hu, F., Cheung, S. H., and Chan, W. S. (2014). Asymptotic properties of multicolor randomly reinforced Pólya urns. *Advances in Applied Probability*, 46(2):585–602.