INTERACTING REINFORCED URN SYSTEMS

ANNA MARIA PAGANONI AND PIERCESARE SECCHI

ABSTRACT. We introduce a class of discrete time stochastic processes generated by interacting systems of reinforced urns. We show that such processes are asymptotically partially exchangeable and we prove a strong law of large numbers. Examples and the analysis of particular cases show that interacting reinforced urn systems are very flexible representations for modelling countable collections of dependent and asymptotically exchangeable sequences of random variables.

1. INTRODUCTION

This paper introduces a class of discrete time stochastic processes generated by interacting systems of reinforced urns.

The prototypical example of a reinforced urn is the *Polya urn* of Eggenberg and Polya (1923); a single urn initially contains a given number of balls of two different colors. At time $n = 0, 1, 2, ...$, a ball is sampled from the urn and reintroduced in it together with a constant, nonnegative number m of balls of the same color. The scheme has suggested a number of interesting variations and extensions: see for instance [2, 6, 7, 8, 14]. Of particular importance for the the present work is the extension in [9] that keeps the idea of reinforcing at each stage only the color currently observed, but assumes that at each time n the number m of balls added to the urn is the realization of a random variable M_n independent of everything observed in the past and of the color currently sampled from the urn; the process of colors generated by this scheme is called *generalized Polya sequence.*

Our typical interacting reinforced urn system consists of a countable number of stochastic processes which, in the absence of interaction, would each be a generalized Polya sequence. The construction of these processes will be carried out in the next section together with the proof of three main results regarding their laws.

1991 *Mathematics Subject Classification.* 60F15,60G09.

We thank an anonymous referee for comments that contributed to improve the paper.

In Bayesian statistics, considerable attention has been dedicated to the Polya urn as a scheme that implements a classical statistical model for infinite sequences of Bernoulli random variables, conditionally independent and identically distributed given their (random) probability of success Θ to which is assigned a Beta prior probability distribution. This has stimulated the study of more complex processes generated by reinforced urns and implementing Bayesian statistical models: see [10, 11, 12, 16]. Some of these processes are specific examples of interacting reinforced urn systems: for instance, the two-color reinforced urn process of [10] that generates an infinite sequence of exchangeable survival times with beta-Stacy prior, a model widely used in Bayesian nonparametric survival analysis. Together with other examples of systems of interacting reinforced urns, two-color reinforced urn processes on the integers will be considered in the third section of the paper while their extension generated by allowing for random reinforcements will be the fourth section's topic. A section on final remarks and open questions concludes the paper.

2. Construction and main results

Let S be a countable set of sites. Every site $s \in S$ labels an urn initially containing $B_0(s)$ balls of color 1 and $W_0(s)$ balls of color 0, with the assumption that $B_0(s)$ and $W_0(s)$ are two nonnegative real numbers summing to a strictly positive number. At time $n = 0, 1, 2, ...,$ a ball is independently sampled from every urn in the system. Call $X_n(s)$ the color generated by the urn with label $s \in S$ and let $X_n = (X_n(s), s \in S)$; next a (real) number $r(s, X_n, M_{n+1}) \geq 0$ of balls of the same color as $X_n(s)$ is introduced in the urn with label s together with the ball extracted. The reinforcement $r(s, X_n, M_{n+1})$ is assumed to be conditionally independent of the color $X_n(s)$ given the past realizations of $X_0, ..., X_{n-1}$ and of $M_1, ..., M_n$; in general, its distribution will depend on the colors generated by the urns labelled by sites different from s as well as on a random element (disturbance) M_{n+1}, independent of $X_0, ..., X_n$ and of $M_1, ..., M_n$ and with values in an appropriate measurable space. In this way we generate a discrete time process $X = (X_n, n = 0, 1, ...)$ with state space $\{0,1\}^S$; X is here baptized *interacting reinforced urn system.*

More formally, given a rich enough probability space $(\Omega, \mathcal{F}, P)$ and a countable set of sites S, an interacting reinforced urn system X, defined on Ω and with values in the state space $\{0,1\}^S$ endowed with its product topology and Borel sigma-field, is obtained by specifying the following three elements:

i. A family $U_0 = ((B_0(s), W_0(s)), s \in S)$ of initial urn compositions, with $B_0(s), W_0(s) \in [0, \infty)$, $B_0(s) + W_0(s) > 0$ for every $s \in S$;
ii. The common probability distribution μ of an infinite sequence $M_1, M_2, \ldots$ of i.i.d. random elements, defined on Ω and with values in a suitable metric, separable and complete space $\mathcal{M}$ endowed with its Borel sigma-field;
iii. A reinforcement rule r, i.e. a bounded, nonnegative, measurable function defined on $S \times \{0,1\}^S \times \mathcal{M}$.

These three elements define the process X whose local dynamics are henceforth described. For $s \in S$, set

$$Z_0(s) = \frac{B_0(s)}{B_0(s) + W_0(s)}$$

and let

$$X_0 = (X_0(s), s \in S)$$

be a collection of independent random variables such that $X_0(s)$ has Bernoulli($Z_0(s)$) distribution; assume that $M_1 \in \mathcal{M}$ is independent of X_0 with probability distribution μ and, for $s \in S$, suppose that $r(s, X_0, M_1)$ is independent of $X_0(s)$. For $n \geq 1$ and $s \in S$, set

$$\begin{aligned} B_n(s) &= B_{n-1}(s) + X_{n-1}(s) r(s, X_{n-1}, M_n), \\ W_n(s) &= W_{n-1}(s) + (1 - X_{n-1}(s)) r(s, X_{n-1}, M_n), \\ Z_n(s) &= \frac{B_n(s)}{B_n(s) + W_n(s)}; \end{aligned}$$

given the sigma-field $\mathcal{F}_n$ generated by $X_0, \ldots, X_{n-1}$ and $M_1, \ldots, M_n$, let

$$X_n = (X_n(s), s \in S)$$

be a collection of conditionally independent random variables such that $X_n(s)$ has Bernoulli($Z_n(s)$) distribution; assume that $M_{n+1} \in \mathcal{M}$ is independent of $\mathcal{F}_n$ and of X_n with probability distribution μ and, for $s \in S$, suppose that $r(s, X_n, M_{n+1})$ is conditionally independent of $X_n(s)$ given $\mathcal{F}_n$.

When the set of sites S consists of a single element s, $\mathcal{M} = [0, \infty)$ and $r(s, x, m) = m$ for every $(s, x, m) \in S \times \{0,1\}^S \times \mathcal{M}$, then the process X reduces to a generalized Polya sequence with parameters $(B_0(s), W_0(s), \mu)$. In fact, the theorems of this section show that the main results true for generalized Polya sequences hold for general interacting reinforced urn systems as well. For $n \geq 0$, let

$$Z_n = (Z_n(s), s \in S)$$

be the random element with values in $[0,1]^S$ describing the proportion of balls of color 1 present in the urns of the system at time n just before they are sampled for the $(n+1)$-th time.

Theorem 2.1. *The process $\{Z_n\}$ is martingale, with respect to the filtration $\{\mathcal{F}_n\}$, with values in $[0,1]^S$. Therefore it converges almost surely to a random element $Z_\infty \in [0,1]^S$.*

Proof. Fix $s \in S$ and compute

$$\begin{aligned}
&\mathbf{E}[Z_{n+1}(s)|\mathcal{F}_n] \\
&\quad = \mathbf{E}\left[\frac{B_n(s)}{B_n(s)+W_n(s)}\frac{B_n(s)+r(s,X_n,M_{n+1})}{B_n(s)+W_n(s)+r(s,X_n,M_{n+1})} + \right. \\
&\qquad\qquad \left. + \frac{W_n(s)}{B_n(s)+W_n(s)}\frac{B_n(s)}{B_n(s)+W_n(s)+r(s,X_n,M_{n+1})}\right|\mathcal{F}_n\Bigg] \\
&\quad = \mathbf{E}\left[\frac{B_n(s)(B_n(s)+W_n(s)+r(s,X_n,M_{n+1}))}{(B_n(s)+W_n(s))(B_n(s)+W_n(s)+r(s,X_n,M_{n+1}))}\middle|\mathcal{F}_n\right] \\
&\quad = \mathbf{E}\left[\frac{B_n(s)}{B_n(s)+W_n(s)}\middle|\mathcal{F}_n\right] = Z_n(s).
\end{aligned}$$

Note that the first equality holds because $r(s,X_n,M_{n+1})$ and $X_n(s)$ are conditionally independent given $\mathcal{F}_n$, while the last is true because Z_n is measurable with respect to $\mathcal{F}_n$, the sigma field generated by $X_0,...,X_{n-1}$ and $M_1,...,M_n$. The previous lines show that, for every $s \in S$, $\{Z_n(s)\}$ is a martingale with respect to the filtration $\mathcal{F}_n$ with values in $[0,1]$. Therefore, there exists $Z_\infty = (Z_\infty(s), s \in S)$ such that, for $s \in S$,

$$\lim_{n\to\infty} Z_n(s) = Z_\infty(s)$$

on a set of probability one. Since S is countable, this implies that $\{Z_n\}$ converges to Z_∞ almost surely. ∎

The sequence of colors generated by a Polya urn is exchangeable. A generalized Polya sequence is asymptotically exchangeable, but it needs not be exchangeable. The next theorem proves that interacting reinforced urn processes are asymptotically partially exchangeable.

For $p \in [0,1]^S$, let $\nu(p)$ be the unique product probability defined on the elements of the product sigma-field of $\{0,1\}^S$ such that, for $j \geq 1$, $t_1,...,t_j \in \{0,1\}$ and $s_1,...,s_j \in S$,

$$\nu(p)(\{y \in \{0,1\}^S : y(s_1)=t_1,...,y(s_j)=t_j\}) = \mathbf{E}[\prod_{i=1}^{j} p(s_i)^{t_i}(1-p(s_i))^{1-t_i}].$$

We say that $Y = (Y_n, n = 1, 2, ...)$ is a partially exchangeable process with state space $\{0,1\}^S$ and de Finetti measure $\nu(\Theta)$ if there exists a random element $\Theta \in \{0,1\}^S$ such that, for every $k \geq 1$ and $A_1, ..., A_k$ belonging to the product sigma-field of $\{0,1\}^S$,

$$P[Y_1 \in A_1, ..., Y_k \in A_k] = \mathbf{E}[\prod_{i=1}^{k} \nu(\Theta)(A_i)]. \tag{2.1}$$

When S is finite, we may equivalently rewrite representation (2.1) as: for every $k \geq 1$ and $y_1, ..., y_k \in \{0,1\}^S$,

$$P(Y_1 = y_1, ..., Y_k = y_k) = \mathbf{E}[\prod_{s \in S} \Theta(s)^{\sum_{i=1}^{k} y_i(s)} (1 - \Theta(s))^{k - \sum_{i=1}^{k} y_i(s)}].$$

If a process $X = (X_n, n = 0, 1, 2, ...)$ with state space $\{0,1\}^S$ is such that, for every $k \geq 1$ and for n going to infinity, the probability distribuiton of $(X_{n+1}, ..., X_{n+k})$ converges to that of $(Y_1, ..., Y_k)$, where $Y = (Y_n, n = 1, 2, ...)$ is a partially exchangeable process on $\{0,1\}^S$ satisfying (2.1), we say that X is asymptotically partially exchangeable and $\nu(\Theta)$ is the de Finetti measure associated to X.

Theorem 2.2. *An interacting reinforced urn system X is asymptotically partially exchangeable and the de Finetti measure associated to X is $\nu(Z_\infty)$.*

Proof. For $n \geq 1$, the conditional probability distribution of $(X_n, M_{n+1}) \in \{0,1\}^S \times \mathcal{M}$ given the sigma field $\mathcal{F}_n$ is the product probability $\nu(Z_n) \times \mu$. Because of Theorem 2.1, $\nu(Z_n) \times \mu$ converges to $\nu(Z_\infty) \times \mu$ on a subset of Ω with probability one. Hence, the thesis follows from Lemma 8.2.b in [1].

■

Therefore, conditionally on Z_∞, the colors generated by the urn labelled by $s \in S$ are asymptotically i.i.d. with distribution Bernoulli$(Z_\infty(s))$ and sequences of colors generated by different urns are asymptotically independent. However, Example 1 shows that the random probabilities $(Z_\infty(s), s \in S)$ need not be independent.

The final result of the section proves that the law of large numbers holds for interacting reinforced urn systems.

Theorem 2.3.

$$\lim_{n \to \infty} \frac{\sum_{i=0}^{n-1} X_i}{n} = Z_\infty$$

on a set of probability one.

Proof. The proof is the same as that for generalized Polya sequences in [9]. For ease of reference, we sketch the argument. Fix $s \in S$, set $L_0(s) = 0$ and, for $n \geq 1$, define

$$L_n(s) = nZ_n(s) - \sum_{i=0}^{n-1} X_i(s).$$

Then $\{L_n(s)\}$ is a martingale with respect to the filtration $\{\mathcal{F}_n\}$. Moreover one can show that

$$\sum_{n=1}^{\infty} \frac{E[(L_n(s) - L_{n-1}(s))^2]}{n^2} < \infty.$$

Then it follows from Burkholder's inequality and Theorem 2.1 that:

$$\lim_{n\to\infty} \frac{L_n(s)}{n} = Z_\infty(s) - \lim_{n\to\infty} \frac{\sum_{i=0}^{n-1} X_i(s)}{n} = 0$$

on a set of probability one. Since S is countable, this proves that

$$\lim_{n\to\infty} \frac{\sum_{i=0}^{n-1} X_i}{n} = Z_\infty$$

on a set of probability one. ∎

Remark 2.4. Theorems 2.1, 2.2 and 2.3 still hold even when reinforcements are generated by a more general rule that computes them as a function of time and of past history of the process. More precisely, for $s \in S$ and $n = 0, 1, \ldots$ we may substitute $r(s, X_n, M_{n+1})$ with a bounded, nonnegative random variable $R_n(s)$ measurable with respect to $\mathcal{F}_{n+1}$ and such that, $R_n(s)$ and $X_n(s)$ are conditionally independent given $\mathcal{F}_n$.

Moreover, we note that martingality of the process $\{Z_n\}$ is not a necessary condition for proving Theorem 2.2 and 2.3: in fact, as long as the dynamics of the process X are such that, for all $s \in S$ and $n \geq 0$, the conditional distribution of $X_{n+1}(s)$ given $\mathcal{F}_n$ is Bernoulli $(Z_n(s))$ and the sequence $\{Z_n(s)\}$ converges to $Z_\infty(s)$ almost surely, Theorem 2.2 and 2.3 still hold. For proving the last result one could, for instance, use a generalized version of Borel-Cantelli Lemma as in Corollaire VII-2-6 in [13].

3. Examples

Example 1. *A coupled Polya urn.* Let $S = \{I, II\}$, $U_0 = ((1,1),(1,1))$, $M_1 \equiv M_2 \equiv \cdots \equiv 1$ and, for all $(x,y) \in \{0,1\}^S$, set $r(I,(x,y),1) = y$ and $r(II,(x,y),1) = x$. Hence each urn is reinforced with a ball of the same color as that currently extracted only at times when the other urn has produced a ball of color 1; otherwise the urn's composition is left unchanged. Notice that, with this reinforcement rule, for $n = 1, 2, \ldots$, $Z_n(I) > Z_{n-1}(I)$ if and only if $Z_n(II) > Z_{n-1}(II)$.

We show that, in this case, the limits $Z_\infty(I)$ and $Z_\infty(II)$ are not independent by proving that $\mathrm{Cov}(Z_\infty(I), Z_\infty(II)) > 0$. Observe that, given $\mathcal{F}_n$, the conditional distribution of the product $Z_{n+1}(I)Z_{n+1}(II)$ concentrates its mass on four different values:

$$\begin{cases} Z_n(I)Z_n(II) & \text{with probability } (1 - Z_n(I))(1 - Z_n(II)) \\ Z_n(I)\frac{B_n(II)}{B_n(II)+W_n(II)+1} & \text{with probability } Z_n(I)(1 - Z_n(II)) \\ \frac{B_n(I)}{B_n(I)+W_n(I)+1}Z_n(II) & \text{with probability } (1 - Z_n(I))Z_n(II) \\ \frac{B_n(I)+1}{B_n(I)+W_n(I)+1}\frac{B_n(II)+1}{B_n(II)+W_n(II)+1} & \text{with probability } Z_n(I)Z_n(II) \end{cases}$$

Hence, for $n \geq 0$,

$$\begin{aligned} &\mathbf{E}[Z_{n+1}(I)Z_{n+1}(II)|\mathcal{F}_n] \\ &\quad = Z_n(I)Z_n(II) + \frac{Z_n(I)Z_n(II)(1 - Z_n(I))(1 - Z_n(II))}{(B_n(I) + W_n(I) + 1)(B_n(II) + W_n(II) + 1)} \end{aligned}$$

which shows that $\{Z_n(I)Z_n(II)\}$ is a submartingale with respect to the filtration $\{\mathcal{F}_n\}$. For $n \geq 0$, set

$$a_n = \mathbf{E}\left[\frac{Z_n(I)Z_n(II)(1 - Z_n(I))(1 - Z_n(II))}{(B_n(I) + W_n(I) + 1)(B_n(II) + W_n(II) + 1)}\right] \geq 0;$$

then

$$\mathbf{E}[Z_{n+1}(I)Z_{n+1}(II)] = Z_0(I)Z_0(II) + \sum_{k=0}^{n} a_k = \frac{1}{4} + \sum_{k=0}^{n} a_k.$$

Therefore, by Dominated Convergence Theorem,

$$\mathbf{E}[Z_\infty(I)Z_\infty(II)] = \frac{1}{4} + \sum_{k=0}^{\infty} a_k$$

while the fact that both $\{Z_n(I)\}$ and $\{Z_n(II)\}$ are martingales implies that

$$E[Z_\infty(I)] = Z_0(I) = \frac{1}{2} = Z_0(II) = E[Z_\infty(II)].$$

Thus

$$\mathrm{Cov}(Z_\infty(I), Z_\infty(II)) = \sum_{k=0}^{\infty} a_k \geq a_0 = \frac{1}{144}.$$

To illustrate the example, we generated by simulation 1000 observations from the joint distribution of $(Z_n(I), Z_n(II))$, with $n = 5000$; an image in heat colors and a perspective plot of the joint frequency distribution of these 1000 couples of numbers, as well as the two marginal frequency distributions of $Z_n(I)$ and of $Z_n(II)$ appear in Figure 1. These graphics illustrate the dependence between $Z_n(I)$ and $Z_n(II)$ when n is large. Sample correlation is equal to 0.4532 while sample means and variances for $Z_n(I)$ and $Z_n(II)$ are summarized in the following table:

	Sample mean	Sample variance
$Z_n(I)$	0.5049	0.0817
$Z_n(II)$	0.4983	0.0851

The simulation supports the conjecture that $Z_\infty(I)$ and $Z_\infty(II)$ are identically distributed with Uniform distribution on $[0,1]$, something we resisted the temptation to prove.

Example 2. *A system of two generalized Polya sequences.* Let $S = \{I, II\}$, $U_0 = ((B_0(I), W_0(I)), (B_0(II), W_0(II)))$, $\{M_n\}$ a sequence of i.i.d. nonnegative random variables, uniformly bounded by $K > 0$ and with probability distribution μ, and, for all $(x,y) \in \{0,1\}^S$ and $m \in [0, K]$, set $r(I, (x,y), m) = m$ and $r(II, (x,y), m) = xm$. Hence, at time $n = 1, 2, ...$, urn I is reinforced with a random number M_n of balls of the same color as that currently sampled. At the same time n, urn II is reinforced with M_n balls of the same color as that currently sampled from II only if the ball currently extracted from urn I is of color 1; otherwise II's composition is left unchanged.

In the particular case when μ is the point mass at $\alpha > 0$ and $B_0(I) > 0$, the process of colors generated by urn I is a Polya sequence, and it is therefore exchangeable, while the process of colors generated by urn II needs not be exchangeable. Moreover, in this case: $Z_\infty(I)$ and $Z_\infty(II)$ are uncorrelated (independent?) with distributions $\text{Beta}(\alpha^{-1}B_0(I), \alpha^{-1}W_0(I))$ and $\text{Beta}(\alpha^{-1}B_0(II), \alpha^{-1}W_0(II))$, respectively.

For a general μ, the processes of colors generated by urn I and II are asymptotically exchangeable and $Z_\infty(I)$ and $Z_\infty(II)$ are uncorrelated (independent?), but we are unable to obtain a closed form for their distributions. In fact, the distribution of $Z_\infty(I)$ is the de Finetti measure associated to a generalized Polya sequence of parameters $(B_0(I), W_0(I), \mu)$. In order to prove that the distribution of $Z_\infty(II)$ is the de Finetti measure associated to a generalized Polya sequence

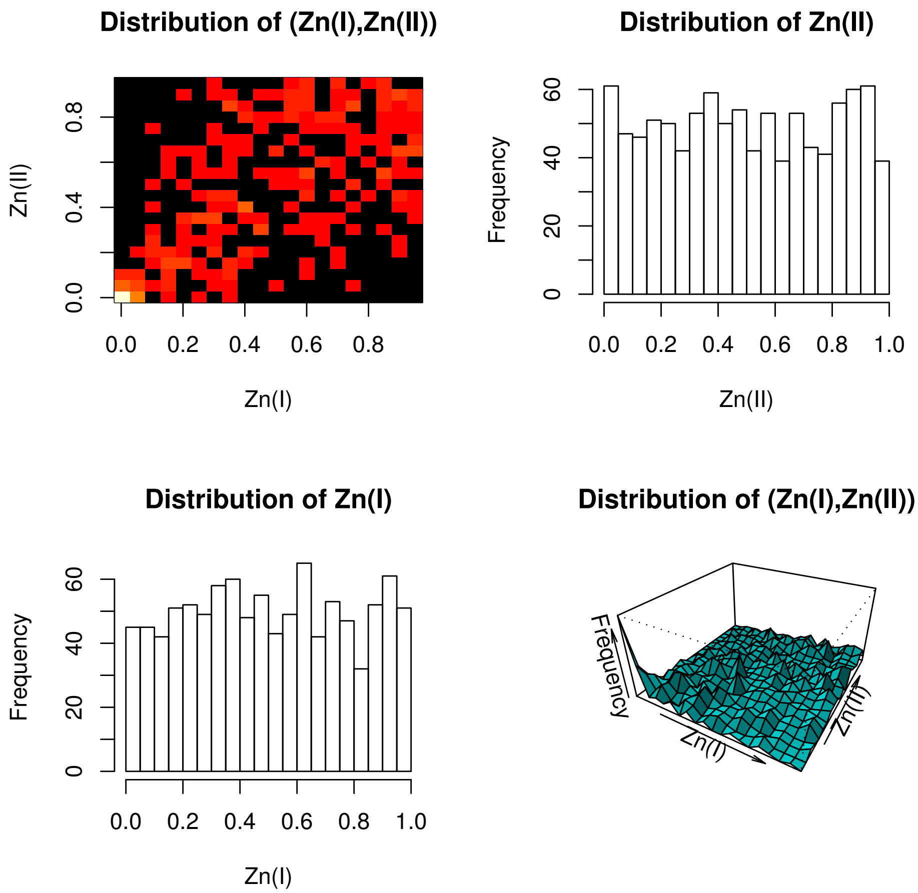

FIGURE 1. Image in heat colors and perspective plot of the joint frequency distribution of 1000 couples of values generated by simulation from the distribution of $(Z_n(I), Z_n(II))$ with $n = 5000$.

with parameters $(B_0(II), W_0(II), \mu)$, set

$$\tau_1 = \inf\{n \geq 0 : X_n(I) = 1\}$$

and, for $i \geq 2$, let

$$\tau_i = \inf\{n > \tau_{i-1} : X_n(I) = 1\}.$$

Lemma 3.1.

$$P\Big[\bigcap_{i=1}^{\infty}\{\tau_i < +\infty\}\Big] = 1 \qquad (3.1)$$

if and only if $B_0(I) > 0$.

Proof. For $B_0(I) = 0$, $P[\tau_1 = \infty] = 1$ hence one implication is trivial.

For proving the other implication, assume that $B_0(I) > 0$. For $n, j \geq 1$, compute

$$\begin{aligned} P[\tau_1 > n + j] &= P[X_1(I) = 0, ..., X_{n+j}(I) = 0] \\ &\leq P[X_{n+1}(I) = 0, ..., X_{n+j}(I) = 0]; \end{aligned}$$

therefore, for every $j \geq 1$,

$$\begin{aligned} &P[\tau_1 < \infty] \\ &\quad = 1 - \lim_{n\to\infty} P[\tau_1 > n + j] \\ &\quad \geq 1 - \lim_{n\to\infty} P[X_{n+1}(I) = 0, ..., X_{n+j}(I) = 0] \\ &\quad = 1 - E[(1 - Z_\infty(I))^j]. \end{aligned}$$

The last equality follows from Theorem 2.2. If $W_0(I) = 0$, $P[Z_\infty(I) = 1] = 1$. If $W_0(I) > 0$, an easy corollary of Theorem 3 in [14] implies that $Z_\infty(I)$ has no atoms (see also Theorem 3.1 in [9]): in particular $P[Z_\infty(I) = 0] = 0$. Thus, in both cases,

$$\lim_{j\to\infty} E[(1 - Z_\infty(I))^j] = 0.$$

Hence $P[\tau_1 < \infty] = 1$. By induction on i it is now easy to prove (3.1). ■

When (3.1) is satisfied, the sequence $\{M_{\tau_i}\}$ is well defined: note that this is the sequence controlling reinforcements for urn II. In fact, one can show that the random variables of the sequence $\{M_{\tau_i}\}$ are i.i.d. with distribution μ. Hence the sequence of colors $\{X_{\tau_i}(II)\}$ is a generalized Polya sequence with parameters $(B_0(II), W_0(II), \mu)$; the fact that the distribution of $Z_\infty(II)$ is the de Finetti measure associated to this generalized Polya sequence follows from the observation that, on a set of probability one,

$$\lim_{i\to\infty} Z_{\tau_i}(II) = Z_\infty(II).$$

Finally, $Z_\infty(I)$ and $Z_\infty(II)$ are uncorrelated because the sequence $\{Z_n(I)Z_n(II)\}$ is a bounded martingale with respect to the filtration

$\{\mathcal{F}_n\}$. In fact, for $n \geq 1$,

$$\begin{aligned}
&\mathbf{E}[Z_{n+1}(I)Z_{n+1}(II)|\mathcal{F}_n] \\
&\quad = \mathbf{E}[Z_n(I)\frac{B_n(I)+M_{n+1}}{B_n(I)+W_n(I)+M_{n+1}} \cdot \\
&\qquad \cdot\{Z_n(II)\frac{B_n(II)+M_{n+1}}{B_n(II)+W_n(II)+M_{n+1}} + \\
&\qquad\qquad +(1-Z_n(II))\frac{B_n(II)}{B_n(II)+W_n(II)+M_{n+1}}\} + \\
&\qquad +(1-Z_n(I))\frac{B_n(I)}{B_n(I)+W_n(I)+M_{n+1}}Z_n(II)|\mathcal{F}_n] \\
&\quad = E[Z_n(I)Z_n(II)|\mathcal{F}_n] \\
&\quad = Z_n(I)Z_n(II).
\end{aligned}$$

Therefore, by Dominated Convergence Theorem,

$$\mathbf{E}[Z_\infty(I)Z_\infty(II)] = \lim_{n\to\infty}\mathbf{E}[Z_n(I)Z_n(II)] = Z_0(I)Z_0(II)$$

while

$$Z_0(I)Z_0(II) = \lim_{n\to\infty}\mathbf{E}[Z_n(I)]\,\mathbf{E}[Z_n(II)] = \mathbf{E}[Z_\infty(I)]\,\mathbf{E}[Z_\infty(II)]$$

where the next to the last equality follows from Theorem 2.1.

Example 3. *Two-color reinforced urn processes.* Let $S = \{0,1,2,...\}$, $U_0 = ((B_0(s), W_0(s)), s \in S)$ with $W_0(0) = 0$, $M_1 \equiv M_2 \equiv \cdots \equiv 1$ and, for all $s \in S$ and $x \in \{0,1\}^S$, set $r(s,x,1) = \prod_{i<s} x(i)$.

For $n = 0,1,2,...$, define

$$T_n = \inf\{s \in S : X_n(s) = 0\}.$$

As in [10], Lemma 3.23, one can prove that

Lemma 3.2.

$$P[\bigcap_{n=0}^{\infty}\{T_n < \infty\}] = 1 \tag{3.2}$$

if and only if

$$\lim_{s\to\infty}\prod_{j=0}^{s}\frac{B_0(j)}{B_0(j)+W_0(j)} = 0 \tag{3.3}$$

When (3.3) is satisfied, for $n = 0,1,2,...$, let

$$\mathcal{B}_n = (0,1,...,T_n)$$

be the sequence of states of S starting from 0 and labelling all the urns until the first that generates a ball of color 0 at time n : call $\mathcal{B}_n$ the n-th

0-vector for the process X. The space of 0-vectors is countable and is given the discrete topology and the Borel sigma-field. The sequence of 0-vectors is exchangeable: in fact the law of $\{\mathcal{B}_n\}$ is the same as that of the sequence of 0-blocks generated by a reinforced urn process with state space S, set of colors $E = \{0, 1\}$, urn composition function U_0 and law of motion $q : S \times \{0, 1\} \to S$ defined by setting

$$q(k, 1) = k + 1 \ \text{ and } \ q(k, 0) = 1.$$

See [10] for more details about reinforced urn processes. In this paper, it is proved that the random variables of the sequence $\{T_n\}$ are conditionally independent and identically distributed, i.e. exchangeable, given a random distribution function F on S whose law is that of a beta-Stacy process on S with parameters $\{(B_0(s), W_0(s))\}$. Beta-Stacy processes are widely used in Bayesian nonparametric studies as prior distributions, say when modelling the law of the survival times $T_0, T_1, T_2, ...$ of an infinite exchangeable sequence of patients: see [15]. The aim of the next section is to extend this example, and Example 2, by allowing randomly generated reinforcements.

4. Two-color Generalized Reinforced Urn Processes

Let $S = \{0, 1, 2, ...\}$, $U_0 = ((B_0(s), W_0(s)), s \in S)$ with $W_0(0) = 0$. Assume that $\{M_n\}$ is a sequence of independent and identically distributed random elements with values in $[0, K]^S$ with $K > 0$ and probability distribution μ. For $s \in S$, $x \in \{0, 1\}^S$ and $m \in [0, K]^S$, define the reinforcement rule

$$r(s, x, m) = m(s) \prod_{i<s} x(i).$$

These assumptions define an interacting reinforced urn system X that is the same as the one in Example 3 except that now we allow random reinforcements, not necessarily equal along 0-vectors.

Set $\omega = \inf\{s \in S : B_0(s) = 0\}$, where as usual $\omega = \infty$ if $B_0(s) > 0$ for every $s \in S$; let $R = \{s \in S : s \leq \omega\}$. Our aim is to study the law of the sequence of integer valued random variables defined by

$$T_n = \inf\{s \in S : X_n(s) = 0\}$$

for $n = 0, 1, 2, ...$ As in Lemma 3.2, it is still true that

$$P[\bigcap_{n=0}^{\infty} \{T_n < \infty\}] = 1$$

if and only if condition (3.3) holds. In particular, this happens if R is strictly contained in S, i.e. if $B_0(s) = 0$ for some $s \in S$.

As in Example 2, it is easy to prove the following result.

Theorem 4.1. *For $s, t \in S$, the sequence $\{\prod_{i=s}^{s+t} Z_n(i)\}$ is a martingale, with respect to the filtration $\{\mathcal{F}_n\}$.*

Associated to the system X, we consider the sequence of random probabilities $Z_\infty = (Z_\infty(s), s \in S)$ introduced in Theorem 2.1. Theorem 4.1 and Dominated Convergence Theorem imply immediately the following

Corollary 4.2. *The random elements of the sequence Z_∞ are mutually uncorrelated.*

Let us now define a random measure F on the subsets of S by setting

$$F(\{0,1,...,k\}) = \begin{cases} 1 - \prod_{s=0}^{k} Z_\infty(s) & \text{if } k \in R, \\ 1 & \text{otherwise.} \end{cases} \tag{4.1}$$

Theorem 4.3. *If (3.3) holds, F is a random probability distribution function on S.*

Proof. We prove that $F(S) = 1$ with probability one. This is true by definition, when $R \subset S$. In order to prove that it is true as well when $R = S$, we show that

$$\lim_{s\to\infty} \prod_{i=1}^{s} Z_\infty(i) = 0 \tag{4.2}$$

almost surely. Compute

$$\begin{aligned} \mathbf{E}[\lim_{s\to\infty} \prod_{i=1}^{s} Z_\infty(i)] &= \lim_{s\to\infty} \mathbf{E}[\prod_{i=1}^{s} Z_\infty(i)] \\ &= \lim_{s\to\infty} \lim_{n\to\infty} \mathbf{E}[\prod_{i=1}^{s} Z_n(i)] \\ &= \lim_{s\to\infty} \prod_{i=1}^{s} Z_0(i) \\ &= \lim_{s\to\infty} \prod_{i=1}^{s} \frac{B_0(i)}{B_0(i) + W_0(i)} \\ &= 0; \end{aligned}$$

the first equality is true because of Monotone Convergence Theorem, the second and third equality hold because, for every $s \geq 1$, $\{\prod_{i=1}^{s} Z_n(i)\}$ is a bounded martingale converging almost surely to $\{\prod_{i=1}^{s} Z_\infty(i)\}$, the last equality is condition (3.3). Hence, $\lim_{s\to\infty} \prod_{i=1}^{s} Z_\infty(i)$ is a nonnegative random variable with mean 0; therefore (4.2) is true. ■

Theorem 4.4. *If condition (3.3) is true, the sequence* $\{T_n\}$ *is asymptotically exchangeable and its de Finetti measure is* F. *This means that, for* $j \geq 1$ *and* $t_1, ..., t_j \in S$,

$$\lim_{N\to\infty} P\left[T_{N+1} = t_1, \cdots, T_{N+j} = t_j\right] = \mathbf{E}[\prod_{i=1}^{j} F(\{t_i\})]. \tag{4.3}$$

Proof. For $t \notin R$,

$$P[T_n = t] = 0 = F(\{t\}) \tag{4.4}$$

for every $n \geq 1$ with probability one.

Let $j \geq 1$, $t_1, ..., t_j \in R$ and set

$$\zeta(s) = \sum_{i=1}^{j} I[s < t_i] \ \text{ and } \ \eta(s) = \sum_{i=1}^{j} I[s = t_i]$$

for $s \in S$. Now compute

$$\begin{aligned}
&\lim_{N\to\infty} P[T_{N+1} = t_1, ..., T_{N+j} = t_j] \\
&\quad = \lim_{N\to\infty} P[X_{N+1}(0) = 1, ..., X_{N+1}(t_1 - 1) = 1, X_{N+1}(t_1) = 0, \\
&\qquad\qquad X_{N+2}(0) = 1, ..., X_{N+2}(t_2 - 1) = 1, X_{N+2}(t_2) = 0, \\
&\qquad\qquad \cdots \\
&\qquad\qquad X_{N+j}(0) = 1, ..., X_{N+j}(t_j - 1) = 1, X_{N+j}(t_j) = 0] \\
&\quad = \mathbf{E}[\prod_{s\in S} Z_\infty^{\zeta(s)}(s)(1 - Z_\infty(s))^{\eta(s)}] \\
&\quad = \mathbf{E}[\prod_{i=1}^{j} Z_\infty(1) \cdots Z_\infty(t_i - 1)(1 - Z_\infty(t_i))] \\
&\quad = \mathbf{E}[\prod_{i=1}^{j} F(\{t_i\})].
\end{aligned}$$

The second equality holds because of Theorem 2.2: note that for only a finite number of $s \in S$ the factor $Z_\infty^{\zeta(s)}(s)(1 - Z_\infty(s))^{\eta(s)}$ is different from one. This and equation (4.4) prove that (4.3) is true for all $j \geq 1$ and $t_1, ..., t_j \in S$. ■

A notable particular case for the system X described in this section is obtained when

$$M_n(0) \equiv M_n(1) \equiv M_n(2) \equiv \cdots \equiv \widetilde{M}_n$$

for every $n = 1, 2, ...,$ where $\widetilde{M}_n$ is a random variable with values in $[0, K]$ and probability distribution $\widetilde{\mu}$. Hence, at each time n, reinforcements are randomly and independently generated, but they are constant along 0-vectors of the system X. An application of this model is when 0-vectors of X represent histories of a sequence of patients with survival times $T_1, T_2, ..., T_n, ...$, respectively; i.e. a ball of color 1 extracted form urn s at time n represents survival of patient n form day s to day $s+1$, while color 0 is for death of the patient at day s. To each patient is associated a nonnegative quantity, randomly generated and independent from the patient's survival time and from the survival time of other patients: this quantity expresses reinforcements for the urns connected with the patient's history. In this case, as in Example 2, it is not difficult to show that, for $s \in S$, $\{X_n(s)\}$ is a generalized Polya sequence with parameters $(B_0(s), W_0(s), \mu)$. However, as shown in [9], the distribution of $Z_\infty(s)$ needs not be a Beta; therefore the law of the random probability distribution F needs not be that of a beta-Stacy process, as it was the case in Example 3.

5. Beta-Stacy blanket, final remarks and open questions

Interacting reinforced urn systems are very flexible representations that accomodate known mathematical models (for applications in statistics, economics, biology) based on the idea of reinforcement and allow their extension to situations where dependence is an important issue: with this respect we consider Example 1 as prototypical. For a more complex example, reminiscent of the construction made in [11] for modelling a countable collection of dependent Polya sequences, we sketch here the description of an interacting system for modelling a countable collection of dependent reinforced urn processes.

Let S be an infinite k-ary tree, that is a connected graph with a distinguished vertex called the root and indicated with the symbol α, no cycles, a countable number of vertices such that every vertex has k children and 1 parent, except for the root that has no parent. Every vertex $s \in S$ labels an urn containing balls of color 1 and 0; initial urn compositions are described by the family $U_0 = ((B_0(s), W_0(s)), s \in S)$ with the assumption that $W_0(\alpha) = 0$. Let

$$M_1 \equiv M_2 \equiv M_3 \equiv \cdots \equiv 1$$

and, for every $s \in S$ and $x \in \{0,1\}^S$ define the reinforcement rule:

$$r(s, x, 1) = \prod_{i=0}^{n} x(s_i)$$

where $\pi(\alpha, s) = (s_0 = \alpha, s_1, ..., s_n, s)$ represents the unique path on the tree connecting the vertex s with the root α. These positions define an interacting reinforced urn system X with state space $\{0,1\}^S$. At time $n = 0, 1, 2...$ a ball is sampled from every urn in the system: the composition of every urn $s \in S$ is reinforced with an extra ball of the same color as that extracted from s if the balls simultaneously extracted from the ancestors of s, i.e. from urns different from s and labelled with vertices belonging to $\pi(\alpha, s)$, are all of color 1; otherwise the composition of urn s is left unchanged.

Let $\epsilon = (s_0 = \alpha, s_1, s_2, ...)$ be an end of the tree S, that is an infinite sequence of vertices of S with the property that s_i is the parent of s_{i+1} for $i = 0, 1, 2, ...$ The process $X(\epsilon) = (X_n(s_i), i = 0, 1, 2, ...)$ has the same law as the interacting reinforced urn system introduced in Example 3. In particular, for $n = 0, 1, 2, ...$, define

$$T_n(\epsilon) = \inf\{i : X_n(s_i) = 0\}.$$

If

$$\lim_{k\to\infty} \prod_{i=1}^{k} \frac{B_0(s_i)}{B_0(s_i) + W_0(s_i)} = 0, \tag{5.1}$$

the sequence $\{T_n(\epsilon)\}$ is exchangeable with de Finetti measure equal to the law of a beta-Stacy process $F(\epsilon)$ on the integers with parameters $((B_0(s_i), W_0(s_i)), i = 0, 1, ...)$.

Now consider two different ends of the tree S, $\epsilon = (s_0 = \alpha, s_1, s_2, ...)$ and $\eta = (t_0 = \alpha, t_1, t_2, ...)$, and assume that condition (5.1) is satisfied for both of them. The exchangeable sequences $\{T_n(\epsilon)\}$ and $\{T_n(\eta)\}$ are not independent, as well as their de Finetti measures, equal to the laws of the beta-Stacy processes $F(\epsilon)$ and $F(\eta)$ respectively. In fact, let $N \geq 0$ be the greatest integer such that $s_i = t_i$ for $i \leq N$: then $F(\epsilon)$ and $F(\eta)$ coincide on subsets of $\{0, ..., N\}$.

Let E be the space of ends of the tree S and assume that condition (5.1) is satisfied for every end $\epsilon \in E$. A future interesting problem will be to study the properties of the sequence of random configurations on the tree S formed by all vertices whose composition changes at time $n = 0, 1, 2, ...$, or, equivalently, the process $T = (T_n(\epsilon), \epsilon \in E)$ and its associated prior process $F = (F(\epsilon), \epsilon \in E)$ that we could call a *beta-Stacy blanket* with parameter family U_0, following [11]. Extensions allowing random reinforcements of the urns labelled with the vertices of the tree S are obvious and we omit to describe them, observing in passing that they generate a class of 'blanket' processes whose law will be completely characterized only when we are able to pinpoint the law of the random probability distribution F defined in 4.1. In particular it would be interesting to know under what conditions F is a neutral to the right process on the integers ([4]) and what class of neutral to the right processes on the integers can be generated by interacting reinforced urn systems as those described in Section 4. In point of fact, the characterization of the distribution of the limit Z_∞ of a generalized Polya sequence remains an open question, although good approximations for it are easily found by means of simulation.

References

1. ALDOUS, D. J. (1985). Exchangeability and related topics. *École d'été de probabilités de Saint-Flour, XIII—1983*, 1–198, Lecture Notes in Math. 1117, Springer, Berlin.
2. ATHREYA K.B. (1969). On a characteristic property of Pólya's urn *Studia Scientiarum Mathematicarum Hungarica*, **4**, 31-35.
3. DIACONIS, P., FREEDMAN, D. (1980). De Finetti's Theorem for Markov chains. *The Annals of Probability,* **8**, 115–130.
4. DOKSUM, K. A. (1974). Tailfree and neutral random probabilities and their posterior distributions. *The Annals of Probability*, **2**, 183–201.
5. EGGENBERGER, F. and POLYA, G. (1923). Über die Statistik verketteter Vorgänge. *Zeitschrift für Angewandte Mathemathik und Mechanik*, **3**, 279–289.
6. FRIEDMAN, B. (1949). A simple urn model. *Communications on Pure and Applied Mathematics*, **2**, 59–70.
7. HILL, B.M., LANE, D. and SUDDERTH, W. (1980). A strong law for some generalized urn processes. *The Annals of Probability*, 8, 214–226.
8. HILL, B.M., LANE, D. and SUDDERTH, W. (1987). Exchangeable urn processes. The Annals of Probability, 15, 1586–1592.
9. MAY, C., PAGANONI, A.M. and SECCHI, P. (2002). On a two-color generalized Pólya urn. *Quaderno di Dipartimento* n. 505/P, Aprile 2002, Dipartimento di Matematica, Politecnico di Milano.
10. MULIERE, P., SECCHI, P. and WALKER, S.G. (2000). Urn schemes and reinforced random walks *Stochastic Processes and their Applications*, **88**, 59–78.
11. MULIERE, P., SECCHI, P. and WALKER, S.G. (2003). Partially exchangeable processes indexed by the vertices of a k tree constructed via reinforcement. *Quaderno di Dipartimento* n. 543/P, Aprile 2003, Dipartimento di Matematica, Politecnico di Milano.
12. MULIERE, P., SECCHI, P. and WALKER, S.G. (2003). Reinforced random processes in continuous time. *Stochastic Processes and their Applications*, 88, 59–78.
13. NEVEU, J. (1972). *Martingales à temps discret*, Masson, Paris.
14. PEMANTLE, R. (1990). A time-dependent version of Pólya's urn. *Jour. Theor. Prob.*, 3, 627 - 637.
15. WALKER, S. and MULIERE, P. (1997). Beta-Stacy processes and a generalization of the Polya urn scheme. *Annals of Statistics*, **25**, 1762–1780.
16. WALKER, S. G. and MULIERE, P. (2002). A bivariate Dirichlet process. *Statistics and Probability Letter*, to appear.

E-mail address: annap@mate.polimi.it

E-mail address: secchi@mate.polimi.it

Dipartimento di Matematica "F. Brioschi", Politecnico di Milano, Piazza Leonardo da Vinci, 32, 20123 Milano - Italy