Monografías del Seminario Matemático García de Galdeano **33**, 309–316 (2006)

ASYMPTOTICS FOR THE NUMBER OF REPLACEMENTS IN A GENERALIZED PÓLYA URN MODEL

José A. Moler, Fernando Plo, Miguel San Miguel
and Henar Urmeneta

Abstract. There are three processes associated to a generalized Pólya urn model. First, the process that represents the proportion of balls of each type in the urn. As in each step a ball is drawn from the urn, its type is noted, and it is placed back in the urn, the second process represents the proportion of balls of each type that have been drawn from the urn. As the replacement policy consists in applying in each step one out of K different actions, the third process represents the proportion of times that each action (replacement) has been applied. This third process has not attracted as much attention as the others in the probabilistic literature. In this work we present conditions under which almost sure convergence results and central limit theorems are obtained for it. We illustrate these results with an application to adaptive clinical trials and random data structures.

Keywords: Generalized Pólya urn models, central limit theorems, Robbins-Monro recurrence equation.

AMS classification: 60F05, 62L20.

§1. Introduction

Generalized Pólya urn models constitute a powerful too to study the evolution of a wide family of random data structures and to perform adaptive clinical trials (see, for instance, [5] and [6]). In the following lines, we give a general framework to introduce generalized Pólya urn models. We consider an urn with a total amount of T_0 balls of L different types. The initial proportion of each kind of ball is represented by the vector $X_0 = (X_{01}, \dots, X_{0L})^t$. At each step n, one out of K different actions can be applied. If action i is applied, r_{nij} balls of type j, $j = 1, \dots, L$, are added to the urn (or extracted, if $r_{nij} < 0$). We assume that $\sum_{j=1}^{L} r_{nij} > 0$. All these values are collected in a matrix $R_n = (r_{nij})$, $i = 1, \dots, K$, $j = 1, \dots, L$, which is called replacement matrix. We denote by $\delta_n = (\delta_{n1}, \dots, \delta_{nK})^t$ a K-dimensional vector of indicator variables such that, if the i-th action is applied, then $\delta_{ni} = 1$ and the rest of components are equal to 0. The total number of balls in the urn after the n-th replacement is denoted by T_n. The process $\{X_n\}$ represents the proportion of balls of each type in the urn, after each replacement n and it takes values in the set $\Delta_{L-1} \subset \mathbb{R}^L$, where $\Delta_{L-1} = \{x : x_1 + \cdots + x_L = 1,\ x_i \geq 0\}$. Throughout this paper the following assumptions are considered:

(A1) For each n, R_n is a deterministic matrix R such that $R\mathbf{1} = s\mathbf{1}$, where $\mathbf{1}$ represents the column vector of ones and $s > 0$.

(A2) There exists a continuous function $G : \Delta_{L-1} \to \Delta_{K-1}$ with components G_i, $i = 1, \ldots, K$, such that

$$P(\delta_{n+1,i} = 1 | \mathscr{F}_n) = G_i(X_n), \quad i = 1, \ldots, K. \tag{1}$$

(A3) The process can not get stuck for impossible removals, so that

$$R^t \delta_n + T_{n-1} X_{n-1} \geq \mathbf{0}.$$

The process $\{U_n\} = \{(X_n, \delta_n)\}$, $n \geq 1$, is a generalized Pólya urn model. We will consider the natural filtration $\{\mathscr{F}_n\}$ where $\mathscr{F}_n = \sigma(U_i : i \leq n)$, $n \geq 1$. The asymptotic behaviour of the process $\{X_n\}$ has been studied in many papers. See, for instance, [1] and [2], where this study is made in the framework of stochastic approximation, and the references therein. In this paper we focus on the establishment of conditions in a generalized Pólya urn model to obtain almost sure convergence results and central limit theorems for the process $\{Z_n\}$ where, for each n, $Z_n = \sum_{k=1}^n \delta_k / n$. This process represents the proportion of times that each action (replacement) has been applied up to the n-th stage. As in each step a ball is drawn from the urn, its type is noted, and it is placed back in the urn, we could be interested in the process, say $\{W_n\}$, that represents the proportion of times that each type of ball has been drawn. We define an L-dimensional vector of indicator variables $\eta_n = (\eta_{n1}, \ldots, \eta_{nL})^t$ such that if the i-th type of ball has been drawn in the n-th step, then $\eta_{ni} = 1$ and the rest of components are equal to 0. We have then that $W_n = \sum_{k=1}^n \eta_k / n$ for each n. Observe that when $G(x) = x$, then $K = L$ and $W_n = Z_n$, for each n. However, if $G(x) \neq x$, the action applied could not coincide with the type of ball drawn from the urn. See, for instance, Application 1 in Section 3, where the processes $\{W_n\}$ and $\{Z_n\}$ provide different information. The paper is organized as follows. In the following section, we give theoretical results that establish conditions on the generalized Pólya urn model to obtain almost sure convergence for the processes $\{Z_n\}$ and $\{W_n\}$ and central limit theorems for the process $\{Z_n\}$. In the final section we illustrate the practical interest of these processes in the framework of clinical trials and in the modelling of random data structures.

§2. Theoretical results

The general procedure to obtain our results is to prove that the process of interest follows a Robbins-Monro scheme and to check that the conditions that guarantee a strong law or a central limit theorem for Robbins-Monro processes hold. We use the following notation. If R is a matrix, then $\|R\| = \sup_i \sum_j |r_{ij}|$. If $x \in \mathbb{R}^n$, then $\|x\| = \sum_{i=1}^n |x_i|$ and the row vector will be noted by x^t. Column vectors of zeros will be noted by $\mathbf{0}$.

Definition 1. Consider a stochastic process $\{Y_n\}$ that takes values in a borelian set $\Gamma \subseteq \mathbb{R}^d$ and that is defined on a probability space $(\Omega, \mathscr{A}, P)$ endowed with a filtration $\mathscr{G} = (\mathscr{G}_n)$. $\{Y_n\}$ follows a Robbins-Monro scheme of stochastic approximation when

$$Y_{n+1} = Y_n + \gamma_n (F(Y_n) + \varepsilon_{n+1} + \boldsymbol{\beta}_{n+1}) \tag{2}$$

In order to obtain asymptotic results for $\{Y_n\}$, several conditions must be assumed. We write down the usual ones:

(C1) $\{\gamma_n\}$ is a sequence of positive random variables such that a.s.:

$$\sum_{n=1}^{\infty}\gamma_n=\infty, \qquad \sum_{n=1}^{\infty}\gamma_n^2<\infty$$

(C2) The sequence $\{\varepsilon_n\}$ is adapted to $\mathscr{G}$ and for a positive constant M and $b>2$:

$$E[\varepsilon_{n+1}\mid\mathscr{G}_n]=0, \qquad E[\|\varepsilon_{n+1}\|^b\mid\mathscr{G}_n]<M, \; a.s.$$

(C3) The sequence $\{\boldsymbol{\beta}_n\}$ converges to 0 a.s.

Consider a generalized Pólya urn model $\{U_n\}$ under hypotheses (A1)–(A3). The process $\{X_n\}$ satisfies the following recurrence equation:

$$X_{n+1}=\frac{T_nX_n+R^t\delta_{n+1}}{T_{n+1}}=X_n+\frac{(I-X_n\mathbf{1}^t)R^t\delta_{n+1}}{T_{n+1}}, \tag{3}$$

where T_n represents the total number of balls in the urn after the n-th replacement and $T_n=T_0+ns$. From (3) we have:

$$\begin{aligned}X_{n+1}&=E[X_{n+1}\mid\mathscr{F}_n]+X_{n+1}-E[X_{n+1}\mid\mathscr{F}_n]\\&=X_n+\frac{1}{T_{n+1}}[(I-X_n\mathbf{1}^t)R^tG(X_n)+(I-X_n\mathbf{1}^t)R^t(\delta_{n+1}-G(X_n))]\\&=X_n+\frac{1}{T_{n+1}}[R^tG(X_n)-sX_n+R^t(\delta_{n+1}-G(X_n))].\end{aligned} \tag{4}$$

Then, $\{X_n\}$ follows a Robbins-Monro scheme with $\Gamma=\Delta_{L-1}$ and, for each n, $\gamma_n=1/(T_0+ns)$, $F(X_n)=R^tG(X_n)-sX_n$, $\varepsilon_{n+1}=R^t(\delta_{n+1}-G(X_n))$ and $\beta_{n+1}=0$. It is easy to check that conditions (C1)–(C3) hold. Now, we focus on the processes $\{Z_n\}$ and $\{W_n\}$. In the following result asymptotic results on the process $\{X_n\}$ are assumed.

Theorem 1. *Consider a generalized Pólya urn model $\{U_n\}$ under assumptions (A1)–(A3) and assume that $X_n\to u$, a.s., where $u\in\Delta_{L-1}$ and $F(u)=0$. Then,*

$$\begin{aligned}Z_n&\to G(u), \quad a.s.,\\W_n&\to u, \quad a.s.\end{aligned}$$

Besides, if $\sqrt{ns}(X_n-u)\to N(\mathbf{0},V)$, $[D]$, then

$$\sqrt{ns}\,\frac{R^t}{s}(Z_n-G(u))\to N(0,V), \quad [D].$$

Proof. For each n, $E[\eta_n\mid\mathscr{F}_{n-1}]=X_{n-1}$, a.s., and $E[\delta_n\mid\mathscr{F}_{n-1}]=G(X_{n-1})$, a.s. From the Lévy's extension of the Borel-Cantelli's Lemma, we have that $\{W_n\}$ converges a.s. to the same limit as $\sum_{k=1}^n X_{k-1}/n$, and that $\{Z_n\}$ converges a.s. to the same limit as $\sum_{k=1}^n G(X_{k-1})/n$. Then, $W_n\to u$, a.s. and, being G continuous, $Z_n\to G(u)$, a.s. In order to prove the central limit theorem, observe that $T_nX_n=T_0X_0+\sum_{k=1}^n R^t\delta_k$, and then

$$X_n=\frac{T_0X_0}{T_n}+\frac{n}{T_n}R^tZ_n. \tag{5}$$

As $F(u) = 0$, then $R^t G(u) = su$, so that

$$\sqrt{ns}(X_n - u) = \sqrt{ns}\,\frac{T_0 X_0}{T_n} + \sqrt{ns}\,\frac{R^t}{s}\Big(\frac{nsZ_n}{T_n} - Z_n\Big) + \sqrt{ns}\,\frac{R^t}{s}(Z_n - G(u)). \tag{6}$$

The first and the second addend in (6) converge a.s. to zero, and therefore the result follows. □

Remark 1. If in Theorem 1 R is a non-singular L-square matrix, then

$$\sqrt{n}(Z_n - G(u)) \to N(0, s(R^t)^{-1} V R^{-1}), \quad [D].$$

In the following result we study directly the asymptotic behaviour of $\{Z_n\}$, without assumptions on the asymptotic behaviour of $\{X_n\}$.

Theorem 2. *Consider a generalized Pólya urn model $\{U_n\}$ under assumptions (A1)–(A3). Assume that $G(X_n) = X_n$, for each n. Then, the process $\{Z_n\}$ follows a Robbins-Monro scheme and conditions (C1)–(C3) hold. Moreover, if s is a simple eigenvalue of R and the rest of eigenvalues have real part strictly lesser than s, then*

$$Z_n \to u, \quad a.s.,$$

where u is the unique vector in Δ_{L-1} such that $u^t R = su^t$. Besides, if s is a simple eigenvalue of R and the rest of eigenvalues have real part strictly lesser than $s/2$, then

$$\sqrt{n}(Z_n - u) \to N(0, V), \quad [D],$$

where V is a singular matrix such that

$$\Big(\frac{R^t}{s} - \frac{1}{2}I\Big)V + V\Big(\frac{R^t}{s} - \frac{1}{2}I\Big)^t = -C,$$

and

$$C = \begin{pmatrix} u_1(1-u_1) & -u_1u_2 & \dots & -u_1u_L \\ -u_1u_2 & u_2(1-u_2) & \dots & -u_2u_L \\ \vdots & \vdots & \ddots & \vdots \\ -u_1u_L & -u_2u_L & \dots & u_L(1-u_L) \end{pmatrix}. \tag{7}$$

Proof. Since $Z_{n+1} = Z_n + (\delta_{n+1} - Z_n)/(n+1)$, taking $\varepsilon_{n+1} = \delta_{n+1} - X_n$, we obtain

$$\begin{aligned} Z_{n+1} &= Z_n + \frac{1}{n+1}[(X_n - Z_n) + \varepsilon_{n+1}] \\ &= Z_n + \frac{1}{n+1}[(\frac{R^t}{s} - I)Z_n + \varepsilon_{n+1} + \beta_{n+1}], \end{aligned} \tag{8}$$

where

$$\beta_{n+1} = \frac{T_0 X_0}{T_n} + \frac{R^t}{s} Z_n \Big(\frac{sn}{T_n} - 1\Big) = o(n^{-\rho}), \quad a.s. \text{ for any } \rho < 1. \tag{9}$$

Therefore, $\{Z_n\}$ fits a Robbins-Monro scheme that satisfies conditions (C1)–(C3) with $F(x) = (R^t/s - I)x$. Let P be the transformation matrix of the real canonical Jordan form of $(R^t/s - I)$. We can take $\mathbf{1}^t$ as the first row of P because it is a left eigenvector associated to the eigenvalue 1 of R^t/s. Consider $P := \begin{pmatrix} \mathbf{1}^t \\ \widetilde{P} \end{pmatrix}$. Then, as $\mathbf{1}^t Z_n = 1$, $\mathbf{1}^t \beta_n = 0$ and $\mathbf{1}^t \varepsilon_n = 0$, we have from (8) that the process $\{\widetilde{Z}_n\} := \{\widetilde{P} Z_n\}$ satisfies

$$\widetilde{Z}_{n+1} = \widetilde{Z}_n + \frac{1}{n+1}[\widetilde{J}\widetilde{Z}_n + \widetilde{\varepsilon}_{n+1} + \widetilde{\beta}_{n+1}], \tag{10}$$

where

$$P\Big(\frac{R^t}{s} - I\Big)P^{-1} = \begin{pmatrix} 1 & \mathbf{0}^t \\ \mathbf{0} & \widetilde{J} \end{pmatrix}, \quad P\varepsilon_n = \begin{pmatrix} 0 \\ \widetilde{\varepsilon}_n \end{pmatrix}, \quad P\beta_n = \begin{pmatrix} 0 \\ \widetilde{\beta}_n \end{pmatrix}.$$

Observe that the ODE associated to the scheme (10) is $\dot{x} = \widetilde{J}x$, and, from the assumptions, all the eigenvalues of $\widetilde{J}$ are negative. So that, 0 is a globally asymptotically stable point for the ODE. We invoke Theorem 5.2.3 in [3] to conclude that $\{\widetilde{Z}_n\}$ converges a.s. to $\widetilde{P}u = \mathbf{0}$. And, then, $Z_n \to u$, a.s. On the other hand, observe that from (5) we deduce that $X_n \to u$, a.s. So that, if $C_n := E[\varepsilon_n \varepsilon_n^t \mid \mathscr{F}_{n-1}]$, then it is easy to check that $C_n \to C$, a.s., where C is as in (7). As all the extra-diagonal components of $-C$ are positive, it follows from Theorem 2.6 in [8] that 0 is a simple eigenvalue of C and the rest of eigenvalues have real part strictly positive. Then, let

$$PCP^t = \begin{pmatrix} 0 & \mathbf{0}^t \\ \mathbf{0} & \widetilde{C} \end{pmatrix},$$

it follows that $\widetilde{C}$ is positive definite. Moreover, we have that $\widetilde{J} + \frac{1}{2}I$ is a stable matrix. From the previous discussion it follows that the conditions in Theorem 1 in [7] hold for the recurrence scheme (10), and then

$$\sqrt{n}(\widetilde{Z}_n) \to N(\mathbf{0}, \Sigma),$$

where Σ is the unique solution of the Lyapunov equation:

$$(\widetilde{J} + \tfrac{1}{2}I)\Sigma + \Sigma(\widetilde{J}^t + \tfrac{1}{2}I) = -\widetilde{C} \tag{11}$$

From the relations between Z_n and $\widetilde{Z}_n$, the result follows with

$$V = P^{-1}\begin{pmatrix} 0 & \mathbf{0}^t \\ \mathbf{0} & \Sigma \end{pmatrix} P^{-1t}. \quad \square$$

§3. Applications

3.1. Application 1: A clinical trial

The randomized Play-The-Winner rule, introduced in [9], and its modifications have been popular adaptive designs used in clinical trials (see, for instance, [6]). This rule is implemented by using a randomized urn model that contains balls of two types (say, type 1 and type 2). Each type is associated with a treatment. When a patient arrives, a ball is drawn. The patient receives the treatment associated with the ball type, i, and the ball is replaced in the urn. It is assumed that the patient gives an immediate and dichotomous response. If the treatment is successful, β balls of type i and α balls of the other type are added into the urn. Otherwise, α balls of type i and β balls of the other type are added into the urn. It is assumed that β and α are non-negative integer numbers such that $\beta > \alpha \geq 0$, so that

$$R = \begin{pmatrix} \beta & \alpha \\ \alpha & \beta \end{pmatrix}.$$

Let p_i be the probability of treatment i being a success, $i = 1, 2$, and $q_i = 1 - p_i$. Let $\{X_n\} = \{(X_{n1}, X_{n2})\}$ be the stochastic process that represents the proportions of balls of each type in the urn at the n-th stage of the experiment. It is well-known (see [9]) that

$$X_{1n} \to u_1 = \frac{\alpha p_2 + \beta q_2}{\alpha(p_1 + p_2) + \beta(q_1 + q_2)}, \quad a.s. \tag{12}$$

In this example, the processes $\{Z_n\}$ and $\{W_n\}$ are not the same. Note that treatment 1 can be rewarded, (that is, action 1 can be applied) even if treatment 2 has been applied. The process $\{Z_n\}$ represents the proportion of times that we have applied each action (replacement) and the process $\{W_n\}$ represents the proportion of times that we have applied each treatment. As $\delta_{ni} = 1$ if treatment i has been applied, we can write

$$P(\delta_{n1} = 1 \mid \mathscr{F}_{n-1}) = p_1 X_{n1} + q_2 X_{n2}, \qquad P(\delta_{n2} = 1 \mid \mathscr{F}_{n-1}) = q_1 X_{n1} + p_2 X_{n2}.$$

And therefore

$$G(x) = \begin{pmatrix} p_1 & q_2 \\ q_1 & p_2 \end{pmatrix} x := Bx, \qquad x^t = (x_1, x_2).$$

As (A1)–(A3) hold, the process $\{X_n\}$ follows a Robbins-Monro scheme and conditions (C1)–(C3) hold. Note that $F(x) = Ax$, where

$$A = (R^t B - (\alpha + \beta) I) = \begin{pmatrix} -\beta q_1 - \alpha p_1 & \beta q_2 + \alpha p_2 \\ \beta q_1 + \alpha p_1 & -\beta q_2 - \alpha p_2 \end{pmatrix}. \tag{13}$$

From (12) we know that $X_n \to u$, a.s., and it is easy to check that $Au = \mathbf{0}$. Then, from Theorem 1 it follows that

$$Z_n \to Bu = \begin{pmatrix} q_2 + (p_1 - q_2) u_1 \\ q_1 + (p_2 - q_1) u_2 \end{pmatrix}.$$

From Corollary 3.1 in [2], it can be proved that, if $2\beta + (\alpha - \beta)(p_1 + p_2) > 1/2$, then $\sqrt{ns}(X_n - u) \to N(\mathbf{0}, V)$. Therefore, applying Theorem 1 and Remak 1 we have that

$$\sqrt{n}(Z_n - G(u)) \to N(0, (\alpha + \beta)(R^t)^{-1} V R^{-1}), \quad [D].$$

3.2. Application 2: Rotations in fringe-balanced binary trees

Rotation is a heuristic that reduces height and path in order to improve the speed of retrieval in a binary search tree. In [4], a generalized Pólya urn model is proposed to model rotations.

Three types of nodes are distinguished. When a node of type 3 is drawn, a rotation is made. The replacement matrix is

$$R = \begin{pmatrix} -2 & 1 & 2 \\ 4 & -1 & -2 \\ 4 & -1 & -2 \end{pmatrix}.$$

In [4] the asymptotic behaviour of the proportion of times that a node of type 3 has been extracted is also obtained using martingale techniques. Studying this model in the setting of this paper, we note that $G(x) = x$, so that $Z_n = W_n$ for each n, and we can apply Theorem 2. The a.s. limit for the process $\{Z_n\}$ is $u = (4/7, 1/7, 2/7)$, and the central limit theorem is:

$$\sqrt{n}(Z_n - \mathbf{u}) \to N\left(\mathbf{0}, \begin{pmatrix} 12/637 & -4/637 & -8/637 \\ -4/637 & 62/637 & -58/637 \\ -8/637 & -58/637 & 66/637 \end{pmatrix}\right).$$

The (singular) covariance matrix has been obtained following the procedure stated in the proof of Theorem 2. First, as $s = 1$, we compute P and J such that $P^{-1}JP = R^t - I$ (where J is the real canonical Jordan form of $R^t - I$). We obtain:

$$P = \begin{pmatrix} 1 & 1 & 1 \\ 881/1079 & -749/688 & -749/688 \\ 0 & -1047/727 & 1402/1947 \end{pmatrix} \quad \text{and} \quad J = \begin{pmatrix} 0 & 0 & 0 \\ 0 & -7 & 0 \\ 0 & 0 & -1 \end{pmatrix}.$$

We compute the matrices C and $\widetilde{C}$:

$$C = \frac{1}{49}\begin{pmatrix} 12 & -4 & -8 \\ -4 & 6 & -2 \\ -8 & -2 & 10 \end{pmatrix} \quad \text{and, then,} \quad \widetilde{C} = \frac{1}{9}\begin{pmatrix} 8 & 0 \\ 0 & 4 \end{pmatrix}.$$

As the matrix $\widetilde{J} + \frac{1}{2}I$ is stable, we obtain the unique solution Σ of (11):

$$\Sigma = \int_0^\infty \exp\big((\widetilde{J} + 1/2)x\big)\widetilde{C}\exp\big((\widetilde{J} + 1/2)^t x\big)\,dx = \begin{pmatrix} 8/117 & 0 \\ 0 & 4/9 \end{pmatrix}.$$

And finally, we obtain the matrix V stated above.

Acknowledgements

The authors thank support from the research group Modelos Estocásticos (DGA) and from the MTM 2004-01175 project of MEC.

References

[1] HIGUERAS, I., MOLER, J., PLO, F., AND SAN MIGUEL, M. Urn models and differential algebraic equations. *Journal of Applied Probability 40* (2003), 401–412.

[2] HIGUERAS, I., MOLER, J., PLO, F., AND SAN MIGUEL, M. Central limit theorems in generalized Pólya urn models. Submitted.

[3] KUSHNER, H. J., AND YIN, G. G. *Stochastic Approximation Algorithms and Applications*. Springer-Verlag, New York, 1997.

[4] MAHMOUD, H. M. On rotations in fringe-balanced binary trees. *Information Processing Letters 65* (1998), 41–46.

[5] MAHMOUD, H. M. The size of random bucket trees via urn models. *Acta Informatica 38* (2002), 813–838.

[6] MOLER, J., PLO, F., AND SAN MIGUEL, M. Adaptive designs and Robbins-Monro algorithm. *Journal of Statistical Planning and Inference 131* (2005), 161–174.

[7] PELLETIER, M. Weak convergence rates for stochastic approximation with application to multiple targets and simulated annealing. *The Annals of Applied Probability 8* (1998), 10–44.

[8] SENETA, E. *Non-negative Matrices and Markov Chains*. Springer-Verlag, New York, 1981.

[9] WEI, L. J., AND DURHAM, S. The randomized play-the-winner rule in medical trials. *Journal of the American Statistical Association 73* (1978), 840–843.

José A. Moler and Henar Urmeneta
Departamento de Estadística e Investigación Operativa
Universidad Pública de Navarra
Campus Arrosadía s/n, 31006-Pamplona, Spain.
`jmoler@unavarra.es` and `henar@unavarra.es`

Fernando Plo and Miguel San Miguel
Departamento de Métodos Estadísticos
Universidad de Zaragoza
Pedro Cerbuna, 12 50009-Zaragoza, Spain
`fplo@unizar.es` and `msm@unizar.es`