Polya's urn and martingales

Daniel Ahlberg, Stockholm University *May 19, 2021*

Consider an urn which at the beginning of time contains one red and one blue ball. In each time step we draw a ball from the urn uniformly at random, and then puts the ball back into the urn together with another ball of the same colour. That is, after k rounds there are $2+k$ balls in the urn, and in round $k+1$ each of these $2+k$ balls is equally likely to be drawn, independently of the outcome of previous rounds, and replaced together with an identical copy. This process is commonly referred to as *Polya's urn*, after the Hungarian mathematician George Pólya.

An interesting question to answer is what happens to the composition of balls in Polya's urn over time. That is, if Y_n denotes the proportion of red balls after n rounds, then we would like to understand how the sequence $(Y_n)_{n\geq 0}$ evolves over time. Does $(Y_n)_{n\geq 0}$ converge, and in what sense? If a limit exists, what can we say about the limit? In order to address these questions we introduce the variables $X_1, X_2, \ldots$ where

$$X_k := \begin{cases} 1 & \text{red ball drawn in round } k, \\ 0 & \text{blue ball drawn in round } k. \end{cases}$$

Since there are $2+n$ balls in the urn after n rounds, we have that the proportion of red balls after n rounds can be expressed as

$$Y_n = \frac{1 + \sum_{k=1}^{n} X_k}{2+n}. \tag{1}$$

From (1) we see that Y_n roughly is the partial sum of the first n variables of the sequence $(X_k)_{k\geq 1}$. Note, however, that the variables in the sequence are not independent, and that the distribution of X_k depends on the outcome of previous rounds. More precisely, given that $Y_k = p$, the variable X_{k+1} is Bernoulli distributed with parameter p. This means that we are outside of the domain where the law of large numbers applies. On the other hand, without reinforcement (that is, if balls are drawn and replaced, without addition of further balls) the variables in the sequence $(X_k)_{k\geq 1}$ would be independent Bernoulli distributed with parameter $1/2$, and the law of large numbers would imply that Y_n converges in probability (and almost surely) to $1/2$ as $n \to \infty$.

Convergence in distribution

If we would allow for reinforcement in the first round, but not in later rounds, then for all future rounds the urn would consist of either one red and two blue, or two red and one blue, and both of the two will occur with equal probability. Given the outcome of the first round, the outcome of the remaining rounds would be i.i.d. Bernoulli distributed with parameter either $1/3$ or $2/3$ (depending on whether the first draw resulted in blue or red). So, by the law of large numbers, Y_n would again converge (in probability and almost surely) to a random variable taking the values 1/3 and 2/3 with equal probability. Allowing for reinforcement in more rounds thus seem to 'spread out' the possible values of Y_n. This is confirmed by the following result.

Theorem 1. *For Polya's urn, the proportion of red balls Y_n converges in distribution, as $n \to \infty$, to the uniform distribution on the interval* $[0, 1]$.

Proof. Let $F_n(x) := \mathbb{P}(Y_n \leq x)$ be the cumulative distribution function of Y_n. We need to show that $F_n(x) \to x$ as $n \to \infty$ for all $x \in [0, 1]$. Since F_n is increasing and takes values in $[0, 1]$, this will imply that $F_n(x)$ is zero for $x < 0$ and one for $x > 1$.

Let A denote the event that among the first n rounds, the first k result in a red ball being drawn and the remaining $n - k$ result in a blue ball being drawn. Since after j rounds there are $2 + j$ balls in the urn, we have

$$\mathbb{P}(A) = \frac{1}{2} \cdot \frac{2}{3} \cdots \frac{k}{k+1} \cdot \frac{1}{k+2} \cdot \frac{2}{k+3} \cdots \frac{n-k}{n+1} = \frac{k!(n-k)!}{(n+1)!}. \tag{2}$$

Note further that for *any* outcome of the first n rounds that result in precisely k red balls drawn has the same probability, since it would only result in a permutation of the numbers in the nominator in (2). Consequently, if R_n denotes the number of red balls in the urn after n rounds, then

$$\mathbb{P}\big(R_n = k + 1\big) = \binom{n}{k}\mathbb{P}(A) = \frac{n!}{k!(n-k)!} \cdot \frac{k!(n-k)!}{(n+1)!} = \frac{1}{n+1} \tag{3}$$

for $k = 0, 1, \dots, n$. That is, any of the $n + 1$ possible compositions of balls in the urn after n rounds is equally likely.

Finally, we rephrase the observation from (3) in terms of F_n. Note that $F_n(1) = 1$ is immediate. For $x \in [0, 1)$, we have from (3) that as $n \to \infty$

$$F_n(x) = \mathbb{P}\big(R_n \leq x(2 + n)\big) = \frac{\lfloor x(2 + n) \rfloor}{n + 1} \to x,$$

where $\lfloor y \rfloor$ denotes the integer part of a real number y. □

The above theorem shows that the sequence $(Y_n)_{n \geq 0}$ converges in distribution to a random variable U which is uniformly distributed on the interval $[0, 1]$. Is it possible that the convergence holds in probability or even almost surely? Let again $R_n = 1 + \sum_{k=1}^{n} X_k$ denote the number of red balls in the urn after n rounds, and note that

$$Y_{n+1} - Y_n = \frac{R_{n+1}}{n+3} - \frac{R_n}{n+2} = \frac{(n+2)(R_{n+1} - R_n) - R_n}{(n+2)(n+3)} = \frac{X_{n+1} - Y_n}{n+3}. \tag{4}$$

That is, $|Y_{n+1} - Y_n|$ is roughy of the order $1/n$, which is non-summable. So that it is certainly *possible* that the sequence $(Y_n)_{n \geq 0}$ converges in distribution, but continues to wander around without stablilizing. Note also that the expected change $\mathbb{E}[Y_{n+1} - Y_n]$ is zero, much like when adding i.i.d. random variables, so it is perhaps *plausible* that the sequence converges also in a stronger sense (in probability or almost surely). However, we do not even seem to have a reasonable candidate for what the limiting random variable should be, apart from having a uniform distribution. We shall therefore aim to prove that the limit exists, without having to determine what it is.

The martingale property

We saw in (4) that the change in each step of the sequence $(Y_n)_{n\geq 0}$ is of the order $1/n$, but that the expected change $\mathbb{E}[Y_{n+1}-Y_n]$ is zero. In fact, conditional on the outcome of previous rounds, the conditional expected change is also zero. Since the conditional distribution of X_{n+1}, given $Y_n = p$, is Bernoulli with parameter p, we obtain via (4) that

$$\begin{aligned}\mathbb{E}[Y_{n+1}-Y_n|Y_0,Y_1,\dots,Y_n] &= \frac{1}{n+3}\,\mathbb{E}[X_{n+1}-Y_n|Y_0,Y_1,\dots,Y_n]\\ &= \frac{1}{n+3}\big(\,\mathbb{E}[X_{n+1}|Y_n]-Y_n\big)\\ &= 0.\end{aligned}$$

The above property is known as *the martingale property*, and a sequence $(M_n)_{n\geq 0}$ of random variables with this property, i.e. which satisfies

$$\mathbb{E}[M_{n+1}-M_n|M_0,M_1,\dots,M_n] = 0 \quad \text{for all } n\geq 0, \tag{5}$$

is known as a *martingale*.

A martingale will not (necessarily) have independent increments, but the martingale property assures that the increments are uncorrelated. This property assures that martingales in many ways enjoy properties similar to sums of i.i.d. random variables.

Lemma 2. *The increments of a martingale $(M_n)_{n\geq 0}$ are uncorrelated, meaning that for all $i\neq j$ we have*

$$\mathrm{Cov}\big(M_{i+1}-M_i, M_{j+1}-M_j\big) = 0.$$

Proof. By symmetry, we may assume that $i<j$. We first note that (since $i+1\leq j$) the martingale property (5) implies that

$$\mathbb{E}\big[M_{j+1}-M_j\big|M_0,\dots,M_{i+1}\big] = \mathbb{E}\big[\,\mathbb{E}\big[M_{j+1}-M_j\big|M_0,\dots,M_j\big]\big|M_0,\dots,M_{i+1}\big] = 0.$$

Consequently, it follows that

$$\begin{aligned}\mathrm{Cov}\big(M_{i+1}-M_i, M_{j+1}-M_j\big) &= \mathbb{E}\big[(M_{i+1}-M_i)(M_{j+1}-M_j)\big]\\ &= \mathbb{E}\big[\,\mathbb{E}\big[(M_{i+1}-M_i)(M_{j+1}-M_j)\big|M_0,\dots,M_{i+1}\big]\big]\\ &= \mathbb{E}\big[(M_{i+1}-M_i)\,\mathbb{E}\big[M_{j+1}-M_j\big|M_0,\dots,M_{i+1}\big]\big]\\ &= 0,\end{aligned}$$

as required. □

From Lemma 2 we obtain an expression for the variance of a martingale:

$$\mathrm{Var}(M_n-M_0) = \mathrm{Var}\left(\sum_{k=0}^{n-1} M_{k+1}-M_k\right) = \sum_{k=0}^{n-1}\mathrm{Var}(M_{k+1}-M_k). \tag{6}$$

For the sequence $(Y_n)_{n\geq 0}$ we obtain, via the conditional variance formula, the martingale property (5) and (4), that

$$\begin{aligned}\mathrm{Var}(Y_{k+1}-Y_k) &= \mathrm{Var}\left(\mathbb{E}[Y_{k+1}-Y_k|Y_k]\right)+\mathbb{E}\left[\mathrm{Var}(Y_{k+1}-Y_k|Y_k)\right]\\ &= \frac{1}{(n+3)^2}\,\mathbb{E}\left[\mathrm{Var}(X_{k+1}-Y_k|Y_k)\right]\\ &= \frac{1}{(n+3)^2}\,\mathbb{E}[Y_k(1-Y_k)],\end{aligned}$$

which is no larger than $1/(n+3)^2$. This is a summable sequence, so (6) implies that $\mathrm{Var}(Y_n)=\mathrm{Var}(Y_n-Y_0)$ is bounded. Moreover, for every $N\geq 1$ and all $n\geq 1$ we have

$$\mathrm{Var}(Y_{N+n}-Y_N)=\sum_{k=N}^{N+n}\mathrm{Var}(Y_{k+1}-Y_k)\leq\int_{N-1}^{\infty}\frac{1}{(x+3)^2}\,dx\leq\frac{1}{N}. \tag{7}$$

So, the 'amount of randomness' that remains after N steps in not more than $1/N$. Using Chabyshev's inequality we find that

$$\mathbb{P}\big(|Y_{N+n}-Y_N|>\varepsilon\big)\leq\frac{1}{\varepsilon^2}\,\mathrm{Var}(Y_{N+n}-Y_N)\leq\frac{1}{\varepsilon^2 N},$$

so it is unlikely that at any *fixed* number of further steps the sequence deviates far from its state after N rounds. The following maximal inequality is a significant strengthening of this observation, showing that it is unlikely that the sequence *ever again* deviates far from its current state.

Lemma 3 (Doob-Kolmogorov maximal inequality)**.** *Let* $(M_n)_{n\geq 0}$ *be a martingale with* $\mathrm{Var}(M_{n+1}-M_n)<\infty$ *for all* $n\geq 0$*. Then, for every* $\varepsilon>0$*,*

$$\mathbb{P}\Big(\max_{1\leq k\leq n}|M_k-M_0|>\varepsilon\Big)\leq\frac{1}{\varepsilon^2}\,\mathrm{Var}(M_n-M_0).$$

Proof. Set $A:=\{\max_{1\leq k\leq n}|M_k-M_0|>\varepsilon\}$, and let A_k be the event that $(|M_n-M_0|)_{n\geq 1}$ exceeds ε for the fist time for $n=k$. That is,

$$A_k=\big\{|M_k-M_0|>\varepsilon,\ \text{but}\ |M_j-M_0|\leq\varepsilon\ \text{for}\ j<k\big\}.$$

Note that $A_1,A_2,\ldots,A_n$ is a partition of A, i.e. pairwise disjoint events whose union equals A. Let $\mathbf{1}_{A_k}$ denote the indicator function of the event A_k, which is the random variable which takes the value 1 if A_k occurs, and 0 otherwise. We then have

$$\mathrm{Var}(M_n-M_0)=\mathbb{E}[(M_n-M_0)^2]\geq\mathbb{E}[(M_n-M_0)^2\,\mathbf{1}_A]=\sum_{k=1}^{n}\mathbb{E}[(M_n-M_0)^2\,\mathbf{1}_{A_k}].$$

By adding and subtracting M_k in the right-hand side we obtain

$$\sum_{k=1}^{n}\mathbb{E}\left[\Big((M_n-M_k)^2+2(M_n-M_k)(M_k-M_0)+(M_k-M_0)^2\Big)\,\mathbf{1}_{A_k}\right]$$

Since A_k is determined by $M_0, M_1, \dots, M_k$ we have

$$\mathbb{E}\left[(M_n - M_k)(M_k - M_0)\,\mathbf{1}_{A_k}\right] = \mathbb{E}\Big[(M_k - M_0)\,\mathbf{1}_{A_k}\,\mathbb{E}[M_n - M_k | M_0, \dots, M_k]\Big],$$

which by the martingale property (5) is zero. Combining the above, and using that $(M_n - M_k)^2 \geq 0$ and that $|M_k - M_0| > \varepsilon$ on the event A_k, we obtain

$$\mathrm{Var}(M_n - M_0) \geq \sum_{k=1}^{n} \mathbb{E}[(M_k - M_0)^2\,\mathbf{1}_{A_k}] \geq \sum_{k=1}^{n} \varepsilon^2\,\mathbb{P}(A_k) = \varepsilon^2\,\mathbb{P}(A),$$

as we set out to prove. □

Convergence almost surely

Having observed some key consequences of the martingale property, which the sequence $(Y_n)_{n\geq 0}$ enjoys, we are finally in position to determine that the convergence in Theorem 1 also holds in the stronger sense of almost sure convergence.

Theorem 4. *For Polya's urn, the proportion of red balls Y_n converges almost surely, as $n \to \infty$, to a random variable uniformly distributed on* $[0, 1]$.

So, we claim that the sequence $(Y_n)_{n\geq 0}$ converges almost surely. However, we have not yet come up with a reasonable candidate for the limit. We know from Theorem 1 that the limiting random variable (if it exists) has to be uniformly distributed, but it is not clear how we could go on and define this random variable. An alternative, when the limit is unknown, is to show that the sequence $(Y_n)_{n\geq 0}$ is a Cauchy sequence.

Recall that a real-valued sequence $(a_n)_{n\geq 0}$ is a *Cauchy sequence* if for every $\varepsilon > 0$ there exists $N \geq 1$ such that $|a_n - a_m| < \varepsilon$ for all $n, m \geq N$. Moreover, a sequence $(a_n)_{n\geq 0}$ is convergent if and only if it is a Cauchy sequence.

Our goal will thus be to show that $(Y_n)_{n\geq 0}$ is a Cauchy sequence with probability one. Before we attend to the proof we make the following observation, usually referred to as the continuity property of probability measures.

Lemma 5. *Let $A_1 \subseteq A_2 \subseteq \dots$ be an increasing sequence of events. Then,*

$$\lim_{n\to\infty} \mathbb{P}(A_n) = \mathbb{P}\bigg(\bigcup_{n\geq 1} A_n\bigg).$$

Proof. Let $B_1 := A_1$ and $B_{k+1} := A_{k+1} \setminus A_k$ for $k \geq 1$. The events $B_1, B_2, \dots$ are then pairwise disjoint, so that by (countable) additivity of probability measures we have

$$\mathbb{P}(A_n) = \mathbb{P}\bigg(\bigcup_{k=1}^{n} B_k\bigg) = \sum_{k=1}^{n} \mathbb{P}(B_k).$$

Sending $n \to \infty$, countable additivity leaves us with

$$\lim_{n\to\infty} \mathbb{P}(A_n) = \sum_{k\geq 1} \mathbb{P}(B_k) = \mathbb{P}\bigg(\bigcup_{k\geq 1} B_k\bigg) = \mathbb{P}\bigg(\bigcup_{n\geq 1} A_n\bigg),$$

as required. □

Proof of Theorem 4. We first note that by the triangle inequality we have

$$|Y_n - Y_m| \le |Y_n - Y_N| + |Y_N - Y_m|,$$

for all $n, m, N \ge 1$, so it will suffice to show that, with probability one, for every $\varepsilon > 0$ there exists $N \ge 1$ such that $|Y_{N+k} - Y_N| < \varepsilon$ for all $k \ge 1$. As a first step we show that for every fixed $\varepsilon > 0$ we have

$$\mathbb{P}\big(\exists N \ge 1 : |Y_{N+k} - Y_N| \le \varepsilon \text{ for all } k \ge 1\big) = 1. \tag{8}$$

Fix $\varepsilon > 0$ and $N \ge 1$. Since $(Y_n)_{n\ge0}$ is a martingale, also the sequence $(Z_n)_{n\ge0}$, where $Z_n := Y_{N+n}$, is a martingale. Hence, by Lemma 3, applied to the sequence $(Z_n)_{n\ge0}$, and (7) we obtain that

$$\mathbb{P}\Big(\max_{k=1,\dots,n} |Y_{N+k} - Y_N| > \varepsilon\Big) \le \frac{1}{\varepsilon^2}\operatorname{Var}(Y_{N+n} - Y_N) \le \frac{1}{\varepsilon^2 N}.$$

Since the events in the left-hand side are increasing, sending n to infinity, Lemma 5 gives that

$$\mathbb{P}\Big(\max_{k\ge1} |Y_{N+k} - Y_N| > \varepsilon\Big) \le \frac{1}{\varepsilon^2 N}. \tag{9}$$

From (9) we have for each fixed $\ell \ge 1$ that

$$\begin{aligned}\mathbb{P}\Big(\bigcup_{N\ge1}\big\{\max_{k\ge1}|Y_{N+k} - Y_N| \le \varepsilon\big\}\Big) &= 1 - \mathbb{P}\Big(\bigcap_{N\ge1}\big\{\max_{k\ge1}|Y_{N+k} - Y_N| > \varepsilon\big\}\Big)\\ &\ge 1 - \mathbb{P}\Big(\max_{k\ge1}|Y_{\ell+k} - Y_\ell| > \varepsilon\Big)\\ &\ge 1 - \frac{1}{\varepsilon^2\ell}.\end{aligned}$$

Since $\ell \ge 1$ was arbitrary, (8) follows.

To complete the proof, let, for $n \ge 1$, A_n denote the event that there exists $N \ge 1$ such that $|Y_{N+k} - Y_N| \le 1/n$ for all $k \ge 1$. By (8) we have $\mathbb{P}(A_n) = 1$ for all $n \ge 1$. Consequently,

$$\begin{aligned}&\mathbb{P}\big(\forall\varepsilon > 0\ \exists N \ge 1 : |Y_{N+k} - Y_k| \le \varepsilon \text{ for all } k \ge 1\big)\\ &\quad = \mathbb{P}\Big(\bigcap_{n\ge1} A_n\Big) = 1 - \mathbb{P}\Big(\bigcup_{n\ge1} A_n^c\Big) \ge 1 - \sum_{n\ge1}\mathbb{P}(A_n^c) = 1.\end{aligned}$$

Consequently, with probability one, the sequence $(Y_n)_{n\ge0}$ is a Cauchy sequence, and the proof is complete. □

Once we know that $(Y_n)_{n\ge0}$ does converge with probability one (almost surely that is), then we can *define* a limiting variable Y as follows. Write $(\Omega, \mathcal{F}, \mathbb{P})$ for the probability space (which we have not defined explicitly) on which the sequence $(Y_n)_{n\ge0}$ is defined. Now let $Y(\omega) := \lim_{n\to\infty} Y_n(\omega)$ for $\omega \in \Omega$ for which the limit exist, and let $Y(\omega) := 0$ otherwise. Since $\lim_{n\to\infty} Y_n(\omega)$ exists with probability one, it follows that $Y_n \to Y$ almost surely, as $n \to \infty$.

Closing remarks

Polya's urn can be generalized in a range of directions. For instance, the initial composition of the urn could be modified, or the rules of reinforcement could be chosen to favour one of the colours. Urn models also appear frequently in applications. Some common techniques, and a long list of applications of urns in other sciences, can be found in [1].

References

[1] Robin Pemantle. A survey of random processes with reinforcement. *Probability Surveys*, 4:1-79, 2007.