Lecture 2

1 Martingales

We now introduce some fundamental tools in martingale theory, which are useful in controlling the fluctuation of martingales.

1.1 Doob's inequality

We have the following maximal inequality of Doob, which allows us to control the fluctuation of the whole path $(X_i)_{0\leq i\leq n}$ by controlling the fluctuation of the end point X_n, when X is a (sub-)martingale.

Theorem 1.1 [Doob's maximal inequality]
Let $(X_i)_{i\in\mathbb{N}}$ be a sub-martingale w.r.t. a filtration $(\mathcal{F}_i)_{i\in\mathbb{N}}$. Let $S_n = \max_{1\leq i\leq n} X_i$ be the running maximum of X_i. Then for any $l > 0$,

$$\mathbb{P}(S_n \geq l) \leq \frac{1}{l}\,\mathbb{E}\big[X_n^+ 1_{\{S_n\geq l\}}\big] \leq \frac{1}{l}\,\mathbb{E}[X_n^+], \tag{1.1}$$

where $X_n^+ = X_n \vee 0$. In particular, if X_i is a martingale and $M_n = \max_{1\leq i\leq n} |X_i|$, then

$$\mathbb{P}(M_n \geq l) \leq \frac{1}{l}\,\mathbb{E}\big[|X_n| 1_{\{M_n\geq l\}}\big] \leq \frac{1}{l}\,\mathbb{E}[|X_n|]. \tag{1.2}$$

Proof. Let $\tau_l = \inf\{i \geq 1 : X_i \geq l\}$. Then

$$\mathbb{P}(S_n \geq l) = \sum_{i=1}^{n} \mathbb{P}(\tau_l = i). \tag{1.3}$$

For each $1 \leq i \leq n$,

$$\mathbb{P}(\tau_l = i) = \mathbb{E}[1_{\{X_i\geq l\}} 1_{\{\tau_l=i\}}] \leq \frac{1}{l}\,\mathbb{E}[X_i^+ 1_{\{\tau_l=i\}}]. \tag{1.4}$$

Note that $\{\tau_l = i\} \in \mathcal{F}_i$, and X_i^+ is a sub-martingale because X_i itself is a sub-martingale while $\phi(x) = x^+$ is an increasing convex function. Therefore

$$\mathbb{E}[X_n^+ 1_{\{\tau_l=i\}}|\mathcal{F}_i] = 1_{\{\tau_l=i\}}\mathbb{E}[X_n^+|\mathcal{F}_i] \geq 1_{\{\tau_l=i\}}\mathbb{E}[X_n|\mathcal{F}_i]^+ \geq 1_{\{\tau_l=i\}} X_i^+,$$

and hence

$$\mathbb{E}[X_i^+ 1_{\{\tau_l=i\}}] \leq \mathbb{E}[X_n^+ 1_{\{\tau_l=i\}}].$$

Substituting this inequality into (1.4) and then summing over $1 \leq i \leq n$ then yields (1.1). (1.2) follows by applying (1.1) to the sub-martingale $|X_i|$. ∎

Corollary 1.2 *Let X, S and M be as in Theorem 1.1. Then for any $p \geq 1$, we have*

$$\begin{aligned}
\mathbb{P}(S_n \geq l) &\leq \frac{1}{l^p}\,\mathbb{E}\big[(X_n^+)^p 1_{\{S_n\geq l\}}\big] \leq \frac{1}{l^p}\,\mathbb{E}[(X_n^+)^p], &(1.5)\\
\mathbb{P}(M_n \geq l) &\leq \frac{1}{l^p}\,\mathbb{E}\big[|X_n|^p 1_{\{M_n\geq l\}}\big] \leq \frac{1}{l^p}\,\mathbb{E}[|X_n|^p]. &(1.6)
\end{aligned}$$

Proof. Apply Theorem 1.1 respectively to the sub-martingales $(X_n^+)^p$ and $|X_n|^p$. ∎

The bound we have on $\mathbb{P}(S_n \geq l)$ allows us to obtain bounds on the moments of S_n^+.

Corollary 1.3 [Doob's L_p maximal inequality]
Let X_i and S_i be as in Theorem 1.1. Then for $p > 1$, we have

$$\mathbb{E}[(S_n^+)^p] \leq \big(\frac{p}{p-1}\big)^p \, \mathbb{E}[(X_n^+)^p]. \tag{1.7}$$

Proof. Note that by the *layer-cake representation* of an integral, we have

$$\begin{aligned} \mathbb{E}[(S_n^+)^p] &= \int_0^\infty \mathbb{P}((S_n^+)^p > t)\mathrm{d}t = p\int_0^\infty l^{p-1}\mathbb{P}(S_n \geq l)\mathrm{d}l \\ &\leq p\int_0^\infty l^{p-2}\mathbb{E}[X_n^+ 1_{\{S_n \geq l\}}]\,\mathrm{d}l = p\,\mathbb{E}\Big[X_n^+\int_0^\infty l^{p-2}1_{\{S_n\geq l\}}\,\mathrm{d}l\Big] \\ &= \frac{p}{p-1}\mathbb{E}\big[X_n^+(S_n^+)^{p-1}\big] \leq \frac{p}{p-1}\mathbb{E}[(X_n^+)^p]^{\frac{1}{p}}\mathbb{E}[(S_n^+)^p]^{\frac{p-1}{p}}. \end{aligned}$$

If $\mathbb{E}[(S_n^+)^p] < \infty$, then (1.7) follows immediately. Otherwise, we can first replace S_n^+ by $S_n^+ \wedge N$ and repeat the above estimates, and then send $N \to \infty$ and apply the monotone convergence theorem. ∎

Example 1.4 *Let $X_n = \sum_{i=1}^n \xi_i$ be a random walk with $X_0 = 0$, where $(\xi_i)_{i\in\mathbb{N}}$ are i.i.d. random variables with $\mathbb{E}[\xi_1] = 0$ and $\mathbb{E}[\xi_1^2] = \sigma^2 < \infty$. Then X_n is a* square-integrable martingale. *Let $S_n = \max_{0\leq i\leq n} X_i$. By Doob's inequalities, we have*

$$\mathbb{P}\Big(\frac{S_n}{\sqrt{n}} \geq a\Big) \leq \frac{\mathbb{E}[X_n^2]}{na^2} = \frac{\sigma^2}{a^2} \qquad \textit{and} \qquad \mathbb{E}\Big[\Big(\frac{S_n}{\sqrt{n}}\Big)^2\Big] \leq \frac{4}{n}\mathbb{E}[X_n^2] = 4\sigma^2. \tag{1.8}$$

Exercise 1.5 *Show that when ξ_i are bounded, the tail bound for S_n in (1.8) can be improved to a Gaussian tail bound as in the Azuma-Hoeffding inequality.*

Exercise 1.6 *Doob's L_p maximal inequality fails for $p = 1$. Indeed, try to construct a non-negative martingale X_n with $\mathbb{E}[X_n] \equiv 1$, and yet $\sup_{n\in\mathbb{N}} \mathbb{E}[S_n] = \infty$. To get a bound on $\mathbb{E}[S_n^+]$, we need a bit more than $\mathbb{E}[X_n^+] < \infty$.*

Exercise 1.7 *Let X_n be a sub-martingale and let $S_n = \max_{1\leq i\leq n} X_i$. By mimicking the proof of Doob's L_p maximal inequality and by using $a\log b \leq a\log a + b/e$ for $a, b > 0$, show that*

$$\mathbb{E}[S_n^+] \leq \frac{1 + \mathbb{E}[X_n^+\log^+(X_n^+)]}{1 - e^{-1}},$$

where $\log^+ x = \max\{0, \log x\}$.

1.2 Stopping times

It is often useful to stop a stochastic process at a time which is determined from past observations of the process. Such times are called *stopping times*.

Definition 1.8 [Stopping time]
Given $(\Omega, \mathcal{F}, \mathbb{P})$, let $(\mathcal{F}_n)_{n\in\mathbb{N}}$ be a filtration in $\mathcal{F}$. A $\{0\}\cup\mathbb{N}\cup\{\infty\}$-valued random variable τ is called a stopping time w.r.t. $(\mathcal{F}_n)_{n\in\mathbb{N}}$, if for every $n \geq 0$, the event $\{\omega \in \Omega : \tau(\omega) = n\} \in \mathcal{F}_n$.

Remark 1.9 *Intuitively, a stopping time τ is a decision on when to stop the stochastic process $(X_n)_{n\in\mathbb{N}}$, using only information up to that time. Examples of stopping times include: $\tau := k$ for some fixed $k \geq 0$; $\tau := \inf\{n \geq 0 : X_n \geq a\}$ for some $a \in \mathbb{R}$, i.e., the first passage time of level a for a process $(X_n)_{n\in\mathbb{N}}$ adapted to $(\mathcal{F}_n)_{n\in\mathbb{N}}$; $\tau := \tau_1 \wedge \tau_2$ or $\tau := \tau_1 \vee \tau_2$, the minimum or maximum of two stopping times τ_1 and τ_2.*

Definition 1.10 [Stopped σ-field]
Let τ be a stopping time w.r.t. the filtration $(\mathcal{F}_n)_{n\in\mathbb{N}}$. The stopped σ-field $\mathcal{F}_\tau$ associated with τ is defined to be

$$\mathcal{F}_\tau := \{A \in \mathcal{F} : A \cap \{\omega : \tau(\omega) = n\} \in \mathcal{F}_n \text{ for all } n \geq 0\}. \tag{1.9}$$

Remark 1.11 *$\mathcal{F}_\tau$ can be interpreted as the information available up to the stopping time τ. For each event $A \in \mathcal{F}_\tau$, we can determine whether it has occurred or not based on what we have observed about the process up to (and including) time τ.*

Exercise 1.12 *Verify that $\mathcal{F}_\tau$ is indeed a σ-algebra, and show that the definition of stopping times and stopped σ-fields are unchanged if we replace $\tau(\omega) = n$ by $\tau(\omega) \leq n$, but not by $\tau(\omega) < n$.*

Theorem 1.13 [Stopped martingales are martingales]
Let $(X_n)_{n\in\mathbb{N}}$ be a martingale, and τ a stopping time, then the stopped martingale $(X_{n\wedge\tau})_{n\in\mathbb{N}}$ is also a martingale. More generally, if θ is another stopping time and $\theta \leq \tau$, then $X_{n\wedge\tau} - X_{n\wedge\theta}$ is a martingale. If X is a sub/super-martingale, then $X_{n\wedge\tau} - X_{n\wedge\theta}$ is also a sub/super-martingale.

Proof. We will verify that $X_{n\wedge\tau} - X_{n\wedge\theta}$ is a super-martingale when X is a super-martingale. The rest then follows. Note that

$$X_{n\wedge\tau} - X_{n\wedge\theta} = \sum_{i=1}^{n} 1_{\{\theta < i \leq \tau\}}(X_i - X_{i-1}),$$

where $\{\theta < i \leq \tau\} \in \mathcal{F}_{i-1}$ since $\{i \leq \tau\} = \{\tau \leq i-1\}^c \in \mathcal{F}_{i-1}$. Therefore, $X_{n\wedge\tau} - X_{n\wedge\theta}$ is a martingale transform of X_n, and

$$\begin{aligned}\mathbb{E}[X_{n\wedge\tau} - X_{n\wedge\theta}|\mathcal{F}_{n-1}] &= X_{(n-1)\wedge\tau} - X_{(n-1)\wedge\theta} + 1_{\{\theta < n \leq \tau\}}\mathbb{E}[X_n - X_{n-1}|\mathcal{F}_{n-1}] \\ &\leq X_{(n-1)\wedge\tau} - X_{(n-1)\wedge\theta}.\end{aligned}$$

Therefore $X_{n\wedge\tau} - X_{n\wedge\theta}$ is a super-martingale when X_n is a super-martingale. ∎

If the price of a stock evolves in time as a martingale, then Theorem 1.13 tells us that no matter when do we decide to buy and sell the stock, as long as our strategy is only based on past observations, the expected payoff will be zero.

1.3 Upcrossing inequality, almost sure Martingale Convergence, and Polya's urn

As an application of the notion of stopped martingales, we prove the upcrossing inequality. Let $(X_i)_{0\leq i\leq n}$ be a super-martingale adapted to the filtration $(\mathcal{F}_i)_{0\leq i\leq n}$ on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $a < b$. An upcrossing by X over the interval (a, b) is a pair of indices

$0 \leq k < l \leq n$ with $X_k \leq a$ and $X_l \geq b$. We are interested in the number U_n of complete upcrossings X makes before time n. Define recursively

$$\begin{aligned}
\tau_1 &:= \inf\{i : X_i \leq a\}, \\
\tau_2 &:= \inf\{i \geq \tau_1 : X_i \geq b\}, \\
&\vdots \\
\tau_{2k+1} &:= \inf\{i \geq \tau_{2k} : X_i \leq a\}, \\
\tau_{2k+2} &:= \inf\{i \geq \tau_{2k+1} : X_i \geq b\}, \\
&\vdots \quad ,
\end{aligned}$$

where the infimum of an empty set is taken to be ∞. Note that τ_i are all stopping times, and the number of completed upcrossings before time n is given by $U_n = \max\{k : \tau_{2k} \leq n\}$.

Theorem 1.14 [Upcrossing inequality]
Let $(X_i)_{0 \leq i \leq n}$ be a (super-)martingale and let U_n be the number of complete upcrossings over (a, b) defined as above. Then

$$\mathbb{E}[U_n] \leq \frac{\mathbb{E}[(a - X_n)^+]}{b - a} \leq \frac{|a| + \mathbb{E}[|X_n^-|]}{b - a}. \tag{1.10}$$

Proof. By Theorem (1.13), $(X_{i \wedge \tau_{2k}} - X_{i \wedge \tau_{2k-1}})_{0 \leq i \leq n}$ is a super-martingale for each $1 \leq k \leq n$. Therefore

$$\mathbb{E}\Big[\sum_{k=1}^{n} \big(X_{n \wedge \tau_{2k}} - X_{n \wedge \tau_{2k-1}}\big)\Big] \leq 0. \tag{1.11}$$

On the other hand, note that

$$X_{n \wedge \tau_{2k}} - X_{n \wedge \tau_{2k-1}} = \begin{cases} X_{\tau_{2k}} - X_{\tau_{2k-1}} \geq b - a & \text{if } 1 \leq k \leq U_n, \\ X_n - X_{\tau_{2U_n+1}} \geq X_n - a & \text{if } k = U_n + 1 \text{ and } \tau_{2U_n+1} \leq n, \\ X_n - X_n = 0 & \text{if } k = U_n + 1 \text{ and } \tau_{2U_n+1} > n, \\ X_n - X_n = 0 & \text{if } k \geq U_n + 2. \end{cases}$$

Therefore,

$$\sum_{k=1}^{n} \big(X_{n \wedge \tau_{2k}} - X_{n \wedge \tau_{2k-1}}\big) \geq (b - a)U_n + 1_{\{\tau_{2U_n+1} \leq n\}}(X_n - a).$$

Taking expectation and combined with (1.11) then yields

$$(b - a)\mathbb{E}[U_n] \leq \mathbb{E}[(a - X_n)1_{\{\tau_{2U_n+1} \leq n\}}] \leq \mathbb{E}[(a - X_n)^+], \tag{1.12}$$

and hence (1.10). To summarize, the number of upcrossings has to be balanced out by X_n^- if X is a super-martingale, which explains why $\mathbb{E}[U_n]$ is controlled by $\mathbb{E}[|X_n^-|]$. ∎

Remark 1.15 *Note that the number of upcrossings and downcrossings differ by at most 1. When X_n is a sub-martingale, the expected number of upcrossings can be bounded in terms of $\mathbb{E}[(X_n - a)^+]$. See Theorem 4.(2.9) in Durrett [1].*

A consequence of the upcrossing inequality is the almost sure martingale convergence theorem for L_1-bounded martingales.

Theorem 1.16 [Martingale convergence theorem]
Let $(X_n)_{n\in\mathbb{N}}$ be a (super)-martingale with $\sup_n \mathbb{E}[|X_n^-|] < \infty$. Then $X = \lim_{n\to\infty} X_n$ exists almost surely, and $\mathbb{E}[|X|] < \infty$. If X_n is a sub-martingale, then the condition becomes $\sup_n \mathbb{E}[X_n^+] < \infty$.

Proof. By the upcrossing inequality, for any $a < b$,

$$\sup_{n\in\mathbb{N}} \mathbb{E}[U_n] \le \sup_{n\in\mathbb{N}} \frac{|a| + \mathbb{E}[|X_n^-|]}{b-a} \le \frac{|a| + \sup_{n\in\mathbb{N}} \mathbb{E}[|X_n^-|]}{b-a} < \infty. \tag{1.13}$$

Since U_n almost surely increases to a limit $U(a,b)$, which is the total number of upcrossings by X over (a,b), by Fatou's Lemma, $\mathbb{E}[U(a,b)] < \infty$ and hence $U(a,b) < \infty$ almost surely. Thus, almost surely, $U(a,b) < \infty$ for all $a < b \in \mathbb{Q}$, which implies that

$$\limsup_{n\to\infty} X_n = \liminf_{n\to\infty} X_n,$$

and hence $X = \lim_{n\to\infty} X_n$ exists in $[-\infty, \infty]$ almost surely.

By Fatou's Lemma,

$$\mathbb{E}[|X^-|] \le \liminf_{n\to\infty} \mathbb{E}[|X_n^-|] < \infty$$

by assumption. Similarly, by Fatou and the fact that X_n is a super-martingale,

$$\mathbb{E}[X^+] \le \liminf_{n\to\infty} \mathbb{E}[X_n^+] = \liminf_{n\to\infty} (\mathbb{E}[X_n] + \mathbb{E}[|X_n^-|]) \le \liminf_{n\to\infty} (\mathbb{E}[X_1] + \mathbb{E}[|X_n^-|]) < \infty.$$

Therefore $\mathbb{E}[|X|] < \infty$. ∎

Corollary 1.17 [Almost sure convergence of a non-negative supermartingale]
If $(X_n)_{n\in\mathbb{N}}$ is a non-negative super-martingale, then $X = \lim_{n\to\infty} X_n$ exists a.s., and $\mathbb{E}[X] \le \mathbb{E}[X_1]$.

Remark 1.18 *In Corollary 1.17, we could have $\mathbb{E}[X] < \mathbb{E}[X_1]$ since part of the measure $X_n \mathrm{d}\mathbb{P}$ could escape to ∞. For example, let $X_n = X_0 + \sum_{i=1}^n \xi_i$ be a symmetric simple random walk on $\mathbb{Z}$ starting from $X_0 = 1$, where ξ_i are i.i.d. with $\mathbb{P}(\xi_i = 1) = \mathbb{P}(\xi_i = -1) = \frac{1}{2}$. Let $\tau = \inf\{n : X_n = 0\}$ be the first hitting time of the origin. Then $X_{n\wedge\tau}$ is a non-negative martingale which converges to a limit X_τ almost surely. The only possible limit for $X_{n\wedge\tau}$ is either 0 or ∞. Since $\mathbb{E}[X_\tau] \le \mathbb{E}[X_1] = 1$ by Corollary 1.17, we must have $X_\tau = 0$ almost surely. Note that this also implies that $\tau < \infty$ almost surely, so that the symmetric simple random walk always visits 0, which is a property called* **recurrence**.

Exercise 1.19 *Show by example that a non-negative sub-martingale need not converge almost surely.*

As an application of Corollary 1.17, we have the following dichotomy between *convergence* and *unbounded oscillation* for martingales with **bounded increments**.

Theorem 1.20 *Let $(X_n)_{n\in\mathbb{N}}$ be a martingale with $|X_{n+1} - X_n| \le M < \infty$ a.s. for all $n \ge 0$. Then almost surely, either $\lim_{n\to\infty} X_n$ exists and is finite, or $\limsup_{n\to\infty} X_n = \infty$ and $\liminf_{n\to\infty} X_n = -\infty$.*

Proof. For $L < 0$, let $\tau_L = \inf\{n : X_n \leq L\}$, which is a stopping time. By assumption, $X_{\tau_L} \geq L - M$. Therefore $X_{n\wedge\tau_L} - L + M$ is a non-negative martingale, which converges a.s. to a finite limit. In particular, on the event $\{\tau_L = \infty\}$, $\lim_{n\to\infty} X_n$ exists and is finite. Letting $L \to -\infty$, it follows that on the event $\{\liminf_{n\to\infty} X_n > -\infty\}$, $\lim_{n\to\infty} X_n$ exists and is finite. Applying the same argument to $-X_n$ implies that $\lim_{n\to\infty} X_n$ exists and is finite on the event that $\{\limsup_{n\to\infty} X_n < \infty\}$. The theorem then follows. ∎

Exercise 1.21 *Construct an example where the conclusion in Theorem 1.20 fails if the bounded increment assumption is removed. (Hint: Let $X_n = \sum_{i=1}^n \xi_i$, where ξ_i are independent with mean 0 but not i.i.d., such that $X_n \to -\infty$ almost surely.)*

Recall that the Borel-Cantelli lemma states: If $(A_n)_{n\in\mathbb{N}} \in \mathcal{F}$ satisfy $\sum_{i=1}^\infty \mathbb{P}(A_n) < \infty$, then almost surely, $(A_n)_{n\in\mathbb{N}}$ occurs only finitely many times. If $(A_n)_{n\in\mathbb{N}}$ are all independent, then $\sum_{i=1}^\infty \mathbb{P}(A_n) = \infty$ guarantees that almost surely, $(A_n)_{n\in\mathbb{N}}$ occurs infinitely often. Using Theorem 1.20, we give a proof of the *second Borel-Cantelli lemma* which allows dependence of events.

Corollary 1.22 [Second Borel-Cantelli Lemma]
Let $(\mathcal{F}_n)_{n\geq 0}$ be a filtration on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with $\mathcal{F}_0 = \{\emptyset, \Omega\}$. Let $A_n \in \mathcal{F}_n$ for $n \geq 1$. Then

$$\Big\{\omega : \sum_{n=1}^{\infty} 1_{A_n}(\omega) = \infty\Big\} = \Big\{\omega : \sum_{n=1}^{\infty} \mathbb{P}(A_n|\mathcal{F}_{n-1}) = \infty\Big\} \tag{1.14}$$

modulo a set of probability 0. When A_n are independent, we retrieve the classic Borel-Cantelli lemma.

Proof. Let $X_n = \sum_{i=1}^n \big(1_{A_i}(\omega) - \mathbb{P}(A_i|\mathcal{F}_{i-1})\big)$. Note that X_n is a martingale with bounded increments. Therefore by Theorem 1.20, either X_n converges to a finite limit, in which case $\sum_{n=1}^\infty 1_{A_n}(\omega) = \infty$ if and only if $\sum_{n=1}^\infty \mathbb{P}(A_n|\mathcal{F}_{n-1}) = \infty$; or X_n oscillates between $\pm\infty$, in which case we have $\sum_{n=1}^\infty 1_{A_n}(\omega) = \sum_{n=1}^\infty \mathbb{P}(A_n|\mathcal{F}_{n-1}) = \infty$. ∎

We now study the example of Polya's urn using martingale techniques.

Example 1.23 [Polya's urn] *Let an urn initially contain b black balls and w white balls. Each time, we pick a ball from the urn with uniform probability, and we put back in the urn 2 balls of the same color. Obviously the number of balls in the urn increase by one each time. A natural question is what is the fraction X_n of black balls in the urn after time step n? What is the asymptotic distribution of X_n as $n \to \infty$?*

Note that $X_0 = \frac{b}{b+w}$. It turns out that X_n is a martingale. To check this, assume that out of the $b + w + n$ balls at time n, j balls are black and $k = b + w + n - j$ balls are white. Then $X_n = \frac{j}{j+k}$, and Then

$$\begin{aligned}
\mathbb{E}\Big[X_{n+1}|X_1, \cdots, X_n\Big] &= \frac{j}{j+k} \cdot \frac{j+1}{j+k+1} + \frac{k}{j+k} \cdot \frac{j}{j+k+1} \\
&= \frac{j}{j+k} = X_n.
\end{aligned}$$

Since X_n is non-negative, by the martingale convergence theorem, almost surely, X_n converges to a limit $X_\infty \in [0, 1]$.

Polya's urn has the special property of exchangeability. If Y_n denotes the indicator event that the ball drawn at the n-th time step is black, then the distribution of $(Y_n)_{n\in\mathbb{N}}$ is invariant under finite permutation of the indices. In particular, if $(Y_i)_{1\leq i\leq n}$ and $(\tilde{Y}_i)_{1\leq i\leq n}$ are two different realizations of Polya's urn up to time n with $\sum_{i=1}^n Y_i = \sum_{i=1}^n \tilde{Y}_i = m$, then observe that

$$\mathbb{P}((Y_i)_{1\leq i\leq n}) = \mathbb{P}((\tilde{Y}_i)_{1\leq i\leq n}) = \prod_{i=1}^{n} \frac{1}{b+w+i-1} \prod_{i=1}^{m}(b+i-1) \prod_{i=1}^{n-m}(w+i-1).$$

Assume $b = w = 1$ for simplicity, then after the n-th time step, for each $0 \leq j \leq n$,

$$\mathbb{P}\Big(X_n = \frac{j+1}{n+2}\Big) = \binom{n}{j}\frac{j!(n-j)!}{(n+1)!} = \frac{1}{n+1}.$$

Therefore X_∞ is uniformly distributed on $[0,1]$.

Check that for general $b, w > 0$, X_∞ follows the beta distribution with parameters b, w and density

$$\frac{\Gamma(b+w)}{\Gamma(b)\Gamma(w)}(1-x)^{b-1}(1-x)^{w-1}, \tag{1.15}$$

where $\Gamma(x)$ is the gamma function.

References

[1] R. Durrett, *Probability: Theory and Examples*, 2nd edition, Duxbury Press, Belmont, California, 1996.

[2] S.R.S. Varadhan, *Probability Theory*, Courant Lecture Notes 7, American Mathematical Society, Providence, Rhode Island, 2001.