Chapter 5

Martingales.

5.1 Definitions and properties

The theory of martingales plays a very important ans ueful role in the study of stochastic processes. A formal definition is given below.

Definition 5.1. *Let $(\Omega, \mathcal{F}, P)$ be a probability space. A* **martingale** *sequence of length n is a chain $X_1, X_2, \cdots, X_n$ of random variables and corresponding sub σ-fields $\mathcal{F}_1, \mathcal{F}_2, \cdots, \mathcal{F}_n$ that satisfy the following relations*

1. *Each X_i is an integrable random variable which is measurable with respect to the corresponding σ-field $\mathcal{F}_i$.*

2. *The σ-field $\mathcal{F}_i$ are increasing i.e. $\mathcal{F}_i \subset \mathcal{F}_{i+1}$ for every i.*

3. *For every $i \in [1, 2, \cdots, n-1]$, we have the relation*

$$X_i = E\{X_{i+1} | \mathcal{F}_i\} \quad \text{a.e.} \quad P.$$

Remark 5.1. We can have an infinite martingale sequence $\{(X_i, \mathcal{F}_i) : i \geq 1\}$ which requires only that for every n, $\{(X_i, \mathcal{F}_i) : 1 \leq i \leq n\}$ be a martingale sequence of length n. This is the same as conditions (i), (ii) and (iii) above except that they have to be true for every $i \geq 1$.

Remark 5.2. From the properties of conditional expectations we see that $E\{X_i\} = E\{X_{i+1}\}$ for every i, and therfore $E\{X_i\} = c$ for some c. We can define $\mathcal{F}_0$ to be the trivial σ-field consisting of $\{\Phi, \Omega\}$ and $X_0 = c$. Then $\{(X_i, \mathcal{F}_i) : \ i \geq 0\}$ is a martingale sequence as well.

Remark 5.3. We can define $Y_{i+1} = X_{i+1} - X_i$ so that $X_j = c + \sum_{1 \leq i \leq j} Y_i$ and property (iii) reduces to

$$E\{Y_{i+1} | \mathcal{F}_i\} = 0 \quad \text{a.e.} \quad P.$$

Such sequences are called martingale differences. If Y_i is a sequence of independent random variables with mean 0, for each i, we can take $\mathcal{F}_i$ to be the σ-field generated by the random variables $\{Y_j : 1 \leq j \leq i\}$ and $X_j = c + \sum_{1 \leq i \leq j} Y_i$, will be a martingale relative to the σ-fields $\mathcal{F}_i$.

Remark 5.4. We can generate martingale sequences by the following procedure. Given any increasing family of σ-fields $\{\mathcal{F}_j\}$, and any integrable random variable X on $(\Omega, \mathcal{F}, P)$, we take $X_i = E\{X | \mathcal{F}_i\}$ and it is easy to check that $\{(X_i, \mathcal{F}_i)\}$ is a martingale sequence. Of course every finite martingale sequence is generated this way for we can always take X to be X_n, the last one. For infinite sequences this raises an important question that we will answer later.

If one participates in a 'fair ' gambling game, the asset X_n of the player at time n is supposed to be a martingale. One can take for $\mathcal{F}_n$ the σ-field of all the results of the game through time n. The condition $E[X_{n+1} - X_n | \mathcal{F}_n] = 0$ is the assertion that the game is neutral irrespective of past history.

A related notion is that of a super or sub-martingale. If, in the definition of a martingale, we replace the equality in (iii) by an inequality we get super or sub-martingales.

For a **sub-martingale** we demand the relation

(iiia) for every i,

$$X_i \leq E\{X_{i+1} | \mathcal{F}_i\} \quad \text{a.e.} \quad P.$$

while for a **super-martingale** the relation is

(iiib) for every i,

$$X_i \geq E\{X_{i+1} | \mathcal{F}_i\} \quad \text{a.e.} \quad P.$$

Lemma 5.1. *If $\{(X_i\,, \mathcal{F}_i)\}$ is a martingale and φ is a convex (or concave) function of one variable such that $\varphi(X_n)$ is integrable for every n, then $\{(\varphi(X_n), \mathcal{F}_n)\}$ is a sub (or super)-martingale.*

Proof. An easy consequence of Jensen's inequality (4.2) for conditional expectations. □

Remark 5.5. A particular case is $\phi(x) = |x|^p$ with $1 \le p < \infty$. For any martingale $(X_n, \mathcal{F}_n)$ and $1 \le p < \infty$, $(|X_n|^p, \mathcal{F}_n)$ is a sub-martingale provided $E[|X_n|^p] < \infty$

Theorem 5.2. (Doob's inequality.) *Suppose $\{X_j\}$ is a martingale sequence of length n. Then*

$$P\left\{\omega : \sup_{1\le j\le n} |X_j| \ge \ell\right\} \le \frac{1}{\ell}\int_{\{\sup_{1\le j\le n}|X_j|\ge \ell\}} |X_n|\,dP \le \frac{1}{\ell}\int |X_n|\,dP \quad (5.1)$$

Proof. Let us define $S(\omega) = \sup_{1\le j\le n}|X_j(\omega)|$. Then

$$\{\omega : S(\omega) \ge \ell\} = E = \cup_j E_j$$

is written as a disjoint union, where

$$E_j = \{\omega : |X_1(\omega)| < \ell, \cdots, |X_{j-1}| < \ell, |X_j| \ge \ell\}.$$

We have

$$P(E_j) \le \frac{1}{\ell}\int_{E_j} |X_j|\,dP \le \frac{1}{\ell}\int_{E_j} |X_n|\,dP. \quad (5.2)$$

The second inequality in (5.2) follows from the fact that $|x|$ is a convex function of x, and therefore $|X_j|$ is a sub-martingale. In particular $E\{|X_n||\mathcal{F}_j\} \ge |X_j|$ a.e. P and $E_j \in \mathcal{F}_j$. Summing up (5.2) over $j = 1, \cdots, n$ we obtain the theorem. □

Remark 5.6. We could have started with

$$P(E_j) \le \frac{1}{\ell^p}\int_{E_j} |X_j|^p\,dP$$

and obtained for $p \ge 1$

$$P(E_j) \le \frac{1}{\ell^p}\int_{E_j} |X_n|^p\,dP. \quad (5.3)$$

Compare it with (3.9) for $p = 2$.

This simple inequality has various implications. For example

Corollary 5.3. (Doob's Inequality.) *Let* $\{X_j : 1 \le j \le n\}$ *be a martingale. Then if, as before,* $S(\omega) = \sup_{1\le j\le n} |X_j(\omega)|$ *we have*

$$E[S^p] \le \left(\frac{p}{p-1}\right)^p E\,[\,|X_n|^p].$$

The proof is a consequence of the following fairly general lemma.

Lemma 5.4. *Suppose* X *and* Y *are two nonnegative random variables on a probability space such that for every* $\ell \ge 0$,

$$P\{Y \ge \ell\} \le \frac{1}{\ell}\int_{Y\ge \ell} X\,dP.$$

Then for every $p > 1$,

$$\int Y^p\,dP \le \left(\frac{p}{p-1}\right)^p \int X^p\,dP.$$

Proof. Let us denote the tail probability by $T(\ell) = P\{Y \ge \ell\}$. Then with $\frac{1}{p} + \frac{1}{q} = 1$, i.e. $(p-1)q = p$

$$\begin{aligned}
\int Y^p\,dP &= -\int_0^\infty y^p dT(y) = p\int_0^\infty y^{p-1}T(y)dy \quad \text{(integrating by parts)}\\
&\le p\int_0^\infty y^{p-1}\frac{dy}{y}\int_{Y\ge y} X\,dP \quad \text{(by assumption)}\\
&= p\int X\left[\int_0^Y y^{p-2}dy\right]dP \quad \text{(by Fubini's Theorem)}\\
&= \frac{p}{p-1}\int X\,Y^{p-1}\,dP\\
&\le \frac{p}{p-1}\left[\int X^p\,dP\right]^{\frac{1}{p}}\left[\int Y^{q(p-1)}\,dP\right]^{\frac{1}{q}} \quad \text{(by Hölder's inequality)}\\
&\le \frac{p}{p-1}\left[\int X^p\,dP\right]^{\frac{1}{p}}\left[\int Y^p\,dP\right]^{1-\frac{1}{p}}
\end{aligned}$$

This simplifies to

$$\int Y^p \, dP \le \left(\frac{p}{p-1}\right)^p \int X^p \, dP$$

provided $\int Y^p \, dP$ is finite. In general given Y, we can truncate it at level N to get $Y_N = \min(Y, N)$ and for $0 < \ell \le N$,

$$P\{Y_N \ge \ell\} = P\{Y \ge \ell\} \le \frac{1}{\ell} \int_{Y \ge \ell} X \, dP = \frac{1}{\ell} \int_{Y_N \ge \ell} X \, dP$$

with $P\{Y_N \ge \ell\} = 0$ for $\ell > N$. This gives us uniform bounds on $\int Y_N^p \, dP$ and we can pass to the limit. So we have the strong implication that the finiteness of $\int X^p \, dP$ implies the finiteness of $\int Y^p \, dP$. □

Exercise 5.1. The result is false for $p = 1$. Construct a nonnegative martingale X_n with $E[X_n] \equiv 1$ such that $\xi = \sup_n X_n$ is not integrable. Consider $\Omega = [0, 1]$, $\mathcal{F}$ is the Borel σ-field and P the Lebesgue measure. Suppose we take $\mathcal{F}_n$ to be the σ-field generated by intervals with end points of the form $\frac{j}{2^n}$ for some integer j. It corresponds to a partition with 2^n sets. Consider the random variables

$$X_n(x) = \begin{cases} 2^n & \text{\it for } 0 \le x \le 2^{-n} \\ 0 & \text{\it for } 2^{-n} < x \le 1. \end{cases}$$

Check that it is a martingale and calculate $\int \xi(x) \, dx$. This is the 'winning ' strategy of doubling one's bets until the losses are recouped.

Exercise 5.2. If X_n is a martingale such that the differences $Y_n = X_n - X_{n-1}$ are all square integrable, show that for $n \ne m$, $E\,[Y_n \, Y_m] = 0$. Therefore

$$E[X_n^2] = E[X_0]^2 + \sum_{j=1}^{n} E\,[Y_j^2].$$

If in addition, $\sup_n E[X_n^2] < \infty$, then show that there is a random variable X such that

$$\lim_{n \to \infty} E\,[\,|X_n - X|^2] = 0.$$

5.2 Martingale Convergence Theorems.

If $\mathcal{F}_n$ is an increasing family of σ-fields and X_n is a martingale sequence with respect to $\mathcal{F}_n$, one can always assume without loss of generality that the full σ-field $\mathcal{F}$ is the smallest σ-field generated by $\cup_n \mathcal{F}_n$. If for some $p \geq 1$, $X \in L_p$, and we define $X_n = E[X|\mathcal{F}_n]$ then X_n is a martingale and by Jensen's inequality, $\sup_n E[|X_n|^p] \leq E[|X|^p]$. We would like to prove

Theorem 5.5. *For $p \geq 1$, if $X \in L_p$, then $\lim_{n\to\infty} \|X_n - X\|_p = 0$.*

Proof. Assume that X is a bounded function. Then by the properties of conditional expectation $\sup_n \sup_\omega |X_n| < \infty$. In particular $E[X_n^2]$ is uniformly bounded. By Exercise 5.2, at the end of last section, $\lim_{n\to\infty} X_n = Y$ exists in L_2. By the properties of conditional expectations for $A \in \mathcal{F}_m$,

$$\int_A Y\,dP = \lim_{n\to\infty} \int_A X_n\,dP = \int_A X\,dP.$$

This is true for all $A \in \mathcal{F}_m$ for any m. Since $\mathcal{F}$ is generated by $\cup_m \mathcal{F}_m$ the above relation is true for $A \in \mathcal{F}$. As X and Y are $\mathcal{F}$ measurable we conclude that $X = Y$ a.e. P. See Exercise 4.1. For a sequence of functions that are bounded uniformnly in n and ω convergence in L_p are all equivalent and therefore convergence in L_2 implies the convergence in L_p for any $1 \leq p < \infty$. If now $X \in L_p$ for some $1 \leq p < \infty$, we can approximate it by $X' \in L_\infty$ so that $\|X' - X\|_p < \epsilon$. Let us denote by X'_n the conditional expectations $E[X'|\mathcal{F}_n]$. By the properties of conditional expectations $\|X'_n - X_n\|_p \leq \epsilon$ for all n, and as we saw earlier $\|X'_n - X'\|_p \to 0$ as $n \to \infty$. It now follows that

$$\limsup_{\substack{n\to\infty\\ m\to\infty}} \|X_n - X_m\|_p \leq 2\epsilon$$

and as $\epsilon > 0$ is arbitrary we are done. □

In general, if we have a martingale $\{X_n\}$, we wish to know when it comes from a random variable $X \in L_p$ in the sense that $X_n = E[X|\mathcal{F}_n]$.

Theorem 5.6. *If for some $p > 1$, a martingale $\{X_n\}$ is bounded in L_p, in the sense that $\sup_n \|X_n\|_p < \infty$, then there is a random variable $X \in L_p$ such that $X_n = E[X|\mathcal{F}_n]$ for $n \geq 1$. In particular $\|X_n - X\|_p \to 0$ as $n \to \infty$.*

Proof. Suppose $\|X_n\|_p$ is uniformly bounded. For $p > 1$, since L_p is the dual of L_q with $\frac{1}{p} + \frac{1}{q} = 1$, bounded sets are weakly compact. See [7] or [3]. We can therefore choose a subsequence X_{n_j} that converges weakly in L_p to a limit in the weak topology. We call this limit X. Then consider $A \in \mathcal{F}_n$ for some fixed n. The function $\mathbf{1}_A(\cdot) \in L_q$.

$$\int_A X\, dP =< \mathbf{1}_A, X >= \lim_{j\to\infty} < \mathbf{1}_A, X_{n_j} >= \lim_{j\to\infty} \int_A X_{n_j}\, dP = \int_A X_n\, dP.$$

The last equality follows from the fact that $\{X_n\}$ is a martingale, $A \in \mathcal{F}_n$ and $n_j > n$ eventually. It now follows that $X_n = E\,[\, X\,|\mathcal{F}_n\,]$. We can now apply the preceeding theorem. ☐

Exercise 5.3. For $p = 1$ the result is false. Example 5.1 gives us at the same time a counterexample of an L_1 bounded martingale that does not converge in L_1 and so cannot be represented as $X_n = E\,[\, X\,|\mathcal{F}_n\,]$.

We can show that the convergence in the preceeding theorems is also valid almost everywhere.

Theorem 5.7. *Let $X \in L_p$ for some $p \geq 1$. Then the martingale $X_n = E\,[X\,|\mathcal{F}_n\,]$ converges to X for almost all ω with respect to P.*

Proof. From Hölder's inequality $\|X\|_1 \leq \|X\|_p$. Clearly it is sufficient to prove the theorem for $p = 1$. Let us denote by $\mathcal{M} \subset L_1$ the set of functions $X \in L_1$ for which the theorem is true. Clearly $\mathcal{M}$ is a linear subset of L_1. We will prove that it is closed in L_1 and that it is dense in L_1. If we denote by $\mathcal{M}_n$ the space of $\mathcal{F}_n$ measurable functions in L_1, then $\mathcal{M}_n$ is a closed subspace of L_1. By standard approximation theorems $\cup_n \mathcal{M}_n$ is dense in L_1. Since it is obvious that $\mathcal{M} \supset \mathcal{M}_n$ for every n, it follows that $\mathcal{M}$ is dense in L_1. Let $Y_j \in \mathcal{M} \subset L_1$ and $Y_j \to X$ in L_1. Let us define $Y_{n,j} = E\,[Y_j\,|\mathcal{F}_n\,]$. With $X_n = E\,[X\,|\mathcal{F}_n\,]$, by Doob's inequality (5.1) and jensen's inequlaity (4.2),

$$\begin{aligned} P\Big\{ \sup_{1\leq n\leq N} |X_n| \geq \ell \Big\} &\leq \frac{1}{\ell} \int_{\{\omega : \sup_{1\leq n\leq N} |X_n| \geq \ell\}} |X_N|\, dP \\ &\leq \frac{1}{\ell} E\,[\,|X_N\,|\,] \\ &\leq \frac{1}{\ell} E\,[\,|X|\,] \end{aligned}$$

and therefore X_n is almost surely a bounded sequence. Since we know that $X_n \to X$ in L_1, it suffices to show that

$$\limsup_n X_n - \liminf_n X_n = 0 \quad \text{a.e.} \quad P.$$

If we write $X = Y_j + (X - Y_j)$, then $X_n = Y_{n,j} + (X_n - Y_{n,j})$, and

$$\begin{aligned}\limsup_n X_n - \liminf_n X_n &\le [\limsup_n Y_{n,j} - \liminf_n Y_{n,j}] \\ &\quad + [\limsup_n (X_n - Y_{n,j}) - \liminf_n (X_n - Y_{n,j})] \\ &= \limsup_n (X_n - Y_{n,j}) - \liminf_n (X_n - Y_{n,j}) \\ &\le 2 \sup_n |X_n - Y_{n,j}|.\end{aligned}$$

Here we have used the fact that $Y_j \in \mathcal{M}$ for every j and hence

$$\limsup_n Y_{n,j} - \liminf_n Y_{n,j} = 0 \quad \text{a.e.} \quad P.$$

Finally

$$\begin{aligned}P\Big\{\limsup_n X_n - \liminf_n X_n \ge \epsilon\Big\} &\le P\Big\{\sup_n |X_n - Y_{n,j}| \ge \frac{\epsilon}{2}\Big\} \\ &\le \frac{2}{\epsilon} E[|X - Y_j|] \\ &= 0\end{aligned}$$

since the left hand side is independent of j and the term on the right on the second line tends to 0 as $j \to \infty$. □

The only case where the situation is unclear is when $p = 1$. If X_n is an L_1 bounded martingale, it is not clear that it comes from an X. If it did arise from an X, then X_n would converge to it in L_1 and in particular would have to be uniformly integrable. The converse is also true.

Theorem 5.8. *If X_n is a uniformly integrable martingale then there is random variable X such that $X_n = E[X|\mathcal{F}_n]$, and then of course $X_n \to X$ in L_1.*

Proof. The uniform integrability of X_n implies the weak compactness in L_1 and if X is any weak limit of X_n [see [7]], it is not difficult to show as in Theorem 5.5, that $X_n = E[X|\mathcal{F}_n]$. □

Remark 5.7. Note that for $p > 1$, a martingale X_n that is bounded in L_p is uniformly integrable in L_p, i.e $|X_n|^p$ is uniformly integrable. This is false for $p = 1$. The L_1 bounded martingale that we constructed earlier in Exercise 5.1 as a counterexample, is not convergent ln L_1 and therefore can not be uniformly integrable. We will defer the analysis of L_1 bounded martingales to the next section.

5.3 Doob Decomposition Theorem.

The simplest example of a submartingale is a sequence of functions that is non decreasing in n for every (almost all) ω. In some sense the simplest example is also the most general. More precisely the decomposition theorem of Doob asserts the following.

Theorem 5.9. (Doob decomposition theorem.) *If $\{X_n : n \geq 1\}$ is a sub-martingale on $(\Omega, \mathcal{F}_n, P)$, then X_n can be written as $X_n = Y_n + A_n$, with the following properties:*

1. *$(Y_n, \mathcal{F}_n)$ is a martingale.*
2. *$A_{n+1} \geq A_n$ for almost all ω and for every $n \geq 1$.*
3. *$A_1 \equiv 0$.*
4. *For every $n \geq 2$, A_n is $\mathcal{F}_{n-1}$ measurable.*

X_n determines Y_n and A_n uniquely .

Proof. Let X_n be any sequence of integrable functions such that X_n is $\mathcal{F}_n$ measurable, and is represented as $X_n = Y_n + A_n$, with Y_n and A_n satisfying (1), (3) and (4). Then

$$A_n - A_{n-1} = E[X_n - X_{n-1}|\mathcal{F}_{n-1}] \tag{5.4}$$

are uniquely determined. Since $A_1 = 0$, all the A_n are uniquely determined as well. Property (2) is then plainly equivalent to the submartingale property. To establish the representation, we define A_n inductively by (5.4). It is routine to verify that $Y_n = X_n - A_n$ is a martingale and the monotonicity of A_n is a consequence of the submartingale property. ☐

Remark 5.8. Actually, given any sequence of integrable functions $\{X_n : n \geq 1\}$ such that X_n is $\mathcal{F}_n$ measurable, equation (5.4) along with $A_1 = 0$ defines $\mathcal{F}_{n-1}$ measurable functions that are integrable, such that $X_n = Y_n + A_n$ and $(Y_n, \mathcal{F}_n)$ is a martingale. The decomposition is always unique. It is easy to verify from (5.4) that $\{A_n\}$ is increasing (or decreasing) if and only if $\{X_n\}$ is a super- (or sub-) maringale. Such a decomposition is called the semi-martingale decomposition.

Remark 5.9. It is the demand that A_n be $\mathcal{F}_{n-1}$ measurable that leads to uniqueness. If we have to deal with continuous time this will become a thorny issue.

We now return to the study of L_1 bounded martingales. A nonnegative martingale is clearly L_1 bounded because $E[\,|X_n|\,] = E[\,X_n\,] = E[\,X_1\,]$. One easy way to generate L_1 bounded martingales is to take the difference of two nonneagtive martingales. We have the converse as well.

Theorem 5.10. *Let X_n be an L_1 bounded martingale. There are two non-negative martingales Y_n and Z_n relative to the same σ-fields $\mathcal{F}_n$, such that $X_n = Y_n - Z_n$.*

Proof. For each j and $n \geq j$, we define

$$Y_{j,n} = E[\,|\,X_n\,|\,|\mathcal{F}_j\,].$$

By the submartingale property of $|\,X_n\,|$

$$Y_{j,n+1} - Y_{j,n} = E[(|X_{n+1}| - |X_n|)\,|\mathcal{F}_j\,] = E[E[(|X_{n+1}| - |X_n|)\,|\mathcal{F}_n\,]|\mathcal{F}_j] \geq 0$$

almost surely. $Y_{j,n}$ is nonnegative and $E[\,Y_{j,n}\,] = E[\,|X_n|\,]$ is bounded in n. By the monotone convergence theorem, for each j, there exists some Y_j in L_1 such that $Y_{j,n} \to Y_j$ in L_1 as $n \to \infty$. Since limits of martingales are again martingales, and $Y_{n,j}$ is a martingale for $n \geq j$, it follows that Y_j is a martingale. Moreover

$$Y_j + X_j = \lim_{n\to\infty} E[\,|\,X_n\,| + X_n\,|\mathcal{F}_j\,] \geq 0$$

and

$$X_j = (Y_j + X_j) - Y_j$$

does it! □

We can always assume that our nonnegative martingale has its expectation equal to 1 because we can always multiply by a suitable constant. Here is a way in which such martingales arise. Suppose we have a probability space $(\Omega\,,\mathcal{F}\,,P\,)$ and and an increasing family of sub σ-fields $\mathcal{F}_n$ of $\mathcal{F}$ that generate $\mathcal{F}$. Suppose Q is another probability measure on $(\Omega\,,\mathcal{F}\,)$ which may or may not be absolutely continuous with respect to P on $\mathcal{F}$. Let us suppose however that $Q << P$ on each $\mathcal{F}_n$, i.e. whenever $A \in \mathcal{F}_n$ and $P(A) = 0$, it follows that $Q(A) = 0$. Then the sequence of Radon-Nikodym derivatives

$$X_n = \frac{dQ}{dP}\bigg|_{\mathcal{F}_n}$$

of Q with respect to P on $\mathcal{F}_n$ is a nonnegative martingale with expectation 1. It comes from an X, if and only if $Q << P$ on $\mathcal{F}$ and this is the uniformly integrable case. By Lebesgue decomposition we reduce our consideration to the case when $Q \perp P$. Let us change the reference measure to $P' = \frac{P+Q}{2}$. The Radon-Nikodym derivative

$$X'_n = \frac{dQ}{dP'}\bigg|_{\mathcal{F}_n} = \frac{2X_n}{1 + X_n}$$

is uniformly integrable with respect to P' and $X'_n \to X'$ a.e. P'. From the orthogonality $P \perp Q$ we know that there are disjoint sets E, E^c with $P(E) = 1$ and $Q(E^c) = 1$. Then

$$\begin{aligned} Q(A) &= Q(A \cap E) + Q(A \cap E^c) = Q(A \cap E^c) \\ &= 2P'(A \cap E^c) = \int_A 2\,\mathbf{1}_E(\omega) dP'. \end{aligned}$$

It is now seen that

$$X' = \frac{dQ}{dP'}\bigg|_{\mathcal{F}} = \begin{cases} 2 & a.e.\ Q \\ 0 & a.e.\ P \end{cases}$$

from which one concludes that

$$P\left\{\lim_{n\to\infty} X_n = 0\right\} = 1.$$

Exercise 5.4. In order to establish that a nonnegative martingale has an almost sure limit (which may not be an L_1 limit) show that we can assume, without loss of generality, that we are in the following situation.

$$\Omega = \otimes_{j=1}^{\infty} R \ ; \ \mathcal{F}_n = \sigma[x_1, \cdots, x_n] \ ; X_j(\omega) = x_j$$

The existence of a Q such that

$$\left.\frac{dQ}{dP}\right|_{\mathcal{F}_n} = x_n$$

is essentially Kolmogorov's consistency theorem (Theorem 3.5.) Now complete the proof.

Remark 5.10. We shall give a more direct proof of almost sure convergence of an L_1 bounded martingale later on by means of the upcrossing inequality.

5.4 Stopping Times.

The notion of stopping times that we studied in the context of Markov Chains is important again in the context of Martingales. In fact the concept of stopping times is relevant whenever one has an ordered sequence of sub σ-fields and is concerned about conditioning with respect to them.

Let $(\Omega, \mathcal{F})$ be a measurable space and $\{\mathcal{F}_t : t \in T\}$ be a family of sub σ-fields. T is an ordred set usually a set of real numbers or integers of the form $T = \{t : a \leq t \leq b\}$ or $\{t : t \geq a\}$. We will assume that $T = \{0, 1, 2, \cdots, \}$ the set of non negative integers. The family $\mathcal{F}_n$ is assumed to be increasing with n. In other words

$$\mathcal{F}_m \subset \mathcal{F}_n \quad \textit{if} \quad m < n$$

An $\mathcal{F}$ measurable random variable $\tau(\omega)$ mapping $\Omega \to \{0, 1, \cdots, \infty\}$ is said to be a stopping time if for every $n \geq 0$ the set $\{\omega : \tau(\omega) \leq n\} \in \mathcal{F}_n$. A stopping time may actually take the value ∞ on a nonempty subset of Ω.

The idea behind the definiton of a stopping time, as we saw in the study of Markov chains is that the decision to stop at time n can be based only on the information available upto that time.

Exercise 5.5. Show that the function $\tau(\omega) \equiv k$ is a stopping time for any admissible value of the constant k.

Exercise 5.6. Show that if τ is a stopping time and $f : T \to T$ is a nondecreasing function that satisfies $f(t) \geq t$ for all $t \in T$, then $\tau'(\omega) = f(\tau(\omega))$ is again a stopping time.

Exercise 5.7. Show that if τ_1, τ_2 are stopping times so are $\max(\tau_1, \tau_2)$ and $\min(\tau_1, \tau_2)$. In particular any stopping time τ is an *increasing* limit of bounded stopping times $\tau_n(\omega) = \min(\tau(\omega), n)$.

To every stopping time τ we associate a stopped σ-field $\mathcal{F}_\tau \subset \mathcal{F}$ defined by

$$\mathcal{F}_\tau = \{A : A \in \mathcal{F} \text{ and } A \cap \{\omega : \tau(\omega) \leq n\} \in \mathcal{F}_n \text{ for every } n\}. \tag{5.5}$$

This should be thought of as the information available upto the stopping time τ. In other words, events in $\mathcal{F}_\tau$ correspond to questions that can be answered with a yes or no, if we stop observing the process at time τ.

Exercise 5.8. Verify that for any stopping time τ, $\mathcal{F}_\tau$ is indeed a sub σ-field i.e. is closed under countable unions and complementations. If $\tau(\omega) \equiv k$ then $\mathcal{F}_\tau \equiv \mathcal{F}_k$. If $\tau_1 \leq \tau_2$ are stopping times $\mathcal{F}_{\tau_1} \subset \mathcal{F}_{\tau_2}$. Finally if τ is a stopping time then it is $\mathcal{F}_\tau$ measurable.

Exercise 5.9. If $X_n(\omega)$ is a sequence of measurable functions on $(\Omega, \mathcal{F})$ such that for every $n \in T$, X_n is $\mathcal{F}_n$ measurable then on the set $\{\omega : \tau(\omega) < \infty\}$, which is an $\mathcal{F}_\tau$ measurable set, the function $X_\tau(\omega) = X_{\tau(\omega)}(\omega)$ is $\mathcal{F}_\tau$ measurable.

The following theorem called Doob's optional stopping theorem is one of the central facts in the theory of martingale sequences.

Theorem 5.11. (Optional Stopping Theorem.) *Let $\{X_n : n \geq 0\}$ be sequence of random variables defined on a probability space $(\Omega, \mathcal{F}, P)$, which is a martingale sequence with respect to the filtration $(\Omega, \mathcal{F}_n, P)$ and $0 \leq \tau_1 \leq \tau_2 \leq C$ be two* **bounded** *stopping times. Then*

$$E[X_{\tau_2} \mid \mathcal{F}_{\tau_1}] = X_{\tau_1} \quad \text{a.e.}$$

Proof. Since $\mathcal{F}_{\tau_1} \subset \mathcal{F}_{\tau_2} \subset \mathcal{F}_C$, it is sufficient to show that for any martingale $\{X_n\}$

$$E\,[\,X_k\,|\mathcal{F}_\tau\,] = X_\tau \quad \text{a.e.} \tag{5.6}$$

provided τ is a stopping time bounded by the integer k. To see this we note that in view of Exercise 4.9,

$$E[X_k|\mathcal{F}_{\tau_1}] = E[E[X_k|\mathcal{F}_{\tau_2}]|\mathcal{F}_{\tau_1}]$$

and if (5.6) holds, then

$$E[X_{\tau_2}|\mathcal{F}_{\tau_1}] = X_{\tau_1} \quad \text{a.e.}$$

Let $A \in \mathcal{F}_\tau$. If we define $E_j = \{\omega : \tau(\omega) = j\}$, then $\Omega = \cup_1^k E_j$ is a disjoint union. Moreover $A \cap E_j \in \mathcal{F}_j$ for every $j = 1, \cdots, k$. By the martingale property

$$\int_{A\cap E_j} X_k\,dP = \int_{A\cap E_j} X_j\,dP = \int_{A\cap E_j} X_\tau\,dP$$

and summing over $j = 1, \cdots, k$ gives

$$\int_A X_k\,dP = \int_A X_\tau\,dP$$

for every $A \in \mathcal{F}_\tau$ and we are done. □

Remark 5.11. In particular if X_n is a martingale sequence and τ is a bounded stopping time then $E[X_\tau] = E[X_0]$. This property, obvious for constant times, has now been extended to bounded stopping times. In a 'fair' game, a policy to quit at an 'opportune' time, gives no advantage to the gambler so long as he or she cannot foresee the future.

Exercise 5.10. The property extends to sub or super-martingales. For example if X_n is a sub-martingale, then for any two bounded stopping times $\tau_1 \leq \tau_2$, we have

$$E\,[X_{\tau_2}\,|\mathcal{F}_{\tau_1}\,] \geq X_{\tau_1} \quad \text{a.e..}$$

One cannot use the earlier proof directly, but one can reduce it to the martingale case by applying the Doob decomposition theorem.

Exercise 5.11. Boundedness is important. Take $X_0 = 0$ and

$$X_n = \xi_1 + \xi_2 \cdots + \xi_n \quad \text{for} \quad n \geq 1$$

where ξ_i are independent identically distributed random variables taking the values ± 1 with probability $\frac{1}{2}$. Let $\tau = \inf\{n : X_n = 1\}$. Then τ is a stopping time, $P[\tau < \infty] = 1$, but τ is not bounded. $X_\tau = 1$ with probability 1 and trivially $E[X_\tau] = 1 \neq 0$.

Exercise 5.12. It does not mean that we can never consider stopping times that are unbounded. Let τ be an unbounded stopping time. For every k, $\tau_k = \min(\tau, k)$ is a bounded stopping time and $E[X_{\tau_k}] = 0$ for every k. As $k \to \infty$, $\tau_k \uparrow \tau$ and $X_{\tau_k} \to X_\tau$. If we can establish uniform integrability of X_{τ_k} we can pass to the limit. In particular if $S(\omega) = \sup_{0 \leq n \leq \tau(\omega)} |X_n(\omega)|$ is integrable then $\sup_k |X_{\tau_k}(\omega)| \leq S(\omega)$ and therefore $E[X_\tau] = 0$.

Exercise 5.13. Use a similar argument to show that if

$$S(\omega) = \sup_{0 \leq k \leq \tau_2(\omega)} |X_k(\omega)|$$

is integrable, then for any $\tau_1 \leq \tau_2$

$$E[X_{\tau_2} | \mathcal{F}_{\tau_1}] = X_{\tau_1} \quad \text{a.e..}$$

Exercise 5.14. The previous exercise needs the fact that if $\tau_n \uparrow \tau$ are stopping times, then

$$\sigma\{\cup_n \mathcal{F}_{\tau_n}\} = \mathcal{F}_\tau.$$

Prove it.

Exercise 5.15. Let us go back to the earlier exercise (Exercise 5.11) where we had

$$X_n = \xi_1 + \cdots + \xi_n$$

as a sum of n idependent random variables taking the values ± 1 with probability $\frac{1}{2}$. Show that if τ is a stopping time with $E[\tau] < \infty$, then $S(\omega) = \sup_{1 \leq n \leq \tau(\omega)} |X_n(\omega)|$ is square integrable and therefore $E[X_\tau] = 0$. [Hint: Use the fact that $X_n^2 - n$ is a martingale.]

5.5 Upcrossing Inequality.

The following inequality due to Doob, that controls the oscillations of a martingale sequence, is very useful for proving the almost sure convergence of L_1 bounded martingales directly. Let $\{X_j : 0 \leq j \leq n\}$ be a martingale sequence with $n+1$ terms. Let us take two real numbers $a < b$. An upcrossing is a pair of terms X_k and X_l, with $k < l$, for which $X_k \leq a < b \leq X_l$. Starting from X_0, we locate the first term that is at most a and then the first term following it that is at least b. This is the first upcrossing. In our martingale sequence there will be a certain number of completed upcrossings (of course over disjoint intervals) and then at the end we may be in the middle of an upcrossing or may not even have started on one because we are still on the way down from a level above b to one below a. In any case there will be a certain number $U(a, b)$ of completed upcrossings. Doob's upcrossing inequlity gives a uniform upper bound on the expected value of $U(a, b)$ in terms of $E[|X_n|]$, i.e. one that does not depend otherwise on n.

Theorem 5.12. Doob's upcrossing inequality *For any n,*

$$E[U(a,b)] \leq \frac{1}{b-a} E[a - X_n]^+ \leq \frac{1}{b-a}[|a| + E\big[|X_n|\big]] \tag{5.7}$$

Proof. Let us define recursively

$$\begin{aligned}
\tau_1 &= n \wedge \inf\{k : X_k \leq a\} \\
\tau_2 &= n \wedge \inf\{k : k \geq \tau_1, X_k \geq b\} \\
&\cdots\cdots\cdot \\
\tau_{2k} &= n \wedge \inf\{k : k \geq \tau_{2k-1}, X_k \geq b\} \\
\tau_{2k+1} &= n \wedge \inf\{k : k \geq \tau_{2k}, X_k \leq a\} \\
&\cdots\cdots\cdot
\end{aligned}$$

Since $\tau_k \geq \tau_{k-1} + 1$, $\tau_n = n$. Consider the quantity

$$D(\omega) = \sum_{j=1}^{n} [X_{\tau_{2j}} - X_{\tau_{2j-1}}]$$

which could very well have lots of 0's at the end. In any case the first few terms correspond to upcrossings and each term is at least $(b - a)$ and there are $U(a,b)$ of them. Before the 0's begin there may be at most one nonzero term which is an incomplete upcrossing, i.e. when $\tau_{2\ell-1} < n = \tau_{2\ell}$ for some ℓ. It is then equal to $(X_n - X_{\tau_{2l-1}}) \geq X_n - a$ for some l. If on the other hand if we end in the middle of a downcrossing, i.e. $\tau_{2\ell} < n = \tau_{2\ell+1}$ there is no incomplete upcrossing. Therefore

$$D(\omega) \geq (b-a)U(a,b) + R_n(\omega)$$

with the remainder $R_n(\omega)$ satisfying

$$\begin{aligned} R_n(\omega) &= 0 && \textit{if } \tau_{2\ell} < n = \tau_{2\ell+1} \\ &\geq (X_n - a) && \textit{if } \tau_{2\ell-1} < n = \tau_{2\ell} \end{aligned}$$

By the optional stopping theorem $E[D(\omega)] = 0$. This gives the bound

$$\begin{aligned} E[U(a,b)] &\leq \frac{1}{b-a} E\big[-R_n(\omega)\big] \leq \frac{1}{b-a} E\big[(a-X_n)^+\big]\big] \\ &\leq \frac{1}{b-a} E\big[|a-X_n|\big] \leq \frac{1}{b-a} E\big[|a| + |X_n|\big]. \end{aligned}$$

□

Remark 5.12. In particular if X_n is an L_1 bounded martingale, then the number of upcrossings of any interval $[a,b]$ is finite with Probability 1. From Doob's inequality, the sequence X_n is almost surely bounded. It now follows by taking a countable number of intervals $[a,b]$ with rational endpoints that X_n has a limit almost surely. If X_n is uniformly integrable then the convergence is in L_1 and then $X_n = E[X|\mathcal{F}_n]$. If we have a uniform L_p bound on X_n, then $X \in L_p$ and $X_n \to X$ in L_p. All of our earlier results on the convergence of martingales now follow.

Exercise 5.16. For the proof it is sufficient that we have a supermartingale. In fact we can change signs and so a submartingale works just as well.

5.6 Martingale Transforms, Option Pricing.

If X_n is a martingale with respect to $(\Omega, \mathcal{F}_n, P)$ and Y_n are the differences $X_n - X_{n-1}$, a martingale transform X'_n of X_n is given by the formula

$$X'_n = X'_{n-1} + a_{n-1} Y_n, \quad \text{for} \quad n \geq 1$$

where a_{n-1} is $\mathcal{F}_{n-1}$ measurable and has enough integrability assumptions to make $a_{n-1}Y_n$ integrable. An elementary calculation shows that

$$E\left[\,X_n'\,|\mathcal{F}_{n-1}\,\right] = X_{n-1}'$$

making X_n' a martingale as well. X_n' is called a martingale transform of X_n. The interpretation is if we have a fair game, we can choose the size and side of our bet at each stage based on the prior history and the game will continue to be fair. It is important to note that X_n may be sums of independent random vaiables with mean zero. But the independence of the increments may be destroyed and X_n' will in general no longer have the independent increments property.

Exercise 5.17. Suppose $X_n = \xi_1 + \cdots + \xi_n$, where ξ_j are independent random variables taking the values ± 1 with probability $\frac{1}{2}$. Let X_n' be the martingale transform given by

$$X_n' = \sum_{j=1}^{n} a_{j-1}(\omega)\xi_j$$

where a_j is $\mathcal{F}_j$ measurable, $\mathcal{F}_j$ being the σ-field generated by $\xi_1, \cdots, \xi_j$. Calculate $E\big\{[X_n']^2\big\}$.

Suppose X_n is a sequence of nonnegative random variables that represent the value of a security that is traded in the market place at a price that is X_n for day n and changes overnight between day n and day $n+1$ from X_n to X_{n+1}. We could at the end of day n, based on any information $\mathcal{F}_n$ that is available to us at the end of that day be either long or short on the security. The quantity $a_n(\omega)$ is the number of shares that we choose to own overnight between day n and day $n+1$ and that could be a function of all the information available to us up to that point. Positive values of a_n represent long positions and negative values represent short positions. Our gain or loss overnight is given by $a_n(X_{n+1} - X_n)$ and the cumulative gain(loss) is the transform

$$X_n' - X_0' = \sum_{j=1}^{n} a_{j-1}(X_j - X_{j-1}).$$

A contingent claim (European Option) is really a gamble or a bet based on the value of X_N at some terminal date N. The nature of the claim is that there is function $f(x)$ such that if the security trades on that day at a price

x then the claim pays an amount of $f(x)$. A **call** is an option to buy at a certain price a and the payoff is $f(x) = (x-a)^+$ whereas a **put** is an option to sell at a fixed price a and therefore has a payoff function $f(x) = (a-x)^+$.

Replicating a claim, if it is possible at all, is determining $a_0, a_1, \cdots, a_N$ and the initial value V_0 such that the transform

$$V_N = V_0 + \sum_{j=1}^{N} a_j(X_{j+1} - X_j)$$

at time N equals the claim $f(X_N)$ under every conceivable behavior of the price movements $X_1, X_2, \cdots, X_N$. If the claim can be exactly replicated starting from an initial capital of V_0, then V_0 becomes the price of that option. Anyone could sell the option at that price, use the proceeds as capital and follow the strategy dictated by the coefficients $a_0, \cdots, a_{N-1}$ and have **exactly** enough to pay off the claim at time N. Here we are ignoring transaction costs as well as interest rates. It is not always true that a claim can be replicated.

Let us assume for simplicity that the stock prices are always some nonnegative integral multiples of some unit. The set of possible prices can then be taken to be the set of nonnegative integers. Let us make a crucial assumption that if the price on some day is x the price on the next day is $x \pm 1$. It has to move up or down a notch. It cannot jump two or more steps or even stay the same. When the stock price hits 0 we assume that the company goes bankrupt and the stock stays at 0 for ever. In all other cases, from day to day, it always moves either up or down a notch.

Let us value the claim f for one period. If the price at day $N-1$ is $x \neq 0$ and we have assets c on hand and invest in a shares we will end up on day N, with either assets of $c+a$ and a claim of $f(x+1)$ or assets of $c-a$ with a claim of $f(x-1)$. In order to make sure that we break even in either case, we need

$$f(x+1) = c + a \; ; f(x-1) = c - a$$

and solving for a and c, we get

$$c(x) = \frac{1}{2}[f(x-1) + f(x+1)] \; ; a(x) = \frac{1}{2}[f(x+1) - f(x-1)]$$

The value of the claim with one day left is

$$V_{N-1}(x) = \begin{cases} \frac{1}{2}[f(x-1) + f(x+1)] & \textit{if } x \geq 1 \\ f(0) & \textit{if } x = 0 \end{cases}$$

and we can proceed by iteration

$$V_{j-1}(x) = \begin{cases} \frac{1}{2}[V_j(x-1) + V_j(x+1)] & \textit{if } x \geq 1 \\ V_j(0) & \textit{if } x = 0 \end{cases}$$

for $j \geq 1$ till we arrive at the value $V_0(x)$ of the claim at time 0 and price x. The corresponding values of $a = a_{j-1}(x) = \frac{1}{2}[V_j(x+1) - V_j(x-1)]$ gives us the number of shares to hold between day $j-1$ and j if the current price at time $j-1$ equals x.

Remark 5.13. The important fact is that the value is determined by arbitrage and is unaffected by the actual movement of the price so long as it is compatible with the model.

Remark 5.14. The value does not depend on any statistical assumptions on the various probabilities of transitions of price levels between successive days.

Remark 5.15. However the value can be interpreted as the expected value

$$V_0(x) = E^{P_x}\Big\{f(X_N)\Big\}$$

where P_x is the random walk starting at x with probability $\frac{1}{2}$ for transitions up or down a level, which is absorbed at 0.

Remark 5.16. P_x can be characterized as the unique probability distribution of $(X_0, \cdots, X_N)$ such that $P_x[X_0 = x] = 1$, $P_x[|X_j - X_{j-1}| = 1 | X_{j-1} \geq 1] = 1$ for $1 \leq j \leq N$ and X_j is a martingale with respect to $(\Omega, \mathcal{F}_j, P_x)$ where $\mathcal{F}_j$ is generated by $X_0, \cdots, X_j$.

Exercise 5.18. It is not necessary for the argument that the set of possible price levels be equally spaced. If we make the assumption that for each price level $x > 0$, the price on the following day can take only one of two possible values $h(x) > x$ and $l(x) < x$ with a possible bankruptcy if the level 0 is reached, a simlar analysis can be worked out. Carry it out.

5.7 Martingales and Markov Chains.

One of the ways of specifying the joint distribution of a sequence $X_0, \cdots, X_n$ of random variables is to specify the distribution of X_0 and for each $j \geq 1$, specify the conditional distribution of X_j given the σ-field $\mathcal{F}_{j-1}$ generated by $X_0, \cdots, X_{j-1}$. Equivalently instead of the conditional distributions one can specify the conditional expectations $E\left[f(X_j)|\mathcal{F}_{j-1}\right]$ for $1 \leq j \leq n$. Let us write

$$h_{j-1}(X_0, \cdots, X_{j-1}) = E\left[f(X_j)|\mathcal{F}_{j-1}\right] - f(X_{j-1})$$

so that, for $1 \leq j \leq n$

$$E\left[\{f(X_j) - f(X_{j-1}) - h_{j-1}(X_0, \cdots, X_{j-1})\}|\mathcal{F}_{j-1}\right] = 0$$

or

$$Z_j^f = f(X_j) - f(X_0) - \sum_{i=1}^{j} h_{i-1}(X_0, \cdots, X_{i-1})$$

is a martingale for every f. It is not difficult to see that the specification of $\{h_i\}$ for each f is enough to determine all the successive conditional expectations and therefore the conditional distributions. If in addition the initial distribution of X_0 is specified then the distribution of $X_0, \cdots, X_n$ is completely determined.

If for each j and f, the corresponding $h_{j-1}(X_0, \cdots, X_{j-1})$ is a function $h_{j-1}(X_{j-1})$ of X_{j-1} only, then the distribution of $(X_0, \cdots, X_n)$ is Markov and the transition probabilities are seen to be given by the relation

$$\begin{aligned} h_{j-1}(X_{j-1}) &= E\big[[f(X_j) - f(X_{j-1})]|\mathcal{F}_{j-1}\big] \\ &= \int [f(y) - f(X_{j-1})]\pi_{j-1,j}(X_{j-1}, dy). \end{aligned}$$

In the case of a stationary Markov chain the relationship is

$$\begin{aligned} h_{j-1}(X_{j-1}) = h(X_{j-1}) &= E\big[[f(X_j) - f(X_{j-1})]|\mathcal{F}_{j-1}\big] \\ &= \int [f(y) - f(X_{j-1})]\pi(X_{j-1}, dy). \end{aligned}$$

If we introduce the linear transformation (transition operator)

$$(\Pi\, f)(x) = \int f(y)\pi(x, dy) \tag{5.8}$$

then

$$h(x) = ([\Pi - I]f)(x).$$

Remark 5.17. In the case of a Markov chain on a countable state space

$$(\Pi f)(x) = \sum_y \pi(x,y) f(y)$$

and

$$h(x) = [\Pi - I](x) = \sum_y [f(y) - f(x)]\pi(x,y).$$

Remark 5.18. The measure P_x on the space $(\Omega, \mathcal{F})$ of sequences $\{x_j : j \geq 0\}$ from the state space X, that corresponds to the Markov Process with transition probability $\pi(x, dy)$, and initial state x, can be characterized as the unique measure on $(\Omega, \mathcal{F})$ such that

$$P_x \left\{ \omega : x_0 = x \right\} = 1$$

and for every bounded measurable function f defined on the state space X

$$f(x_n) - f(x_0) - \sum_{j=1}^{n} h(x_{j-1})$$

is a martingale with respect to $(\Omega, \mathcal{F}_n, P_x)$ where

$$h(x) = \int_X [f(y) - f(x)]\pi(x, dy).$$

Let $A \subset X$ be a measurable subset and let $\tau_A = \inf\{n \geq 0 : x_n \in A\}$ be the first entrance time into A. It is easy to see that τ_A is a stopping time. It need not always be true that $P_x\{\tau_A < \infty\} = 1$. But $U_A(x) = P_x\{\tau_A < \infty\}$ is a well defined measurable function of x, that satisfies $0 \leq U(x) \leq 1$ for all x and is the exit probability from the set A^c. By its very definition $U_A(x) \equiv 1$ on A and if $x \notin A$, by the Markov property,

$$U_A(x) = \pi(x, A) + \int_{A^c} U_A(y)\pi(x, dy) = \int_X U_A(y)\pi(x, dy).$$

In other words U_A satisfies $0 \leq U_A \leq 1$ and is a solution of

$$\begin{aligned}(\Pi - I)V &= 0 \text{ on } A^c \\ V &= 1 \text{ on } A \end{aligned} \tag{5.9}$$

Theorem 5.13. *Among all nonnegative solutions V of the equation (5.9) $U_A(x) = P_x\{\tau_A < \infty\}$ is the smallest. If $U_A(x) = 1$, then any bounded solution of the equation*

$$\begin{aligned}(\Pi - I)V &= 0 \text{ on } A^c \\ V &= f \text{ on } A \end{aligned} \tag{5.10}$$

is equal to

$$V(x) = E^{P_x}\big\{f(x_{\tau_A})\big\}. \tag{5.11}$$

In particular if $U_A(x) = 1$ for all $x \notin A$, then any bounded solution V of equation (5.10) is unique and is given by the formula (5.11).

Proof. First we establish that any nonnegative solution V of (5.10) dominates U_A. Let us replace V by $W = \min(V, 1)$. Then $0 \leq W \leq 1$ everywhere, $W(x) = 1$ for $x \in A$ and for $x \notin A$,

$$(\Pi W)(x) = \int_X W(y)\pi(x, dy) \leq \int_X V(y)\pi(x, dy) = V(x).$$

Since $\Pi W \leq 1$ as well we conclude that $\Pi W \leq W$ on A^c. On the otherhand it is obvious that $\Pi W \leq 1 = W$ on A. Since we have shown that $\Pi W \leq W$ everywhere it follows that $\{W(x_n)\}$ is a supermartingale with with repect to $(\Omega, \mathcal{F}_n, P_x)$. In particular for any bounded stopping time τ

$$E^{P_x}\big\{W(x_\tau)\big\} \leq E^{P_x}\big\{W(x_0)\big\} = W(x).$$

While we cannot take $\tau = \tau_A$ (since τ_A may not be bounded), we can always take $\tau = \tau_N = \min(\tau_A, N)$ to conclude

$$E^{P_x}\big\{W(x_{\tau_N})\big\} \leq E^{P_x}\big\{W(x_0)\big\} = W(x).$$

Let us let $N \to \infty$. On the set $\{\omega : \tau_A(\omega) < \infty\}$, $\tau_N \uparrow \tau_A$ and $W(x_{\tau_N}) \to W(x_{\tau_A}) = 1$. Since W is nonnegative and bounded,

$$\begin{aligned} W(x) &\geq \limsup_{N\to\infty} E^{P_x}\big\{W(x_{\tau_N})\big\} \geq \limsup_{N\to\infty} \int_{\tau_A<\infty} W(x_{\tau_N}) dP_x \\ &= P_x\{\tau_A < \infty\} = U_A(x). \end{aligned}$$

Since $V(x) \geq W(x)$ it follows that $V(x) \geq U_A(x)$.

For a bounded solution V of (5.10), let us define $h = (\Pi - I)V$ which will be a function vanishing on A^c. We know that

$$V(x_n) - V(x_0) - \sum_{j=1}^{n} h(x_{j-1})$$

is a martingale with rsepect to $(\Omega, \mathcal{F}_n, P_x)$ and let us use the stopping theorem with $\tau_N = \min(\tau_A, N)$. Since $h(x_{j-1}) = 0$ for $j \leq \tau_A$, we obtain

$$V(x) = E^{P_x}\big\{V(x_{\tau_N})\big\}.$$

If we now make the assumption that $U_A(x) = P_x\{\tau_A < \infty\} = 1$, let $N \to \infty$ and use the bounded convergence theorem it is easy to see that

$$V(x) = E^{P_x}\big\{f(x_{\tau_A})\big\}$$

which proves (5.11) and the rest of the theorem. □

Such arguments are powerful tools for the study of qualitative properties of Markov chains. Solutions to equations of the type $[\Pi - I]V = f$ are often easily constructed. They can be used to produce martingales, submartingales or supermartingales that have certain behavior and that in turn implies certain qualitative behavior of the Markov chain. We will now see several illustrations of this method.

Example 5.1. Consider the symmetric simple random walk in one dimension.

We know from recurrence that the random walk exits the interval $(-R, R)$ in a finite time. But we want to get some estimates on the exit time τ_R. Consider the function $u(x) = \cos \lambda x$. The function $f(x) = [\Pi u](x)$ can be calculated and

$$\begin{aligned} f(x) &= \frac{1}{2}[\cos \lambda(x-1) + \cos \lambda(x+1)] \\ &= \cos \lambda \, \cos \lambda x \\ &= \cos \lambda \, u(x). \end{aligned}$$

If $\lambda < \frac{\pi}{2R}$, then $\cos \lambda x \geq \cos \lambda R > 0$ in $[-R, R]$. Consider $Z_n = e^{\sigma n} \cos \lambda x_n$ with $\sigma = -\log \cos \lambda$.

$$E^{P_x}\{Z_n|\mathcal{F}_{n-1}\} = e^{\sigma\, n} f(x_{n-1}) = e^{\sigma\, n} \cos \lambda \cos \lambda x_{n-1} = Z_{n-1}.$$

If τ_R is the exit time from the interval $(-R, R)$, for any N, we have

$$E^{P_x}\{Z_{\tau_R \wedge N}\} = E^{P_x}\{Z_0\} = \cos \lambda x.$$

Since $\sigma > 0$ and $\cos \lambda x \geq \cos \lambda R > 0$ for $x \in [-R, R]$, if R is an integer, we can claim that

$$E^{P_0}\{e^{\sigma\,[\tau_R \wedge N]}\} \leq \frac{\cos \lambda x}{\cos \lambda R}.$$

Since the estimate is uniform we can let $N \to \infty$ to get the estimate

$$E^{P_0}\{e^{\sigma\,\tau_R}\} \leq \frac{\cos \lambda x}{\cos \lambda R}.$$

Exercise 5.19. Can you prove equality above? What is range of validity of the equality? Is $E^{P_x}\{e^{\sigma \tau_R}\} < \infty$ for all $\sigma > 0$?

Example 5.2. Let us make life slightly more complicated by taking a Markov chain in Z^d with transition probabilities

$$\pi(x, y) = \begin{cases} \frac{1}{2d} + \delta(x, y) & \textit{if } |x - y| = 1 \\ 0 & \textit{if } |x - y| \neq 0 \end{cases}$$

so that we have slightly perturbed the random walk with perhaps even a possible bias.

Exact calculations like in Eaxmple 5.1 are of course no longer possible. Let us try to estimate again the exit time from a ball of radius R. For $\sigma > 0$ consider the function

$$F(x) = \exp[\sigma \sum_{i=1}^{d} |x_i|]$$

defined on Z^d. We can get an estimate of the form

$$(\Pi F)(x_1, \cdots, x_d) \geq \theta F(x_1, \cdots, x_d)$$

for some choices of $\sigma > 0$ and $\theta > 1$ that may depend on R. Now proceed as in Example 5.1.

Example 5.3. We can use these methods to show that the random walk is transient in dimension $d \geq 3$.

For $0 < \alpha < d-2$ consider the function $V(x) = \frac{1}{|x|^\alpha}$ for $x \neq 0$ with $V(0) = 1$. An approximate calculation of $(\Pi V)(x)$ yields, for sufficiently large $|x|$ (i.e $|x| \geq L$ for some L), the estimate

$$(\Pi V)(x) - V(x) \leq 0$$

If we start initially from an x with $|x| > L$ and take τ_L to be the first entrance time into the ball of radius L, one gets by the stopping theorem, the inequality

$$E^{P_x}\big\{V(x_{\tau_L \wedge N})\big\} \leq V(x).$$

If $\tau_L \leq N$, then $|x_{\tau_L}| \leq L$. In any case $V(x_{\tau_L \wedge N}) \geq 0$. Therefore,

$$P_x\big\{\tau_L \leq N\big\} \leq \frac{V(x)}{\inf_{|y| \leq L} V(y)}$$

valid uniformly in N. Letting $N \to \infty$

$$P_x\big\{\tau_L < \infty\big\} \leq \frac{V(x)}{\inf_{|y| \leq L} V(y)} \ .$$

If we let $|x| \to \infty$, keeping L fixed, we see the transience. Note that recurrence implies that $P_x\big\{\tau_L < \infty\big\} = 1$ for all x. The proof of transience really only required a function V defined for large $|x|$, that was strictly positive for each x, went to 0 as $|x| \to \infty$ and had the property $(\Pi V)(x) \leq V(x)$ for large values of $|x|$.

Example 5.4. We will now show that the random walk is recurrent in $d = 2$.

This is harder because the recurrence of random walk in $d = 2$ is right on the border. We want to construct a function $V(x) \to \infty$ as $|x| \to \infty$ that satisfies $(\Pi V)(x) \leq V(x)$ for large $|x|$. If we succeed, then we can estimate by a stopping argument the probability that the chain starting from a point x in the annulus $\ell < |x| < L$ exits at the outer circle before getting inside the inner circle.

$$P_x\big\{\tau_L < \tau_\ell\big\} \leq \frac{V(x)}{\inf_{|y| \geq L} V(y)} \ .$$

We also have for every L,

$$P_x\big\{\tau_L < \infty\big\} = 1.$$

This proves that $P_x\big\{\tau_\ell < \infty\big\} = 1$ thereby proving recurrence. The natural candidate is $F(x) = \log|x|$ for $x \neq 0$. A computation yields

$$(\Pi F)(x) - F(x) \le \frac{C}{|x|^4}$$

which does not quite make it. On the other hand if $U(x) = |x|^{-1}$, for large values of $|x|$,

$$(\Pi U)(x) - U(x) \ge \frac{c}{|x|^3}$$

for some $c > 0$. The choice of $V(x) = F(x) - CU(x) = \log x - \frac{C}{|x|}$ works with any $C > 0$.

Example 5.5. We can use these methods for proving positive recurrence as well.

Suppose X is a countable set and we can find $V \ge 0$, a finite set F and a constant $C \ge 0$ such that

$$(\Pi V)(x) - V(x) \le \begin{cases} -1 & \textit{for } x \notin F \\ C & \textit{for } x \in F \end{cases}$$

Let us let $U = \Pi V - V$, and we have

$$\begin{aligned}
-V(x) &\le E^{P_x}\big\{V(x_n) - V(x)\big\} \\
&= E^{P_x}\big\{\sum_{j=1}^{n} U(x_{j-1})\big\} \\
&\le E^{P_x}\big\{\sum_{j=1}^{n} C\,\mathbf{1}_F(x_{j-1}) - \sum_{j=1}^{n} \mathbf{1}_{F^c}(x_{j-1}) \\
&= -E^{P_x}\big\{\sum_{j=1}^{n}[1 - (1+C)\mathbf{1}_F(x_{j-1})]\big\} \\
&= -n + (1+C)\sum_{j=1}^{n}\sum_{y\in F}\pi^n(x,y) \\
&= -n + o(n) \quad \textit{as} \quad n \to \infty.
\end{aligned}$$

if the process is not positive recurrent. This is a contradiction.

For instance if $X = Z$, the integers, and we have a little bit of bias towards the origin in the random walk

$$\pi(x, x+1) - \pi(x, x-1) \geq \frac{a}{|x|} \quad \textit{if} \quad x \leq -\ell$$
$$\pi(x, x-1) - \pi(x, x+1) \geq \frac{a}{|x|} \quad \textit{if} \quad x \geq \ell$$

with $V(x) = x^2$, for $x \geq \ell$

$$\begin{aligned}(\Pi V)(x) &\leq (x+1)^2\frac{1}{2}(1 - \frac{a}{|x|}) + (x-1)^2\frac{1}{2}(1 + \frac{a}{|x|}) \\ &= x^2 + 1 - 2a\end{aligned}$$

If $a > \frac{1}{2}$, we can multiply V by a constant and it works.

Exercise 5.20. What happens when

$$\pi(x, x+1) - \pi(x, x-1) = -\frac{1}{2x}$$

for $|x| \geq 10$? (See Exercise 4.16)

Example 5.6. Let us return to our example of a branching process Example 4.4. We see from the relation

$$E[X_{n+1}|\mathcal{F}_n] = mX_n$$

that $\frac{X_n}{m^n}$ is a martingale. If $m < 1$ we saw before quite easily that the population becomes extinct. If $m = 1$, X_n is a martingale. Since it is nonnegative it is L_1 bounded and must have an almost sure limit as $n \to \infty$. Since the population is an integer, this means that the size eventually stabilizes. The limit can only be 0 because the population cannot stabilize at any other size. If $m > 1$ there is a probability $0 < q < 1$ such that $P[X_n \to 0|X_0 = 1] = q$, We can show that with probability $1 - q$, $X_n \to \infty$. To see this consider the function $u(x) = q^x$. In the notation of Example 4.4

$$\begin{aligned}E[q^{X_{n+1}}|\mathcal{F}_n] &= [\sum q^j p_j]^{X_n} \\ &= [P(q)]^{X_n} \\ &= q^{X_n}\end{aligned}$$

so that q^{X_n} is a non negative martingale. It then has an almost sure limit, which can only be 0 or 1. If q is the probabbility that it is 1 i.e that $X_n \to 0$, then $1 - q$ is the probability that it is 0, i.e. that $X_n \to \infty$.