Basic Stochastic Processes

Benjamin Arora

Department of Mathematics,
University of Connecticut (UConn)

Directed Reading Program Presentations
December 9th, 2020

Outline

Outline

σ-fields

Definition

Given a non-empty set Ω, we define a *σ-field* $\mathcal{F}$ on Ω as a collection of subsets of Ω such that

- $\emptyset \in \mathcal{F}$;
- If $A \in \mathcal{F}$, then $A^c \in \mathcal{F}$;
- If $A_1, A_2, \ldots$ are in $\mathcal{F}$, then $A_1 \cup A_2 \cup \cdots$ is in $\mathcal{F}$.

Some commonly used σ-fields are $\mathcal{P}(\mathbb{N})$ for discrete random variables and $\mathcal{B}(\mathbb{R})$, the family of *Borel sets* on $\mathbb{R}$, for continuous set-ups.

Probability Measures

Definition

A *probability measure* P on a σ-field $\mathcal{F}$ is a function $P : \mathcal{F} \to [0, 1]$ such that

- $P(\Omega) = 1$
- If $A_1, A_2, \ldots$ are pairwise disjoint sets, then

$$P(A_1 \cup A_2 \cup \cdots) = P(A_1) + P(A_2) + \cdots$$

These are the axioms of probability presented in Math 3160.

Probability Spaces

Definition

A *probability space* is a triple $(\Omega, \mathcal{F}, P)$ such that:

- Ω is any non-empty set
- $\mathcal{F}$ is a σ-field on Ω
- P is a probability measure

Basic Probability Exercises

Lemma (Borel-Cantelli)

Let $A_1, A_2, \ldots$ be a sequence of events such that $P(A_1) + P(A_2) + \cdots < \infty$ and let $B_n = A_n \cup A_{n+1} \cup \cdots$. Then

$$P(B_1 \cap B_2 \cap \cdots) = 0$$

.

Proof.

We know $\lim_{n\to\infty} P(A_n) = 0$, and that $B_1, B_2, \ldots$ is a contracting series of events, so that $P(B_1 \cap B_2 \cap \cdots) = \lim_{n\to\infty} P(B_n)$. Since $P(A_n \cup A_{n+1} \cup \cdots) \leq \sum_{i=n}^{\infty} P(A_i)$ for all n, it follows

$$\lim_{n\to\infty} P(B_n) = \lim_{n\to\infty} P(A_n \cup A_{n+1} \cup \cdots) \leq \lim_{n\to\infty} \sum_{i=n}^{\infty} P(A_i) = 0.$$

Random Variables

In Math 3160, we often defined random variables by their densities. For example, we defined a Gaussian random variable $X \sim N(\mu, \sigma^2)$ by the density function

$$f_X(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}.$$

Here, we will take a more abstract approach.

Random Variables

Definition

Given a σ-field $\mathcal{F}$ on Ω, a function $X : \Omega \to \mathbb{R}$ is $\mathcal{F}$*-measurable* if

$$\forall B \in \mathcal{B}(\mathbb{R}), \quad X^{-1}(B) \in \mathcal{F}.$$

We often write $\{X \in B\}$ in place of $X^{-1}(B)$.

Definition

If $(\Omega, \mathcal{F}, P)$ form a probability space, then X is called a *random variable*.

Random Variables

Example (Bernoulli Random Variable)

We define a Bernoulli random variable $X : \Omega \to \{0, 1\}$ by

$$P(X = 1) = p; \quad P(X = 0) = 1 - p.$$

Useful Ideas about Random Variables

Definition

The *σ-field $\sigma(X)$ generated by a random variable X* is σ-field generated by the sets of the form $\{X \in B\}$ for all $B \in \mathcal{B}(\mathbb{R})$.

Lemma (Doob-Dynkin)

If X is a random variable, then every $\sigma(X)$-measurable random variable Y can be written $Y = f(X)$ for some Borel function $f : \mathbb{R} \to \mathbb{R}$.

Expectation

In 3160, we defined the expectation $E(X)$ of a random variable X with density f as

$$E(X) = \int_{-\infty}^{\infty} xf(x)\,dx$$

provided this integral is absolutely convergent. With our more abstract definition of a random variable, we will be able to emulate this definition:

Definition

The *expectation* of a random variable X is defined

$$E(X) = \int_{\Omega} X\,dP$$

provided X is integrable.

Expectation

These definitions don't *quite* look the same. We will attempt to recover the formula from 3160:

We start by "pushing" our measure P to the real line, which gives us

$$E(X) = \int_{\Omega} (X(\omega))\, dP(\omega) = \int_{\mathbb{R}} x\, dP_X(x).$$

With a continuous random variable with density f_X,

$$P_X([x_i, x_{i+1}]) = \int_{x_i}^{x_{i+1}} f_X(x)\, dx.$$

Expectation

If we "differentiate" this expression (using the Radon-Nikodym derivative), we get

$$dP_X(x) = f_X(x)\,dx.$$

So, we have

$$E(X) = \int_{\mathbb{R}} x\,dP_X(x) = \int_{\mathbb{R}} x f_X(x)\,dx.$$

□

Outline

Conditional Expectation on a Discrete RV

Definition

Given a random variable X and a discrete random variable Y, the *conditional expectation of X given Y* is a random variable $E(X|Y)$ such that

$$E(X|Y)(\omega) = E(X|\{Y = y_n\}) \text{ if } Y(\omega) = y_n$$

for $n = 0, 1, 2, \ldots$.

Conditional Expectation on a Discrete RV

Example

Say we flip three coins, worth 5c, 10c, and 25c, respectively. Let X denote the total value of the face-up coins. We want to find the conditional expectation $E(X|Y)$ given the amount Y shown by just the 5c and 10c coins.

We see that Y takes on the values 0c, 5c, 10c, and 15c. So, we find each of the following

$$E(X|\{Y=0\}),\ E(X|\{Y=5\}),\ E(X|\{Y=10\}),\ E(X|\{Y=15\}).$$

Note for any event B,

$$E(X|B) = \frac{1}{P(B)} \int_B X \, dP.$$

Conditional Expectation on a Discrete RV

So, we get the following result for $E(X|Y)$:

$$E(X|Y)(\omega) = \begin{cases} 12.5c & \text{if } Y(\omega) = 0 \\ 17.5c & \text{if } Y(\omega) = 5 \\ 22.5c & \text{if } Y(\omega) = 10 \\ 27.5c & \text{if } Y(\omega) = 15. \end{cases}$$

Conditional Expectation on a Random Variable

Definition

Given an integrable random variable X and an *arbitrary* random variable Y, the *conditional expectation of X given Y* is a random variable $E(X|Y)$ that satisfies the following properties:

1. $E(X|Y)$ is $\sigma(Y)$-measurable
2. For any $A \in \sigma(Y)$,

$$\int_A E(X|Y)\, dP = \int_A X\, dP.$$

Remark

Notice, the definition of $E(X|Y)$ in both of these cases does not really depend on Y, rather, on $\sigma(Y)$. This leads us to the following, more general definition:

Conditional Expectation on a σ-field

Definition

Given an integrable random variable X on a probability space $(\Omega, \mathcal{F}, P)$ and a σ-field $\mathcal{G} \subset \mathcal{F}$, the *conditional expectation of* X *given* $\mathcal{G}$ is a random variable $E(X|\mathcal{G})$ that satisfies the following properties:

1. $E(X|\mathcal{G})$ is $\mathcal{G}$-measurable
2. For any $A \in \mathcal{G}$,
$$\int_A E(X|\mathcal{G})\, dP = \int_A X\, dP.$$

Conditional Expectation on a σ-field

We will state these useful properties without proof:

Proposition

Given integrable random variables X and Y and a σ-field $\mathcal{G}$,

1. $E(aX + bY|\mathcal{G}) = aE(X|\mathcal{G}) + bE(Y|\mathcal{G})$;
2. $E(E(X|\mathcal{G})) = E(X)$;
3. $E(XY|\mathcal{G}) = XE(Y|\mathcal{G})$ *if* X *is* $\mathcal{G}$*-measurable;*
4. $E(X|\mathcal{G}) = E(X)$ *if* X *is independent of* $\mathcal{G}$;
5. $E(E(X|\mathcal{G})|\mathcal{H}) = E(X|\mathcal{H})$ *if* $\mathcal{H} \subset \mathcal{G}$;
6. *if* $X \geq 0$, *then* $E(X|\mathcal{G}) \geq 0$.

Moving on...

These definitions and properties will be central as we move into our discussion of *martingales*.

Outline

Martingale Basics

Definition

Given a probability space $(\Omega, \mathcal{F}, P)$, a *filtration* is a family of σ-fields $\{\mathcal{F}_n\}$ such that

$$\mathcal{F}_1 \subset \mathcal{F}_2 \subset \cdots \subset \mathcal{F}.$$

We can think of $\mathcal{F}_n$ as the set of events we can measure up to time n.

Definition

A sequence of random variables $X_1, X_2, \ldots$ is *adapted* to a filtration $\mathcal{F}_1, \mathcal{F}_2, \ldots$ if X_n is $\mathcal{F}_n$-measurable.

Martingale Basics

Definition

A *martingale* on $(\Omega, \mathcal{F}, \mathcal{F}_n, P)$ is a collection of random variables $\{X_n\}$ such that

1. X_n is integrable for all n;
2. $X_1, X_2, \ldots$ is adapted to $\mathcal{F}_1, \mathcal{F}_2, \ldots$;
3. $E(X_{n+1}|\mathcal{F}_n) = X_n$

This last point tells us the following: what will happen at time $t = n + 1$ given all that has happened at $t = 0, \ldots, n$ only depends on time $t = n$.

Some Interesting Things about Martingales

Proposition

Given a martingale $X_1, X_2, \ldots$ adapted to a filtration $\mathcal{F}_1, \mathcal{F}_2, \ldots,$

$$E(X_{n+1}|\mathcal{F}_n) = E(X_{n+1}|X_n).$$

Some Interesting Things about Martingales

Proof.

Since $\sigma(X_n) \subseteq \mathcal{F}_n$ for all n, we have

$$E(E(X_{n+1}|\mathcal{F}_n)|X_n) = E(X_{n+1}|X_n)$$

by the tower property. Also,

$$E(E(X_{n+1}|\mathcal{F}_n)|X_n) = E(X_n|X_n) = X_n$$

because we can take out what is known. So,

$$E(X_{n+1}|\mathcal{F}_n) = X_n = E(X_{n+1}|X_n).$$

□

Some Interesting Things about Martingales

Proposition

If X_n is a martingale with respect to $\mathcal{F}_n$, then

$$E(X_1) = E(X_2) = \cdots .$$

Some Interesting Things about Martingales

Proof.

For any n, we have

$$E(E(X_{n+1}|\mathcal{F}_n)) = E(X_{n+1})$$

by the properties of conditional expectation. Also,

$$E(E(X_{n+1}|\mathcal{F}_n)) = E(X_n)$$

as X_n is a martingale with respect to $\mathcal{F}_n$. So,

$$E(X_n) = E(X_{n+1})$$

for all n. □

Examples with Martingales

Problem

Let X_n be a symmetric random walk, that is,

$$X_n = Y_1 + Y_2 + \cdots + Y_n,$$

where $Y_1, Y_2, \ldots$ is a sequence of of independent identically distributed random variables such that

$$P\{Y_n = 1\} = P\{Y_n = -1\} = \frac{1}{2}.$$

Show that $X_n^2 - n$ is a martingale with respect to the filtration

$$\mathcal{F}_n = \sigma(Y_1, \ldots, Y_n).$$

Examples with Martingales

Proof.

It first needs to be shown that $X_n^2 - n$ is integrable for all n. Since

$$|X_n| = |Y_1 + Y_2 + \cdots + Y_n| \leq |Y_1| + |Y_2| + \cdots + |Y_n| < n,$$

we know that $E(|X_n^2 - n|) \leq n^2 + n < \infty$.

Now, since $X_n^2 - n = (Y_1 + Y_2 + \cdots + Y_n)^2 - n$, by the Doob-Dynkin Lemma $X_n^2 - n$ is $\sigma(Y_1, Y_2, \ldots, Y_n)$-measurable as it is a function of $Y_1, Y_2, \ldots, Y_n$. □

Examples with Martingales

We finally want to show

$$E(X_{n+1}^2 - (n+1)|\mathcal{F}_n) = X_n^2 - n.$$

Writing the term X_{n+1} as $X_n + Y_{n+1}$, we get

$$X_{n+1}^2 - (n+1) = X_n^2 + 2X_nY_{n+1} + Y_{n+1}^2 - (n+1), \quad \text{so}$$

$$E(X_{n+1}^2 - (n+1)|\mathcal{F}_n) = E(X_n^2 + 2X_nY_{n+1} + Y_{n+1}^2 - (n+1)|\mathcal{F}_n)$$

$$= E(X_n^2|\mathcal{F}_n) + 2E(X_nY_{n+1}|\mathcal{F}_n) + E(Y_{n+1}^2|\mathcal{F}_n) - E(n+1|\mathcal{F}_n)$$

$$= X_n^2 \cdot E(1|\mathcal{F}_n) + 2X_nE(Y_{n+1}|\mathcal{F}_n) + E(Y_{n+1}^2|\mathcal{F}_n) - (n+1)$$

$$= X_n^2 + 2X_nE(Y_{n+1}) + E(Y_{n+1}^2) - (n+1)$$

Examples with Martingales

A quick calculation shows $E(Y_{n+1}) = 0$ and $E(Y_{n+1}^2) = 1$, so we have

$$\begin{aligned} E(X_{n+1}^2 - (n+1)|\mathcal{F}_n) &= X_n^2 + 2X_n \cdot 0 + 1 - (n+1) \\ &= X_n^2 - n. \end{aligned}$$

Examples with Martingales

Theorem (Optional Stopping Theorem)

If X_n is a martingale and τ a stopping time such that

1. $\tau < \infty$
2. X_n *is integrable*
3. $E(X_n 1_{\tau > n}) \to 0$ *as* $n \to \infty$

then $E(X_\tau) = E(X_1)$.

Problem

Let X_n be a random walk as in the previous example, and let K be a positive integer. We define the first hitting time to be

$$\tau = \min\{n : |X_n| = K\}.$$

Then $E(\tau) = K^2$.

Examples with Martingales

In the previous example, we showed $X_n^2 - n$ is a martingale. The first two conditions of the Optional Stopping Theorem hold. The proof of the third condition is too lengthy to include here, but it holds as well.
So, we can say

$$E(X_\tau^2 - \tau) = E(X_1^2 - 1) = E(Y_1^2 - 1) = 0.$$

By the linearity of expectation,

$$E(X_\tau^2) = E(\tau).$$

Since $X_\tau = K$, we see $E(\tau) = E(K^2) = K^2$.
This tells us something interesting about gambling with finite capital- you can't beat the house!

Examples with Martingales

I wrote a script to verify this result empirically; Here is the plot of the results:

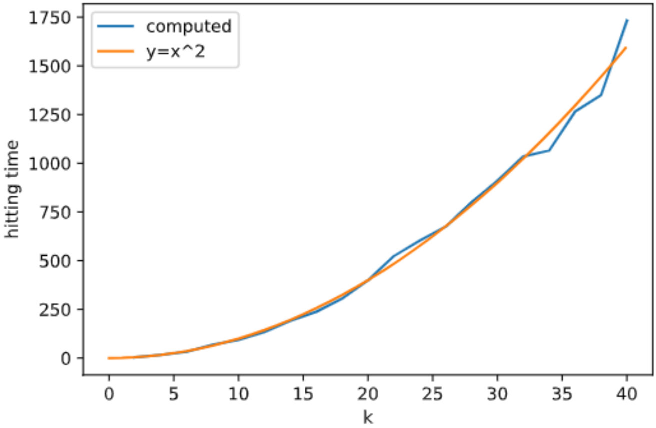

Examples with Martingales

Here is one sample path:

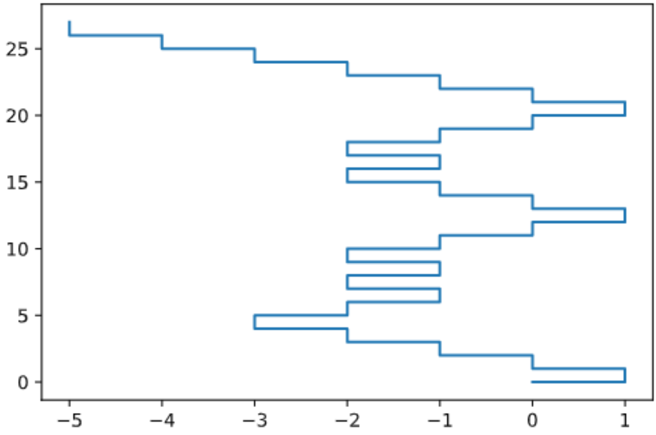

A Gambling Strategy

Problem

This gambling strategy is called "the martingale". Suppose we flip a coin and denote the outcomes $h_1, h_2, \ldots,$ *where* h_n *can take on the values* $+1$ *for heads and* -1 *for tails. We start by betting $1 on heads. If* h_n *is heads, we quit. Otherwise, we double our bet and keep playing.*

So, if we let a_n *denote our bet for flip* n*, we have the following strategy:*

$$a_n = \begin{cases} 2^{n-1}, & \textit{if} \;\; h_1 = h_2 = \cdots = h_{n-1} = -1 \\ 0, & \textit{otherwise.} \end{cases}$$

A Gambling Strategy

Now, if we let $G_n = h_1 + 2h_2 + \cdots + 2^{n-1}h_n$ be our winnings at time n, and $\tau = \min\{n : h_n = +1\}$, then we have the following properties:

- $G_{\tau \wedge n}$ is a martingale with respect to $\sigma(h_1, h_2, \cdots, h_n)$
- $P\{\tau < \infty\} = 1$
- $G_\tau = 1$

… *if* we have infinite time and are allowed to go infinitely into debt!

A Gambling Strategy

Proposition

If a gambler plays "the martingale" their expected loss right before the win is infinite:

$$E(G_{\tau-1}) = -\infty.$$

A Gambling Strategy

Proof.

Clearly, $P\{\tau = n\} = \frac{1}{2^n}$. Also,

$$\begin{aligned} G_{\tau-1} &= h_1 + 2h_2 + \cdots + 2^{\tau-2} h_{\tau-1} \\ &= -1(1 + 2 + \cdots + 2^{\tau-2}) \\ &= 1 - 2^{\tau-1}. \end{aligned}$$

So, we have

$$\begin{aligned} E(G_{\tau-1}) &= \sum_n (1 - 2^{n-1}) \cdot \frac{1}{2^n} \\ &= \sum_n \left(\frac{1}{2^n} - \frac{1}{2} \right) = -\infty. \end{aligned}$$

□

References

Z. Brzeiniak, T. Zastawniak, *Basic Stochastic Processes*, Springer, 2000