

Topics:

- Probability spaces
- Discrete random variables
- Conditional expectations

Probability spaces

A *random experiment* is an experiment with a *a priori* unknown outcome, and which can be repeated under identical conditions.

Example: *Toss a coin twice and report the two ordered outcomes.*

A *sample space* is the set of possible outcomes of the experiment, denoted by Ω . Individual elements (the outcomes) are denoted by ω . A sample space is *finite* if $\Omega = \{\omega_i, i = 1, \dots, n\}$ for some integer n , and *countable* if $\Omega = \{\omega_i, i = 1, 2, \dots\}$; otherwise it is uncountable.

Example: *The sample space of the previous example is $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\} = \{HH, HT, TH, TT\}$, where, e.g., HT means “first toss heads (H), second toss tails (T)”.*

A σ -field or σ -algebra, often denoted by \mathcal{F} , is a collection of subsets of Ω (*events*), such that if the sets $\{A_1, A_2, \dots\}$ are in \mathcal{F} , then also their complements $\{A_1^c, A_2^c, \dots\}$, their intersection $A_1 \cap A_2 \cap \dots$, and their union $A_1 \cup A_2 \cup \dots$ are in \mathcal{F} . A σ -algebra always contains the null set \emptyset (the “impossible event”), and Ω (the “sure event”).

Example: *One possible σ -field in the example is*

$$\begin{aligned} \mathcal{F} = & \{\emptyset, \{\omega_1\}, \{\omega_2\}, \{\omega_3\}, \{\omega_4\}, \\ & \{\omega_1, \omega_2\}, \{\omega_1, \omega_3\}, \{\omega_1, \omega_4\}, \{\omega_2, \omega_3\}, \{\omega_2, \omega_4\}, \{\omega_3, \omega_4\}, \\ & \{\omega_1, \omega_2, \omega_3\}, \{\omega_1, \omega_2, \omega_4\}, \{\omega_1, \omega_3, \omega_4\}, \{\omega_2, \omega_3, \omega_4\}, \Omega\} \end{aligned}$$

This contains all possible events. For example, the set $\{\omega_1, \omega_4\}$ signifies the event “both tosses have the same outcome”. The smaller collection

$$\mathcal{F}_1 = \{\emptyset, \{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}, \Omega\}$$

is also a σ -field. Here the only interesting events are “first toss is heads” or “first toss is tails”.

A *probability measure* \mathbb{P} is a function which gives a number between 0 and 1 to all events A in \mathcal{F} , with the additional properties $\mathbb{P}(\emptyset) = 0$, $\mathbb{P}(\Omega) = 1$, $\mathbb{P}(A^c) = 1 - \mathbb{P}(A)$, and for disjoint sets $\{A_i\}$,

$$\mathbb{P}\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mathbb{P}(A_i).$$

Example: *The most natural probability measure in our example satisfies $\mathbb{P}(\text{HH}) = p^2$, $\mathbb{P}(\text{HT}) = \mathbb{P}(\text{TH}) = p(1 - p)$, $\mathbb{P}(\text{TT}) = (1 - p)^2$ for some number $0 < p < 1$. If the coin is fair, then $p = 0.5$.*

The triple $(\Omega, \mathcal{F}, \mathbb{P})$ is a *probability space*. It fully characterizes the random experiment.

Discrete random variables

A random variable X a real-valued function on Ω . Thus, we assign to each outcome a real number. The function should be such that events of the form $\{\omega : X(\omega) \in B\}$ are in \mathcal{F} , for all sets B which are intervals of \mathbb{R} or can be constructed from unions and intersections of intervals; this defines the Borel σ -field \mathcal{B} consisting of Borel sets B .

When events are in \mathcal{F} we say that they are *measurable*, i.e., they can be given a probability measure. The function X satisfying the property $\{\omega : X(\omega) \in B\} \in \mathcal{F}$ for all $B \in \mathcal{B}$ is called *Borel-measurable*. Thus all events $\{X \in B\}$ can be given a probability measure.

The function X transforms the original probability space $(\Omega, \mathcal{F}, \mathbb{P})$ into a new probability space $(\mathbb{R}, \mathcal{B}, P_X)$, where P_X is the *distribution* of X , defined by $P_X(B) = \mathbb{P}(X(\omega) \in B)$. The *distribution function* F_X is defined by $F_X(x) = P_X((-\infty, x]) = \mathbb{P}(X(\omega) \leq x)$.

Example: Two possible random variables on the previous space are, for some number $a > 0$,

$$X_1(\omega) = \begin{cases} a, & \text{if } \omega \in \{\text{HH}, \text{HT}\}, \\ -a, & \text{if } \omega \in \{\text{TH}, \text{TT}\}, \end{cases}$$

and

$$X_2(\omega) = \begin{cases} a, & \text{if } \omega \in \{\text{HH}, \text{TH}\}, \\ -a, & \text{if } \omega \in \{\text{HT}, \text{TT}\}. \end{cases}$$

Note that X_1 depends on the outcome of the first toss, and X_2 on the second toss. From this example one can construct a multiplicative binomial tree with $u = 1/d = \exp(a)$, by choosing a fixed S_0 and defining $S_1(\omega) = S_0 \exp(X_1(\omega))$ and $S_2(\omega) = S_1(\omega) \exp(X_2(\omega)) = S_0 \exp(X_1(\omega) + X_2(\omega))$. The distribution function of both X_1 and X_2 is

$$F_X(x) = \begin{cases} 0, & \text{if } x < -a, \\ p, & \text{if } -a \leq x < a, \\ 1, & \text{if } x \geq a. \end{cases}$$

A random vector is simply a vector-valued function of Ω . Borel-measurability can be defined analogously. In the above example, $\mathbf{X}(\omega) = (X_1(\omega), X_2(\omega))$ is a random vector.

Expectations

Consider a discrete random variable X , defined on a countable probability space $\Omega = \{\omega_1, \omega_2, \dots\}$. Let $x_i = X(\omega_i)$. Then the *mean* or *expected value* is simply given by

$$\begin{aligned}\mathbb{E}(X) &= \sum_{i=1}^{\infty} X(\omega_i)\mathbb{P}(\omega_i) \\ &= \sum_{i=1}^{\infty} x_i P_X(x_i).\end{aligned}$$

The two sums may alternatively be denoted as $\int_{\Omega} X(\omega)d\mathbb{P}(\omega)$, or $\int_{\mathbb{R}} x dP_X(x)$. For more general (continuous) random variables, the definition of these integrals is more complicated. For the present case, they may simply be seen as an alternative way to denote the sums.

Sometimes we need to make the explicit the relevant probability measure, leading to the notation $\mathbb{E}_{\mathbb{P}}$, $\mathbb{E}_{\mathbb{Q}}$, etc.

We may also take expectations of functions of X . Of particular importance is the *variance* of X :

$$\text{Var}(X) = \mathbb{E}\{(X - \mathbb{E}[X])^2\} = \mathbb{E}(X^2) - \{\mathbb{E}(X)\}^2.$$

Example: *In the example, we have*

$$\begin{aligned}\mathbb{E}(X_1) = \mathbb{E}(X_2) &= ap + (-a)(1-p) = a(2p-1) \\ \text{Var}(X_1) = \text{Var}(X_2) &= a^2p + a^2(1-p) - a^2(2p-1)^2 \\ &= a^2[1 - (2p-1)^2].\end{aligned}$$

Furthermore, remembering $u = 1/d = \exp(a)$,

$$\mathbb{E}(\exp(X_1)) = \mathbb{E}(\exp(X_2)) = up + d(1-p).$$

If we define another probability measure \mathbb{Q} by the property $\mathbb{Q}(\omega_1) = q^2$, $\mathbb{Q}(\omega_2) = \mathbb{Q}(\omega_3) = q(1-q)$, $\mathbb{Q}(\omega_4) = (1-q)^2$, with $q = (1-d)/(u-d)$, then

$$\mathbb{E}_{\mathbb{Q}}(\exp(X_1)) = \mathbb{E}_{\mathbb{Q}}(\exp(X_2)) = uq + d(1-q) = 1.$$

Independence

Two events A and B are *independent* if $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B)$.

Two σ -fields \mathcal{F}_1 and \mathcal{F}_2 , which are both contained in \mathcal{F} , are independent if every $A_1 \in \mathcal{F}_1$ and $A_2 \in \mathcal{F}_2$ are independent.

Two random variables X_1 and X_2 are independent if $\mathbb{P}(X_1 \in B_1 \cap X_2 \in B_2) = \mathbb{P}(X_1 \in B_1)\mathbb{P}(X_2 \in B_2)$ for all Borel sets B_1, B_2 .

X_i is called \mathcal{F}_i -measurable, if $\{\omega : X_i(\omega) \in B\} \in \mathcal{F}_i, i = 1, 2$. That is, all events relevant for X_i are in \mathcal{F}_i . The above implies that X_1 and X_2 are independent if X_i is \mathcal{F}_i -measurable and \mathcal{F}_1 and \mathcal{F}_2 are independent.

Example: The events $A_1 = \{HH, HT\}$ (“first toss is heads”) and $A_2 = \{HH, TH\}$ (“second toss is heads”) are independent under \mathbb{P} . The σ -fields

$$\mathcal{F}_1 = \{\emptyset, \{HH, HT\}, \{TH, TT\}, \Omega\}$$

and

$$\mathcal{F}_2 = \{\emptyset, \{HH, TH\}, \{HT, TT\}, \Omega\}$$

are independent, because they represent the outcomes of the two tosses. The random variables X_1 and X_2 are independent, because X_i is \mathcal{F}_i -measurable.

Conditioning

For two events A and B with $\mathbb{P}(B) > 0$, the *conditional probability* of A given B is defined by

$$\mathbb{P}(A|B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}.$$

This may be interpreted as the probability of the event A , based on the information that B occurred. If A and B are independent, then $\mathbb{P}(A|B) = \mathbb{P}(A)$; in that case the event B is irrelevant for the probability of A .

For a discrete random variable X , the *conditional expectation* of X given B is given by

$$\begin{aligned}\mathbb{E}(X|B) &= \int_{\Omega} X(\omega) d\mathbb{P}(\omega|B) \\ &= \sum_{i=1}^{\infty} X(\omega_i) \mathbb{P}(\omega_i|B).\end{aligned}$$

The conditional expectation given a σ -field \mathcal{F} , $\mathbb{E}(X|\mathcal{F})$ is a random variable, defined by

$$\mathbb{E}(X|\mathcal{F}) = \mathbb{E}(X|A_i) \text{ if } \omega \in A_i,$$

for all A_i in \mathcal{F} . When \mathcal{F}_1 is the σ -field *generated* by X_1 (i.e., the smallest σ -field with respect to which X_1 is measurable),

denoted by $\mathcal{F}_1 = \sigma(X_1)$, then we obtain the conditional expectation of Y given X_1 :

$$\mathbb{E}(Y|X_1) = \mathbb{E}(Y|\mathcal{F}_1).$$

Here are some rules for manipulating conditional expectations:

- $\mathbb{E}(X|\mathcal{F}) = X$ if X is \mathcal{F} -measurable
- $\mathbb{E}(XY|\mathcal{F}) = X\mathbb{E}(Y|\mathcal{F})$ if X is \mathcal{F} -measurable
- $\mathbb{E}(X|\mathcal{F}_0) = \mathbb{E}(X)$, where $\mathcal{F}_0 = \{\emptyset, \Omega\}$
- $\mathbb{E}(aX + bY|\mathcal{F}) = a\mathbb{E}(X|\mathcal{F}) + b\mathbb{E}(Y|\mathcal{F})$
- $\mathbb{E}[\mathbb{E}(X|\mathcal{F})] = \mathbb{E}(X)$ (*law of iterated expectation*)
- $\mathbb{E}[\mathbb{E}(X|\mathcal{F})|\mathcal{G}] = \mathbb{E}[\mathbb{E}(X|\mathcal{G})|\mathcal{F}] = \mathbb{E}(X|\mathcal{F})$ if $\mathcal{F} \subset \mathcal{G}$ (*tower property*).

Often the σ -fields upon which one conditions are interpreted as *information sets*.

Example: In the binomial tree example, with $S_1 = S_0 \exp(X_1)$, we have $\mathcal{F}_1 = \sigma(X_1) = \sigma(S_1)$. The conditional expectation of $S_2 = S_1 \exp(X_2)$ given S_1 is

$$\begin{aligned}\mathbb{E}(S_2|\mathcal{F}_1) &= S_1\mathbb{E}[\exp(X_2)|\mathcal{F}_1] \\ &= S_1\mathbb{E}[\exp(X_2)],\end{aligned}$$

because X_2 is independent of X_1 and hence S_1 . Similarly, $E(S_1|\mathcal{F}_0) = S_0E[\exp(X_1)]$. Hence, if we use the measure \mathbb{Q} defined earlier, we have

$$\mathbb{E}_{\mathbb{Q}}(S_i|\mathcal{F}_{i-1}) = S_{i-1}.$$

This final property defines the sequence $\{S_i, \mathcal{F}_i\}$ to be a *martingale*. The probability measure \mathbb{Q} under which this is the case is the *equivalent martingale measure*.

Exercises

1. Define Y to be the number of heads in the example. Derive the σ -field generated by Y .
2. Find the expectation of random variable Y from the previous exercise, and also the conditional expectation of Y given \mathcal{F}_1 . Check that in this case, $\mathbb{E}[\mathbb{E}(Y|\mathcal{F}_1)] = \mathbb{E}(Y)$.
3. In the example, let $Y = 1$ if the two tosses have the same result, and zero otherwise. Show that if $p = \frac{1}{2}$, Y and X_1 are independent, and also that Y and X_2 are independent. Is Y also independent of (X_1, X_2) ?
4. In the binomial tree example, derive an expression for the conditional variance $\text{Var}(S_2|S_1) = \mathbb{E}(S_2^2|S_1) - \{\mathbb{E}(S_2|S_1)\}^2$.
5. For any random variable X and increasing sequence of σ -fields $\{\mathcal{F}_i\}$, show that $X_i = \mathbb{E}(X|\mathcal{F}_i)$ is a martingale.

Multi-step binomial tree

Ingredients:

- Time horizon T , n steps, hence time step is $\delta t = T/n$;
- Sample space Ω , with element $\omega = (\omega^1, \omega^2, \dots, \omega^n)$, where ω^j is H or T;
- When k is # heads in ω , $\mathbb{P}(\omega) = p^k(1-p)^{n-k}$; since there are $\binom{n}{k}$ combinations possible with k heads, $\mathbb{P}(k \text{ heads}) = \binom{n}{k} p^k(1-p)^{n-k}$;
- $X_i(\omega) = \sigma\sqrt{\delta t}$ if $\omega^i = \text{H}$, and $-\sigma\sqrt{\delta t}$ if $\omega^i = \text{T}$;
- $S_i(\omega) = S_{i-1}(\omega)e^{X_i(\omega)}$; S_0 is fixed, and furthermore $\mathcal{F}_i = \sigma(S_1, \dots, S_i) = \sigma(X_1, \dots, X_i)$;
- $B_i = B_{i-1}e^{r\delta t} = B_0e^{ir\delta t}$ (with $B_0 = 1$);
- $q = \frac{e^{r\delta t} - d}{u - d}$, and $\mathbb{Q}(k \text{ heads}) = \binom{n}{k} q^k(1-q)^{n-k}$, where $u = 1/d = \exp(\sigma\sqrt{\delta t})$.

Define the discounted stock price $Z_i = B_i^{-1}S_i = e^{-ir\delta t}S_i$, so that

$$\begin{aligned} Z_i &= e^{-ir\delta t}S_{i-1}e^{X_i} \\ &= e^{-(i-1)r\delta t}S_{i-1}e^{-r\delta t}e^{X_i} \\ &= Z_{i-1}e^{X_i - r\delta t}. \end{aligned}$$

Then

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}}(Z_i | \mathcal{F}_{i-1}) &= Z_{i-1} \mathbb{E}_{\mathbb{Q}}(e^{X_i - r\delta t} | \mathcal{F}_{i-1}) \\ &= Z_{i-1} \left[ue^{-r\delta t} \frac{e^{r\delta t} - d}{u - d} + de^{-r\delta t} \frac{u - e^{r\delta t}}{u - d} \right] \\ &= Z_{i-1}. \end{aligned}$$

Therefore, under \mathbb{Q} , the discounted stock price is a martingale. We therefore call \mathbb{Q} the *equivalent martingale measure* (\mathbb{Q} equivalent to \mathbb{P} means $\mathbb{Q}(A) = 0 \Leftrightarrow \mathbb{P}(A) = 0$).

Because

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}} \left(\frac{S_i - S_{i-1}}{S_{i-1}} \middle| \mathcal{F}_{i-1} \right) &= \mathbb{E}_{\mathbb{Q}} (e^{X_i} - 1 | \mathcal{F}_{i-1}) \\ &= e^{r\delta t} - 1 \\ &= \frac{B_i - B_{i-1}}{B_{i-1}}, \end{aligned}$$

we also call \mathbb{Q} the *risk-neutral measure*.

Construction of a claim

1. Claim (payoff) X at time T ; define $E_T = B_T^{-1}X$;
2. Define $E_i = \mathbb{E}_{\mathbb{Q}}(B_T^{-1}X|\mathcal{F}_i)$; then E_i is a \mathbb{Q} -martingale with respect to $\mathcal{F}_i = \sigma(S_1, \dots, S_i)$.
3. Binomial representation theorem:

$$E_i = E_0 + \sum_{k=1}^i \phi_k(Z_k - Z_{k-1}),$$

or $\Delta E_i = \phi_i \Delta Z_i$, with $\phi_i \in \mathcal{F}_{i-1}$ (*previsible*).

4. Define $\phi_i(S_{i-1}) = \frac{E_i(\text{up}) - E_i(\text{down})}{Z_i(\text{up}) - Z_i(\text{down})}$.
5. Strategy: portfolio of ϕ_{i+1} stocks and $\psi_{i+1} = (E_i - \phi_{i+1}B_i^{-1}S_i)$ bonds. At $i = 0$, $V_0 = \phi_1 S_0 + \psi_1 B_0$. At any i , the value becomes

$$\begin{aligned} \phi_i S_i + \psi_i B_i &= \phi_i S_i + E_{i-1} B_i - \phi_i B_{i-1}^{-1} S_{i-1} B_i \\ &= B_i [E_{i-1} + \phi_i (B_i^{-1} S_i - B_{i-1}^{-1} S_{i-1})] \\ &= B_i [E_{i-1} + \phi_i (Z_i - Z_{i-1})] \\ &= B_i E_i \\ &= \phi_{i+1} S_i + \psi_{i+1} B_i =: V_i. \end{aligned}$$

This is *self-financing*: $V_i - V_{i-1} = \phi_i (S_i - S_{i-1}) + \psi_i (B_i - B_{i-1})$, and replicating: $V_T = B_T E_T = X$.

Early exercise

American-type call option:

- at time T : $V_T = X = [S_T - K]^+$
- at time $T - 1$ (take $\delta t = 1$):

$$V_{T-1} = \max \{ [S_{T-1} - K]^+, e^{-r} \mathbb{E}_{\mathbb{Q}}([S_T - K]^+ | \mathcal{F}_{T-1}) \}.$$
- at time $T - 2$:

$$\begin{aligned} V_{T-2} &= \max \{ [S_{T-2} - K]^+, e^{-r} \mathbb{E}_{\mathbb{Q}}(V_{T-1} | \mathcal{F}_{T-2}) \} \\ &= \max \{ [S_{T-2} - K]^+, e^{-r} \mathbb{E}_{\mathbb{Q}}([S_{T-1} - K]^+ | \mathcal{F}_{T-2}), \\ &\quad e^{-2r} \mathbb{E}_{\mathbb{Q}}([S_T - K]^+ | \mathcal{F}_{T-2}) \}. \end{aligned}$$
- at time i : choice between $T - i + 1$ European call options, expiring at time $i, i + 1, \dots, T$. The value of the American option is the maximum of the values of these European options.
- In the absence of dividend payments on the stock, it can be shown that for a *call* option, it never pays to exercise early (so value American option = value European option). For in-the-money *put* options however ($X = [K - S_T]^+$, $K > S_0$), the values will differ.