

# Math 5760/6890: Introduction to Mathematical Finance

## The Cox-Ross-Rubinstein model

See Petters and Dong 2016, Section 5.2

Akil Narayan<sup>1</sup>

<sup>1</sup>Department of Mathematics, and Scientific Computing and Imaging (SCI) Institute  
University of Utah

Fall 2024



We have modeled a security's price  $S_j = S(t_j)$  via,

$$S_{j+1} = G_{j+1}S_j, \quad G_j = \begin{cases} u, & \text{with probability } p \\ d, & \text{with probability } 1 - p \end{cases}$$

From this model, we've concluded:

- $L := \log(S_n/S_0)$  is a scaled/shifted Binomial( $n, p$ ) random variable.
- $S_n = S_0 e^L$  is the exponential of a scaled/shifted Binomial random variable
- The triple  $(p, u, d)$  determines the distribution entirely.

It's worth pointing out some terminology that relates to our model

- The return is determined by the sequence of iid (scaled/shifted) Bernoulli random variables  $\{L_j\}_{j=1}^n$ . This sequence is called a scaled/shifted **Bernoulli process**.

It's worth pointing out some terminology that relates to our model

- The return is determined by the sequence of iid (scaled/shifted) Bernoulli random variables  $\{L_j\}_{j=1}^n$ . This sequence is called a scaled/shifted **Bernoulli process**.
- The cumulative log-return  $L = \sum_{j=1}^n L_j$  is a sum of independent scaled/shifted Bernoulli random variables. This is called a scaled/shifted **Bernoulli counting process**. (We've identified it as a sequence of Binomial random variables.)

It's worth pointing out some terminology that relates to our model

- The return is determined by the sequence of iid (scaled/shifted) Bernoulli random variables  $\{L_j\}_{j=1}^n$ . This sequence is called a scaled/shifted **Bernoulli process**.
- The cumulative log-return  $L = \sum_{j=1}^n L_j$  is a sum of independent scaled/shifted Bernoulli random variables. This is called a scaled/shifted **Bernoulli counting process**. (We've identified it as a sequence of Binomial random variables.)
- Scaled/shifted Bernoulli counting processes are called (additive) **random walks**.

It's worth pointing out some terminology that relates to our model

- The return is determined by the sequence of iid (scaled/shifted) Bernoulli random variables  $\{L_j\}_{j=1}^n$ . This sequence is called a scaled/shifted **Bernoulli process**.
- The cumulative log-return  $L = \sum_{j=1}^n L_j$  is a sum of independent scaled/shifted Bernoulli random variables. This is called a scaled/shifted **Bernoulli counting process**. (We've identified it as a sequence of Binomial random variables.)
- Scaled/shifted Bernoulli counting processes are called (additive) **random walks**.
- The exponential of this walk,  $e^L$ , is called a **multiplicative** or **geometric random walk**.

It's worth pointing out some terminology that relates to our model

- The return is determined by the sequence of iid (scaled/shifted) Bernoulli random variables  $\{L_j\}_{j=1}^n$ . This sequence is called a scaled/shifted **Bernoulli process**.
- The cumulative log-return  $L = \sum_{j=1}^n L_j$  is a sum of independent scaled/shifted Bernoulli random variables. This is called a scaled/shifted **Bernoulli counting process**. (We've identified it as a sequence of Binomial random variables.)
- Scaled/shifted Bernoulli counting processes are called (additive) **random walks**.
- The exponential of this walk,  $e^L$ , is called a **multiplicative** or **geometric random walk**.

Hence, we have constructed a geometric random walk model for an asset's price.

How can we choose  $(p, u, d)$  to accurately emulate real stock price behavior?

We will follow the *Cox-Ross-Rubinstein* (CRR) model, which in part appeals to limiting continuous-time arguments to set the discrete-time values  $(p, u, d)$ .



How can we choose  $(p, u, d)$  to accurately emulate real stock price behavior?

We will follow the *Cox-Ross-Rubinstein* (CRR) model, which in part appeals to limiting continuous-time arguments to set the discrete-time values  $(p, u, d)$ . We need to determine 3 unknowns:  $p, u, d$ . One suspects we need three constraints to determine these.

At a high level the CRR model imposes the following constraints:

- The recombination constraint  $ud = 1$ . This implies that at every even time index, there is non-zero probability of returning to the original share price.

How can we choose  $(p, u, d)$  to accurately emulate real stock price behavior?

We will follow the *Cox-Ross-Rubinstein* (CRR) model, which in part appeals to limiting continuous-time arguments to set the discrete-time values  $(p, u, d)$ . We need to determine 3 unknowns:  $p, u, d$ . One suspects we need three constraints to determine these.

At a high level the CRR model imposes the following constraints:

- The recombination constraint  $ud = 1$ . This implies that at every even time index, there is non-zero probability of returning to the original share price.
- The per-period mean return (or log return) should match the real-world statistical mean

How can we choose  $(p, u, d)$  to accurately emulate real stock price behavior?

We will follow the *Cox-Ross-Rubinstein* (CRR) model, which in part appeals to limiting continuous-time arguments to set the discrete-time values  $(p, u, d)$ . We need to determine 3 unknowns:  $p, u, d$ . One suspects we need three constraints to determine these.

At a high level the CRR model imposes the following constraints:

- The recombination constraint  $ud = 1$ . This implies that at every even time index, there is non-zero probability of returning to the original share price.
- The per-period mean return (or log return) should match the real-world statistical mean
- The per-period log return variance should match the real-world statistical variance

How can we choose  $(p, u, d)$  to accurately emulate real stock price behavior?

We will follow the *Cox-Ross-Rubinstein* (CRR) model, which in part appeals to limiting continuous-time arguments to set the discrete-time values  $(p, u, d)$ . We need to determine 3 unknowns:  $p, u, d$ . One suspects we need three constraints to determine these.

At a high level the CRR model imposes the following constraints:

- The recombination constraint  $ud = 1$ . This implies that at every even time index, there is non-zero probability of returning to the original share price.
- The per-period mean return (or log return) should match the real-world statistical mean
- The per-period log return variance should match the real-world statistical variance

These are all reasonable assumptions, and one of the major appeals of the CRR Model is that these rather simple assumptions yield probabilistic models that behave like real-world stock prices.

A more abstract reason for these constraints is that they can be used to construct a well-posed underlying continuous-time mathematical model.

The first CRR constraint is that,

$$ud = 1 \quad \implies \quad d = \frac{1}{u}$$

One visually appealing result of this assumption is that the tree has a symmetry around  $S_n = S_0$ . (It is easier to see this for log-returns.)

The first CRR constraint is that,

$$ud = 1 \quad \implies \quad d = \frac{1}{u}$$

One visually appealing result of this assumption is that the tree has a symmetry around  $S_n = S_0$ . (It is easier to see this for log-returns.)

There are also some other reasons why this is an attractive choice:

- It's useful when using the pricing model to compute valuation metrics (e.g., the “theta”)

The first CRR constraint is that,

$$ud = 1 \quad \implies \quad d = \frac{1}{u}$$

One visually appealing result of this assumption is that the tree has a symmetry around  $S_n = S_0$ . (It is easier to see this for log-returns.)

There are also some other reasons why this is an attractive choice:

- It's useful when using the pricing model to compute valuation metrics (e.g., the “theta”)
- It's a sensible condition if we adopt the hypothesis that sequential upward and downward market movements are geometrically symmetric.

The first CRR constraint is that,

$$ud = 1 \quad \implies \quad d = \frac{1}{u}$$

One visually appealing result of this assumption is that the tree has a symmetry around  $S_n = S_0$ . (It is easier to see this for log-returns.)

There are also some other reasons why this is an attractive choice:

- It's useful when using the pricing model to compute valuation metrics (e.g., the “theta”)
- It's a sensible condition if we adopt the hypothesis that sequential upward and downward market movements are geometrically symmetric.
- I.e., it makes the log-return a symmetric (not necessarily centered) random walk:

$$\log L_j = \begin{cases} \log u, & \text{with probability } p, \\ -\log u, & \text{with probability } 1 - p. \end{cases}$$



We seek to impose that the average of the geometric random walk process emulate an expected return rate.

We *could* do this by using historical data to compute an inter-period average return. However, there are some problems with this:

We seek to impose that the average of the geometric random walk process emulate an expected return rate.

We *could* do this by using historical data to compute an inter-period average return. However, there are some problems with this:

- Stock data typically computes average (say annual) return rate, which in other finance contexts is treated essentially as a continuous-time rate.
- It's somewhat awkward to understand what should happen when a more complicated model is desired:
  - ▶ Suppose we want the random model to have 4 coin flips per day.
  - ▶ We have to determine an appropriate quarter-day return rate.
  - ▶ This is not too difficult, but it's easier if we assume a continuous return rate and use that to determine discrete-time rates.

Because our random walk is geometric, it's easier to calibrate the mean of the log-returns.

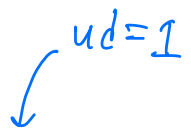
With  $T$  fixed, suppose we choose a number of equal periods  $n$  that are used to divide  $[0, T]$ .

From our previous notation:  $h = h_n := T/n$ .

Because our random walk is geometric, it's easier to calibrate the mean of the log-returns.

With  $T$  fixed, suppose we choose a number of equal periods  $n$  that are used to divide  $[0, T]$ .  
From our previous notation:  $h = h_n := T/n$ .

The mean of the log-return is given by:

$$\mathbb{E}L_j = \mathbb{E} \log G_j = p \log u + (1 - p) \log d = (2p - 1) \log u$$


This mean is the expected (log-)return over a time period of length  $h_n$ .

Because our random walk is geometric, it's easier to calibrate the mean of the log-returns.

With  $T$  fixed, suppose we choose a number of equal periods  $n$  that are used to divide  $[0, T]$ . From our previous notation:  $h = h_n := T/n$ .

The mean of the log-return is given by:

$$\mathbb{E}L_j = \mathbb{E} \log G_j = p \log u + (1 - p) \log d = (2p - 1) \log u$$

This mean is the expected (log-)return over a time period of length  $h_n$ .

To standardize this mean relative to its time interval, we consider

$$\mu_n := \frac{\mathbb{E}L_j}{h_n} = \frac{n}{T} (2p - 1) \log u.$$

Because our random walk is geometric, it's easier to calibrate the mean of the log-returns.

With  $T$  fixed, suppose we choose a number of equal periods  $n$  that are used to divide  $[0, T]$ .  
From our previous notation:  $h = h_n := T/n$ .

The mean of the log-return is given by:

$$\mathbb{E}L_j = \mathbb{E} \log G_j = p \log u + (1 - p) \log d = (2p - 1) \log u$$

This mean is the expected (log-)return over a time period of length  $h_n$ .

To standardize this mean relative to its time interval, we consider

$$\mu_n := \frac{\mathbb{E}L_j}{h_n} = \frac{n}{T} (2p - 1) \log u.$$

The assumption in a CRR model is that, as  $n \rightarrow \infty$ ,  $\mu_n$  converges to a constant value:

$$\lim_{n \rightarrow \infty} \mu_n = \mu_{\text{RW}}, \quad \text{"RW" : "real world"}$$

i.e.,  $\mu_n \approx \mu_{\text{RW}}$  for large  $n$ .

Note that  $\mu$  is in units of one over time. As with most rates, the unit of time used in practice is typically a year.

The constant  $\mu$  is an instantaneous log-return: it is frequently called the real-world/instantaneous/continuous-time **drift**.

To approximate  $\mu$ , we use data in order to realize the approximation,

$$\mu \approx \mathbb{E} \frac{L}{h} = \mathbb{E} \frac{1}{t_1 - t_0} \log \frac{S(t_1)}{S(t_0)}.$$

Note that  $\mu$  is in units of one over time. As with most rates, the unit of time used in practice is typically a year.

The constant  $\mu$  is an instantaneous log-return: it is frequently called the real-world/instantaneous/continuous-time **drift**.

To approximate  $\mu$ , we use data in order to realize the approximation,

$$\mu \approx \mathbb{E} \frac{L}{h} = \mathbb{E} \frac{1}{t_1 - t_0} \log \frac{S(t_1)}{S(t_0)}.$$

## Example

Suppose we are given (deterministic) daily security prices  $S_0, S_1, S_2, \dots, S_n$ .

(E.g., this could come from historical data, and there are  $\approx 252$  trading days per year.)

Compute (an approximation to) the continuous-time drift.

A single sample of the (time-normalized) log-return is  $\frac{1}{t_{j+1} - t_j} \log \frac{S_{j+1}}{S_j}$



$1 \approx \underbrace{252}_{n} \text{ periods} \Rightarrow h_n = \frac{1}{252} = t_{j+1} - t_j$  (assume 1 year w/ 252 samples)

$$\Rightarrow \left\{ \underbrace{\frac{1}{t_{j+1} - t_j}}_{\left(\frac{1}{252}\right)} \log\left(\frac{S_{j+1}}{S_j}\right) \right\}_{j=0}^{251} = \left\{ 252 \log\left(\frac{S_{j+1}}{S_j}\right) \right\}_{j=0}^{251}$$

empirical mean of these is an approximation (from data) to  $\mu = \mu_{RW}$ .

The final CRR constraint is similar to the mean-matching condition: we assume a well-defined instantaneous rate of change for the variance.

The discrete-time variance, normalized by the time period, is

$$\sigma_n^2 := \frac{1}{h} \text{Var} \log \frac{S_{j+1}}{S_j},$$

which is independent of  $j$  due to the iid property of the log-returns.

The final CRR constraint is similar to the mean-matching condition: we assume a well-defined instantaneous rate of change for the variance.

The discrete-time variance, normalized by the time period, is

$$\sigma_n^2 := \frac{1}{h} \text{Var} \log \frac{S_{j+1}}{S_j},$$

which is independent of  $j$  due to the iid property of the log-returns.

The CRR model makes the assumption that  $n \rightarrow \infty$ , then  $\sigma_n^2$  converges to a constant,

$$\sigma_n^2 = \lim_{n \rightarrow \infty} \sigma_n^2.$$

The constant  $\sigma$  is the real-world/instantaneous/continuous-time **volatility**.

The final CRR constraint is similar to the mean-matching condition: we assume a well-defined instantaneous rate of change for the variance.

The discrete-time variance, normalized by the time period, is

$$\sigma_n^2 := \frac{1}{h} \text{Var} \log \frac{S_{j+1}}{S_j},$$

which is independent of  $j$  due to the iid property of the log-returns.

The CRR model makes the assumption that  $n \rightarrow \infty$ , then  $\sigma_n^2$  converges to a constant,

$$\sigma^2 = \lim_{n \rightarrow \infty} \sigma_n^2.$$

The constant  $\sigma$  is the real-world/instantaneous/continuous-time **volatility**.

Unlike  $\mu$ , the limiting process for  $\sigma$  (not  $\sigma^2$ ) scales like  $1/\sqrt{n}$ .

Just like the continuous-time drift  $\mu$ , estimate the volatility  $\sigma$  is typically accomplished through access to data with the finite-time approximation,

$$\sigma^2 \approx \frac{1}{h} \text{Var}L = \frac{1}{t_1 - t_0} \text{Var} \log \frac{S(t_1)}{S(t_0)}.$$

### Example

Suppose we are given (deterministic) daily security prices  $S_0, S_1, S_2, \dots, S_n$ .  
(E.g., this could come from historical data, and there are  $\approx 252$  trading days per year.)  
Compute (an approximation to) the continuous-time volatility.

The CRR model places the following additional constraints on our standard Binomial tree model:

- Geometric symmetry of tree prices:  $u = 1/d$
- The continuous-time limit of the expected log-return matches the real-world drift:

$$\mu = \lim_{n \rightarrow \infty} \frac{1}{h_n} \mathbb{E}L_j$$

- The continuous-time limit of the variance of the log-return matches the real-world (squared) volatility:

$$\sigma^2 = \lim_{n \rightarrow \infty} \frac{1}{h_n} \text{Var}L_j$$

The CRR model places the following additional constraints on our standard Binomial tree model:

- Geometric symmetry of tree prices:  $u = 1/d$
- The continuous-time limit of the expected log-return matches the real-world drift:

$$\mu = \lim_{n \rightarrow \infty} \frac{1}{h_n} \mathbb{E}L_j$$

- The continuous-time limit of the variance of the log-return matches the real-world (squared) volatility:

$$\sigma^2 = \lim_{n \rightarrow \infty} \frac{1}{h_n} \text{Var}L_j$$

Of course, we are not (yet) going to take the limit as  $n \uparrow \infty$ , so what do we do for finite but large  $n$ ?

In the last two bullets above, both  $\mathbb{E}L_j$  and  $\text{Var}L_j$  depend on  $(p, u, d)$ .

The CRR model places the following additional constraints on our standard Binomial tree model:

- Geometric symmetry of tree prices:  $u = 1/d$
- The continuous-time limit of the expected log-return matches the real-world drift:

$$\mu = \lim_{n \rightarrow \infty} \frac{1}{h_n} \mathbb{E}L_j$$

- The continuous-time limit of the variance of the log-return matches the real-world (squared) volatility:

$$\sigma^2 = \lim_{n \rightarrow \infty} \frac{1}{h_n} \text{Var}L_j$$

Of course, we are not (yet) going to take the limit as  $n \uparrow \infty$ , so what do we do for finite but large  $n$ ?

In the last two bullets above, both  $\mathbb{E}L_j$  and  $\text{Var}L_j$  depend on  $(p, u, d)$ .

Hence, for finite  $n$ ,  $(p, u, d)$  should depend on the time discretization parameter  $n$ . I.e., for finite  $n$ ,

$$(p, u, d) = (p_n, u_n, d_n).$$

Next time: how do we choose  $(p_n, u_n, d_n)$  to match the constraints above?





Petters, Arlie O. and Xiaoying Dong (2016). *An Introduction to Mathematical Finance with Applications: Understanding and Building Financial Intuition*. Springer. ISBN: 978-1-4939-3783-7.