

Math 5760/6890: Introduction to Mathematical Finance

Continuous-time limits, I

See Petters and Dong 2016, Section 5.2, 5.3

Akil Narayan¹

¹Department 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.

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.

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

This results (after some approximation) in the following *real-world CRR equations*:

$$u_n = \exp(\sigma\sqrt{h_n}), \quad d_n = \exp(-\sigma\sqrt{h_n}), \quad p_n = \frac{1}{2} \left(1 + \frac{\mu}{\sigma} \sqrt{h_n} \right).$$

The distribution of S_n

D20-S03(a)

What kind of distribution does S_n have? It will be useful to write S_n in terms of *standardizations* of L_j .

The distribution of S_n

What kind of distribution does S_n have? It will be useful to write S_n in terms of *standardizations* of L_j .

A standardization of L_j (or of any random variable) is

$$\tilde{L}_j = \frac{L_j - \mathbb{E}L_j}{\sqrt{\text{Var}L_j}}, \quad \text{i.e. } L_j = \mathbb{E}L_j + \sqrt{\text{Var}L_j} \tilde{L}_j$$

i.e., it is a centered version of L_j , inversely scaled by its standard deviation: standardizations of random variables are mean-0 and variance-1.

In terms of the Binomial tree parameters p_n , we have that,

$$\mathbb{E}L_j = \mu h_n,$$

$$\text{Var}L_j = 4p_n(1 - p_n)\sigma^2 h_n.$$

$$L_j = \begin{cases} \log u_n, & \text{w/ prob } p_n \\ \log d_n, & \text{w/ prob } 1-p_n \end{cases}$$

Note that this agrees with our real-world CRR approximation for large n : $\text{Var}L_j \sim \sigma^2 h_n$.

What kind of distribution does S_n have? It will be useful to write S_n in terms of *standardizations* of L_j .

A standardization of L_j (or of any random variable) is

$$\tilde{L}_j = \frac{L_j - \mathbb{E}L_j}{\sqrt{\text{Var}L_j}},$$

i.e., it is a centered version of L_j , inversely scaled by its standard deviation: standardizations of random variables are mean-0 and variance-1.

In terms of the Binomial tree parameters p_n , we have that,

$$\mathbb{E}L_j = \mu h_n, \quad \text{Var}L_j = 4p_n(1 - p_n)\sigma^2 h_n.$$

Note that this agrees with our real-world CRR approximation for large n : $\text{Var}L_j \sim \sigma^2 h_n$. Hence, the \tilde{L}_j variables have distribution:

$$(L_j = \mathbb{E}L_j + \sqrt{\text{Var}L_j} \tilde{L}_j)$$

With the standardization of the L_j variables, we have,

$$S_n = S_0 \exp \left(\sum_{j=1}^n \left(\mathbb{E}L_j + \tilde{L}_j \sqrt{\text{Var}L_j} \right) \right).$$

$$S_n = S_0 \exp \left(\sum_{j=1}^n \left(\mathbb{E}L_j + \tilde{L}_j \sqrt{\text{Var}L_j} \right) \right).$$

This expression allows us to understand the large- n behavior of S_n .

$$\mathbb{E}L_j = \mu h_n$$

$$\text{Var}L_j = 4p_n(1-p_n)\sigma^2 h_n$$

$$\Rightarrow \sum_{j=1}^n \mathbb{E}L_j = \mu h_n n = \mu T \quad (T = nh_n)$$

$$\sqrt{\text{Var}L_j} = \sqrt{4p_n(1-p_n)} \sigma \sqrt{h_n} = \sqrt{4p_n(1-p_n)} \sigma \sqrt{\frac{T}{n}}$$

$$\sum_{j=1}^n (\mathbb{E}L_j + \tilde{L}_j \sqrt{\text{Var}L_j}) = \mu T + \sigma \sqrt{T} \frac{1}{\sqrt{n}} \sum_{j=1}^n \underbrace{\sqrt{4p_n(1-p_n)}}_{\rightarrow 1} \tilde{L}_j$$

for large n : $S_n \approx S_0 \exp(\mu T + \sigma \sqrt{T} \cdot \frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j)$

$$S_n = S_0 \exp \left(\sum_{j=1}^n \left(\mathbb{E}L_j + \tilde{L}_j \sqrt{\text{Var}L_j} \right) \right).$$

This expression allows us to understand the large- n behavior of S_n .

After some manipulation, we find that

$$S_n = S_0 \exp \left(\underbrace{\mu T + \sqrt{4p_n(1-p_n)}\sigma\sqrt{T}}_{\rightarrow 1} \frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j \right)$$

The goal is to take $n \uparrow \infty$.

$$S_n = S_0 \exp \left(\sum_{j=1}^n \left(\mathbb{E}L_j + \tilde{L}_j \sqrt{\text{Var}L_j} \right) \right).$$

This expression allows us to understand the large- n behavior of S_n .

After some manipulation, we find that

$$S_n = S_0 \exp \left(\mu T + \sqrt{4p_n(1-p_n)} \sigma \sqrt{T} \frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j \right)$$

The goal is to take $n \uparrow \infty$.

Note that:

$$\lim_{n \uparrow \infty} \exp(\mu T) = \exp(\mu T),$$

$$\lim_{n \uparrow \infty} \exp(\sigma \sqrt{T} \sqrt{4p_n(1-p_n)}) = \exp(\sigma \sqrt{T}).$$

But what about $\exp \left(\frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j \right)$?

The form of the quantity,

$$\frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j,$$

is reminiscent of a classical result in probability theory.

The form of the quantity,

$$\frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j,$$

is reminiscent of a classical result in probability theory.

Theorem (Central Limit Theorem)

Let $\{X_j\}_{j=1}^{\infty}$ be iid random variables with zero mean and variance σ^2 . Then,

$$\lim_{n \rightarrow \infty} \frac{1}{\sqrt{n}} \sum_{j=1}^n X_j \sim \mathcal{N}(0, \sigma^2).$$

"is distributed like"

normal RV, mean 0, variance σ^2

The form of the quantity,

$$\frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j,$$

is reminiscent of a classical result in probability theory.

Theorem (Central Limit Theorem)

Let $\{X_j\}_{j=1}^{\infty}$ be iid random variables with zero mean and variance σ^2 . Then,

$$\lim_{n \rightarrow \infty} \frac{1}{\sqrt{n}} \sum_{j=1}^n X_j \sim \mathcal{N}(0, \sigma^2).$$

Remarks:

- This result is convergence *in distribution*, but is not stronger than that.

The form of the quantity,

$$\frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j,$$

is reminiscent of a classical result in probability theory.

Theorem (Central Limit Theorem)

Let $\{X_j\}_{j=1}^{\infty}$ be iid random variables with zero mean and variance σ^2 . Then,

$$\lim_{n \rightarrow \infty} \frac{1}{\sqrt{n}} \sum_{j=1}^n X_j \sim \mathcal{N}(0, \sigma^2).$$

Remarks:

- This result is convergence *in distribution*, but is not stronger than that.
- A direct corollary: the error of the empirical “Monte Carlo” mean scales like $\sqrt{\text{Var} X_j} / \sqrt{n}$.

The form of the quantity,

$$\frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j,$$

is reminiscent of a classical result in probability theory.

Theorem (Central Limit Theorem)

Let $\{X_j\}_{j=1}^{\infty}$ be iid random variables with zero mean and variance σ^2 . Then,

$$\lim_{n \rightarrow \infty} \frac{1}{\sqrt{n}} \sum_{j=1}^n X_j \sim \mathcal{N}(0, \sigma^2).$$

Remarks:

- This result is convergence *in distribution*, but is not stronger than that.
- A direct corollary: the error of the empirical “Monte Carlo” mean scales like $\sqrt{\text{Var} X_j} / \sqrt{n}$.
- It is important that the X_j random variables *not* depend on n .

We need to determine the n -asymptotic behavior of

$$\frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j,$$

where the \tilde{L}_j are indeed iid.

We need to determine the n -asymptotic behavior of

$$\frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j,$$

where the \tilde{L}_j are indeed iid.

The problem: The distribution of \tilde{L}_j *does* depend on n .

To more formally understand why this is an issue: for each fixed n , we have the collection of random variables,

$$\tilde{L}_{n,1}, \tilde{L}_{n,2}, \dots, \tilde{L}_{n,n}, \quad \tilde{L}_{n,j} = \frac{L_j - \mathbb{E}L_j}{\sqrt{\text{Var}L_j}}$$

($\tilde{L}_{n+1,1}$ has a different distribution than $\tilde{L}_{n,1}$)

We need to determine the n -asymptotic behavior of

$$\frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j,$$

where the \tilde{L}_j are indeed iid.

The problem: The distribution of \tilde{L}_j *does* depend on n .

To more formally understand why this is an issue: for each fixed n , we have the collection of random variables,

$$\tilde{L}_{n,1}, \tilde{L}_{n,2}, \dots, \tilde{L}_{n,n}, \quad \tilde{L}_{n,j} = \frac{L_j - \mathbb{E}L_j}{\sqrt{\text{Var}L_j}}$$

However, the parameter (p_n, u_n, d_n) depend on n , and therefore the distribution of $\tilde{L}_{n,j}$ depends on n .

We need to determine the n -asymptotic behavior of

$$\frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j,$$

where the \tilde{L}_j are indeed iid.

The problem: The distribution of \tilde{L}_j *does* depend on n .

To more formally understand why this is an issue: for each fixed n , we have the collection of random variables,

$$\tilde{L}_{n,1}, \tilde{L}_{n,2}, \dots, \tilde{L}_{n,n}, \quad \tilde{L}_{n,j} = \frac{L_j - \mathbb{E}L_j}{\sqrt{\text{Var}L_j}}$$

However, the parameter (p_n, u_n, d_n) depend on n , and therefore the distribution of $\tilde{L}_{n,j}$ depends on n .

The Central Limit Theorem as we've stated it does *not* directly tell us about the $n \rightarrow \infty$ limit of,

$$\frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_{n,j}.$$

Although the distribution of $\tilde{L}_{n,j}$ depends on n , for large n and m the distributions of $\tilde{L}_{n,j}$ and $\tilde{L}_{m,j}$ are actually quite similar.

Although the distribution of $\tilde{L}_{n,j}$ depends on n , for large n and m the distributions of $\tilde{L}_{n,j}$ and $\tilde{L}_{m,j}$ are actually quite similar.

In our case, for example, we could write the $(n + 1)$ st summation as,

$$\frac{1}{\sqrt{n+1}} \sum_{j=1}^{n+1} \tilde{L}_{n+1,j} = \underbrace{\frac{1}{\sqrt{n+1}} \sum_{j=1}^n \tilde{L}_{n,j} + \frac{1}{\sqrt{n+1}} \tilde{L}_{n+1,n+1}}_{(a)} + \frac{1}{\sqrt{n+1}} \sum_{j=1}^n \left[\tilde{L}_{n+1,j} - \tilde{L}_{n,j} \right].$$

Although the distribution of $\tilde{L}_{n,j}$ depends on n , for large n and m the distributions of $\tilde{L}_{n,j}$ and $\tilde{L}_{m,j}$ are actually quite similar.

In our case, for example, we could write the $(n + 1)$ st summation as,

$$\frac{1}{\sqrt{n+1}} \sum_{j=1}^{n+1} \tilde{L}_{n+1,j} = \underbrace{\frac{1}{\sqrt{n+1}} \sum_{j=1}^n \tilde{L}_{n,j} + \frac{1}{\sqrt{n+1}} \tilde{L}_{n+1,n+1}}_{(a)} + \frac{1}{\sqrt{n+1}} \sum_{j=1}^n [\tilde{L}_{n+1,j} - \tilde{L}_{n,j}].$$

Term (a) is a sum of $n + 1$ independent random variables scaled by $1/\sqrt{n+1}$, but the $(n + 1)$ st summand is *not* identically distributed.

Hence, if we had a Central Limit Theorem for *non*-identically distributed random variables, we could tackle this case.

The Lindeberg Condition

D20-S08(a)

Let $\{X_j\}_{j=1}^{\infty}$ be independent and mean-zero, but not identically distributed.

Definition (Lindeberg's condition)

Let $\Sigma_n^2 := \sum_{j=1}^n \text{Var} X_j$.

Lindeberg's condition is the following on the sequence $\{X_j\}_{j=1}^{\infty}$: For every $\epsilon > 0$, we have,

$$\lim_{n \uparrow \infty} \frac{\sum_{j=1}^n \mathbb{E} \left[X_j^2 \mathbb{1}_{|X_j| > \epsilon \Sigma_n} \right]}{\Sigma_n^2} = 0.$$

$\mathbb{1}_A$: indicator function on set A

$$X_j^2 \mathbb{1}_{|X_j| > \epsilon \Sigma_n} = \begin{cases} 0 & \text{if } |X_j| \leq \epsilon \Sigma_n \\ X_j^2 & \text{else} \end{cases}$$

The Lindeberg Condition

D20-S08(b)

Let $\{X_j\}_{j=1}^{\infty}$ be independent and mean-zero, but not identically distributed.

Definition (Lindeberg's condition)

Let $\Sigma_n^2 := \sum_{j=1}^n \text{Var} X_j$.

Lindeberg's condition is the following on the sequence $\{X_j\}_{j=1}^{\infty}$: For every $\epsilon > 0$, we have,

$$\lim_{n \uparrow \infty} \frac{\sum_{j=1}^n \mathbb{E} \left[X_j^2 \mathbb{1}_{|X_j| > \epsilon \Sigma_n} \right]}{\Sigma_n^2} = 0.$$

Theorem ((Lindeberg) Central Limit Theorem)

Suppose $\{X_j\}_{j=1}^{\infty}$ are independent and mean-zero, and satisfy Lindeberg's condition. Then,

$$\lim_{n \uparrow \infty} \frac{1}{\Sigma_n} \sum_{j=1}^n X_j \sim \mathcal{N}(0, 1).$$

The Lindeberg Condition

D20-S08(c)

Let $\{X_j\}_{j=1}^{\infty}$ be independent and mean-zero, but not identically distributed.

Definition (Lindeberg's condition)

Let $\Sigma_n^2 := \sum_{j=1}^n \text{Var}X_j$.

Lindeberg's condition is the following on the sequence $\{X_j\}_{j=1}^{\infty}$: For every $\epsilon > 0$, we have,

$$\lim_{n \uparrow \infty} \frac{\sum_{j=1}^n \mathbb{E} \left[X_j^2 \mathbb{1}_{|X_j| > \epsilon \Sigma_n} \right]}{\Sigma_n^2} = 0.$$

Theorem ((Lindeberg) Central Limit Theorem)

Suppose $\{X_j\}_{j=1}^{\infty}$ are independent and mean-zero, and satisfy Lindeberg's condition. Then,

$$\lim_{n \uparrow \infty} \frac{1}{\Sigma_n} \sum_{j=1}^n X_j \sim \mathcal{N}(0, 1).$$

The upshot for us: so long as our random variables satisfy the appropriate version of Lindeberg's condition, then we can use the Central Limit Theorem.

For our (“triangular”) sequence of random variables, $\{\tilde{L}_{n,j}\}_{j=1}^n$ with $n \in \mathbb{N}$, Lindeberg's condition for this setup is:
For every $\epsilon > 0$,

$$\lim_{n \uparrow \infty} \frac{1}{n} \sum_{j=1}^n \mathbb{E} \left[\tilde{L}_{n,j}^2 \mathbb{1}_{|\tilde{L}_{n,j}| > \sqrt{n}\epsilon} \right] = 0.$$

For our (“triangular”) sequence of random variables, $\{\tilde{L}_{n,j}\}_{j=1}^n$ with $n \in \mathbb{N}$, Lindeberg's condition for this setup is:
For every $\epsilon > 0$,

$$\lim_{n \uparrow \infty} \frac{1}{n} \sum_{j=1}^n \mathbb{E} \left[\tilde{L}_{n,j}^2 \mathbb{1}_{|\tilde{L}_{n,j}| > \sqrt{n}\epsilon} \right] = 0.$$

This holds in our particular case, which implies:

$$\lim_{n \uparrow \infty} \frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_{n,j} \sim \mathcal{N}(0, 1).$$

Finally, recall that we started with the assertions:

$$S_n = S_0 \exp \left(\mu T + \sqrt{4p_n(1-p_n)} \sigma \sqrt{T} \frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j \right)$$

$$\lim_{n \uparrow \infty} \exp(\mu T) = \exp(\mu T),$$

$$\lim_{n \uparrow \infty} \exp(\sigma \sqrt{T} \sqrt{4p_n(1-p_n)}) = \exp(\sigma \sqrt{T}).$$

Finally, recall that we started with the assertions:

$$S_n = S_0 \exp \left(\mu T + \sqrt{4p_n(1-p_n)} \sigma \sqrt{T} \frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j \right)$$

$$\lim_{n \uparrow \infty} \exp(\mu T) = \exp(\mu T),$$

$$\lim_{n \uparrow \infty} \exp(\sigma \sqrt{T} \sqrt{4p_n(1-p_n)}) = \exp(\sigma \sqrt{T}).$$

We now add to this:

$$\lim_{n \uparrow \infty} \frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_{n,j} \sim \mathcal{N}(0, 1).$$

Recall: \tilde{L}_j actually depends on n , so denote it as $\tilde{L}_{n,j}$.

Finally, recall that we started with the assertions:

$$S_n = S_0 \exp \left(\mu T + \sqrt{4p_n(1-p_n)} \sigma \sqrt{T} \frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j \right)$$

$$\lim_{n \uparrow \infty} \exp(\mu T) = \exp(\mu T),$$

$$\lim_{n \uparrow \infty} \exp(\sigma \sqrt{T} \sqrt{4p_n(1-p_n)}) = \exp(\sigma \sqrt{T}).$$

We now add to this:

$$\lim_{n \uparrow \infty} \frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_{n,j} \sim \mathcal{N}(0, 1).$$

equals in distribution.

Therefore, if $X \sim \mathcal{N}(0, 1)$, then

$$\lim_{n \uparrow \infty} S_n \sim S_0 \exp(\mu T + X \sigma \sqrt{T}).$$

Finally, recall that we started with the assertions:

$$S_n = S_0 \exp \left(\mu T + \sqrt{4p_n(1-p_n)} \sigma \sqrt{T} \frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_j \right)$$

$$\lim_{n \uparrow \infty} \exp(\mu T) = \exp(\mu T),$$

$$\lim_{n \uparrow \infty} \exp(\sigma \sqrt{T} \sqrt{4p_n(1-p_n)}) = \exp(\sigma \sqrt{T}).$$

We now add to this:

$$\lim_{n \uparrow \infty} \frac{1}{\sqrt{n}} \sum_{j=1}^n \tilde{L}_{n,j} \sim \mathcal{N}(0, 1).$$

Therefore, if $X \sim \mathcal{N}(0, 1)$, then

$$\lim_{n \uparrow \infty} S_n \sim S_0 \exp(\mu T + X \sigma \sqrt{T}).$$

Put another way: if $Z \sim \mathcal{N}(\mu T, \sigma^2 T)$, then

$$\lim_{n \uparrow \infty} S_n \sim S_0 \exp(Z),$$

i.e., the continuous-time limit of S_n is the exponential of a normally distributed random variable.

If we let $S(T)$ denote the $n \uparrow \infty$ limit of S_n , we conclude that,

$$S(T) \sim S_0 \exp(Z),$$

$$Z \sim \mathcal{N}(\mu T, \sigma^2 T).$$

If we let $S(T)$ denote the $n \uparrow \infty$ limit of S_n , we conclude that,

$$S(T) \sim S_0 \exp(Z), \quad Z \sim \mathcal{N}(\mu T, \sigma^2 T).$$

A random variable that is the exponential of a normal random variable is called a **lognormal** random variable.

I.e., our continuous-time security price is a lognormal random variable, which is typically written as,

$$S(T) \sim \text{lognormal}(\underbrace{\mu T + \log S_0}_{\text{mean of } \log S(T)}, \underbrace{\sigma^2 T}_{\text{variance of } \log S(T)}).$$

mean of
 $\log S(T)$

variance of
 $\log S(T)$.

If we let $S(T)$ denote the $n \uparrow \infty$ limit of S_n , we conclude that,

$$S(T) \sim S_0 \exp(Z), \quad Z \sim \mathcal{N}(\mu T, \sigma^2 T).$$

A random variable that is the exponential of a normal random variable is called a **lognormal** random variable.

I.e., our continuous-time security price is a lognormal random variable, which is typically written as,

$$S(T) \sim \text{lognormal}(\mu T + \log S_0, \sigma^2 T).$$

- Note that T is arbitrary; e.g., the same rationale implies that $S(T/2)$ is also a lognormal random variable.

If we let $S(T)$ denote the $n \uparrow \infty$ limit of S_n , we conclude that,

$$S(T) \sim S_0 \exp(Z), \quad Z \sim \mathcal{N}(\mu T, \sigma^2 T).$$

A random variable that is the exponential of a normal random variable is called a **lognormal** random variable.

I.e., our continuous-time security price is a lognormal random variable, which is typically written as,

$$S(T) \sim \text{lognormal}(\mu T + \log S_0, \sigma^2 T).$$

- Note that T is arbitrary; e.g., the same rationale implies that $S(T/2)$ is also a lognormal random variable.
- It is not true that $\mathbb{E}S(T) = \mu T$ or $\mathbb{E}S(T) = \exp(\mu T)$. In fact, one can show that

$$\mathbb{E}S(T) = \exp\left(\mu T + \frac{\sigma^2}{2} T\right).$$

Note that this matches our expression for the mean from last time.

Modeling continuous-time prices

D20-S12(a)

For any $t > 0$, our continuous-time CRR model states:

$$S(t) \sim \text{lognormal}(\mu t + \log S_0, \sigma^2 t).$$

How would we simulate a trajectory given (S_0, μ, σ^2) ?

For any $t > 0$, our continuous-time CRR model states:

$$S(t) \sim \text{lognormal}(\mu t + \log S_0, \sigma^2 t).$$

How would we simulate a trajectory given (S_0, μ, σ^2) ? Well, for each t , we could:

1. Generate $Z \sim \mathcal{N}(\mu t + \log S_0, \sigma^2 t)$
2. Set $S(t) = \exp(Z)$

(It's true that above $S(t)$ as the correct distribution.)

For any $t > 0$, our continuous-time CRR model states:


$$S(t) \sim \text{lognormal}(\mu t + \log S_0, \sigma^2 t).$$

How would we simulate a trajectory given (S_0, μ, σ^2) ? Well, for each t , we could:

1. Generate $Z \sim \mathcal{N}(\mu t + \log S_0, \sigma^2 t)$
2. Set $S(t) = \exp(Z)$

(It's true that above $S(t)$ as the correct distribution.)

This, unfortunately, does not produce what we expect:



The missing piece of the puzzle for us is the temporal structure of the signal: consider the model

$$S(t) \sim \text{lognormal}(\mu t + \log S_0, \sigma^2 t).$$

for very small $t \ll 1$.

In this case, $S(t)$ is “very close” to $\exp \log S_0 = S_0$. This fact is reflected in the generated image.

The missing piece of the puzzle for us is the temporal structure of the signal: consider the model

$$S(t) \sim \text{lognormal}(\mu t + \log S_0, \sigma^2 t).$$

for very small $t \ll 1$.

In this case, $S(t)$ is “very close” to $\exp \log S_0 = S_0$. This fact is reflected in the generated image.

What is not accounted for is the Markovian structure of this process: I.e., while $S(T)$ has the distributed specified, if we are provided that the price at time $T - \epsilon$ is $S(T - \epsilon) = s$, then $S(T)$ *conditioned* on this value has a lognormal distribution with small parameters:

$$S(T) \mid S(T - \epsilon) \sim \text{lognormal}(\log S(T - \epsilon) + \mu \epsilon, \sigma^2 \epsilon)$$

The missing piece of the puzzle for us is the temporal structure of the signal: consider the model

$$S(t) \sim \text{lognormal}(\mu t + \log S_0, \sigma^2 t).$$

for very small $t \ll 1$.

In this case, $S(t)$ is “very close” to $\exp \log S_0 = S_0$. This fact is reflected in the generated image.

What is not accounted for is the Markovian structure of this process: I.e., while $S(T)$ has the distributed specified, if we are provided that the price at time $T - \epsilon$ is $S(T - \epsilon) = s$, then $S(T)$ *conditioned* on this value has a lognormal distribution with small parameters:

$$S(T) \mid S(T - \epsilon) \sim \text{lognormal}(\log S(T - \epsilon) + \mu\epsilon, \sigma^2\epsilon)$$

I.e., $S(T)$ should be constrained to lie “close” to $S(T - \epsilon)$, and in particular the asset price should be continuous in time.

We have not captured this structure by only inspecting the distribution.

To more formally understand these concepts, we'll need to introduce stochastic processes (next time).

It's worth considering one more specialization of the (finite- n) binomial tree: the *risk neutral* tree.

Recall the principle of risk neutrality: a probabilistic model is risk neutral if the model's expected value of the asset equals the future value of today's price.

It's worth considering one more specialization of the (finite- n) binomial tree: the *risk neutral* tree.

Recall the principle of risk neutrality: a probabilistic model is risk neutral if the model's expected value of the asset equals the future value of today's price.

When in the context of probabilistic modeling, risk neutrality assumes that the outcomes are the same as in the marketplace.

For binomial trees, the family of outcomes is determined entirely by (u_n, d_n) , which in turn are estimated from real (marketplace) (μ, σ) data.

It's worth considering one more specialization of the (finite- n) binomial tree: the *risk neutral* tree.

Recall the principle of risk neutrality: a probabilistic model is risk neutral if the model's expected value of the asset equals the future value of today's price.

When in the context of probabilistic modeling, risk neutrality assumes that the outcomes are the same as in the marketplace.

For binomial trees, the family of outcomes is determined entirely by (u_n, d_n) , which in turn are estimated from real (marketplace) (μ, σ) data.

Hence, in a risk-neutral "world", the values u_n and d_n should match their values in the marketplace, i.e., due to the real-world CRR equations:

$$u_n = \exp(\sigma\sqrt{h_n}), \quad d_n = \exp(-\sigma\sqrt{h_n}), \quad (\text{risk neutral})$$

The risk-neutral probability

D20-S15(a)

What does change in a risk-neutral world is the probabilistic structure, i.e., p_n .

We choose p_n so that,

$$\text{FV}(S(t_0)) = \mathbb{E}S(t_1).$$

What does change in a risk-neutral world is the probabilistic structure, i.e., p_n .

We choose p_n so that,

$$\text{FV}(S(t_0)) = \mathbb{E}S(t_1).$$

To determine future value, we need the analogue of a risk-free (interest) rate for securities.

Recall that this is provided by

- the actual risk-free rate $r > 0$ (e.g., from risk-free securities)
- the dividend rate $-q < 0$ (a negative rate because paying dividends decreases capital/worth)

What does change in a risk-neutral world is the probabilistic structure, i.e., p_n .

We choose p_n so that,

$$\text{FV}(S(t_0)) = \mathbb{E}S(t_1).$$

To determine future value, we need the analogue of a risk-free (interest) rate for securities.

Recall that this is provided by

- the actual risk-free rate $r > 0$ (e.g., from risk-free securities)
- the dividend rate $-q < 0$ (a negative rate because paying dividends decreases capital/worth)

Hence: $\text{FV}(S(t_0)) = e^{(r-q)h_n} S(t_0)$. Using this in the risk-neutrality condition, we have,

$$e^{(r-q)h_n} = p_n u_n + (1 - p_n) d_n,$$

i.e.,

$$p_n = \frac{e^{(r-q)h_n} - d_n}{u_n - d_n}.$$

(Recall a convenient fact: assuming a no-arbitrage market implies $0 < p_n < 1$.)

The risk-neutral CRR model

D20-S16(a)

The risk-neutral CRR model has the conditions:

$$u_n = \exp(\sigma\sqrt{h_n}),$$

$$d_n = \exp(-\sigma\sqrt{h_n}),$$

$$p_n = \frac{e^{(r-q)h_n} - d_n}{u_n - d_n}.$$

The risk-neutral CRR model has the conditions:

$$u_n = \exp(\sigma\sqrt{h_n}), \quad d_n = \exp(-\sigma\sqrt{h_n}), \quad p_n = \frac{e^{(r-q)h_n} - d_n}{u_n - d_n}.$$

With some analysis (similar to the standard CRR model), one can determine that for large n , one has the valid approximations

$$u_n = \exp(\sigma_{RN}\sqrt{h_n}), \quad d_n = \exp(-\sigma_{RN}\sqrt{h_n}), \quad p_n = \frac{1}{2} \left(1 + \frac{\mu_{RN}}{\sigma_{RN}} \sqrt{h_n} \right),$$

where (μ_{RN}, σ_{RN}) are the risk-neutral drift and volatility, which satisfy:

$$\sigma_{RN} = \sigma, \quad \mu_{RN} = r - q - \frac{\sigma^2}{2}$$

Hence, one can use these equations to set (p_n, u_n, d_n) for a risk-neutral CRR tree.

The risk-neutral CRR model has the conditions:

$$u_n = \exp(\sigma\sqrt{h_n}), \quad d_n = \exp(-\sigma\sqrt{h_n}), \quad p_n = \frac{e^{(r-q)h_n} - d_n}{u_n - d_n}.$$

With some analysis (similar to the standard CRR model), one can determine that for large n , one has the valid approximations

$$u_n = \exp(\sigma_{RN}\sqrt{h_n}), \quad d_n = \exp(-\sigma_{RN}\sqrt{h_n}), \quad p_n = \frac{1}{2} \left(1 + \frac{\mu_{RN}}{\sigma_{RN}} \sqrt{h_n} \right),$$

where (μ_{RN}, σ_{RN}) are the risk-neutral drift and volatility, which satisfy:

$$\sigma_{RN} = \sigma, \quad \mu_{RN} = r - q - \frac{\sigma^2}{2}$$

Hence, one can use these equations to set (p_n, u_n, d_n) for a risk-neutral CRR tree.

Note that under this model,

$$\mathbb{E}S(T) = S_0 \exp(\mu_{RN}T + T\sigma_{RN}^2/2) = S_0 \exp(r - q)T$$

precisely as expected from a risk-neutral model.



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.