Math 5760/6890: Introduction to Mathematical Finance

A (rather non-technical) primer on Stochastic Processes

See Petters and Dong 2016, Section 6.1, 6.6

Akil Narayan[1]

[1] **Department of Mathematics, and Scientific Computing and Imaging (SCI) Institute**
University of Utah

Fall 2024

A brief intro to stochastic processes

In order to gain further understanding of our continuous-time binomial tree construction, we require some language of stochastic processes.

The high-level idea: we'll be formalizing random variables that are functions of time t.

In order to gain further understanding of our continuous-time binomial tree construction, we require some language of stochastic processes.

The high-level idea: we'll be formalizing random variables that are functions of time t.

To get there we'll go through some abstract formalization:

Definition

Let Ω be (probabilistic) event space, and let $\mathcal{I}$ be a set.
A stochastic process $X = X(t, \omega)$ is a map,

$$X : \mathcal{I} \times \Omega \to \mathbb{R}, \qquad X(t, \omega) \in \mathbb{R} \text{ for every } t \in \mathcal{I} \text{ and } \omega \in \Omega.$$

In order to gain further understanding of our continuous-time binomial tree construction, we require some language of stochastic processes.

The high-level idea: we'll be formalizing random variables that are functions of time t.

To get there we'll go through some abstract formalization:

Definition

Let Ω be (probabilistic) event space, and let $\mathcal{I}$ be a set.
A stochastic process $X = X(t, \omega)$ is a map,

$$X : \mathcal{I} \times \Omega \to \mathbb{R}, \qquad X(t, \omega) \in \mathbb{R} \text{ for every } t \in \mathcal{I} \text{ and } \omega \in \Omega.$$

The set $\mathcal{I}$ is called the *index set* of the process X, and $\mathbb{R}$ is the *state space*.

For fixed $t \in \mathcal{I}$, $X(t, \cdot)$ is a scalar random variable (e.g., and has a distribution).

For fixed $\omega \in \Omega$, $X(\cdot, \omega)$ is a deterministic scalar-valued function on the domain $\mathcal{I}$.

Example

Let $\boldsymbol{X} \in \mathbb{R}^N$ be a random vector. (Say the unknown time-1 price in Markowitz portfolio analysis.)

Then $\boldsymbol{X}$ is a stochastic process with index set $\mathcal{I} = \{1, 2, \ldots, N\}$.

Example

Let $\boldsymbol{X} \in \mathbb{R}^N$ be a random vector. (Say the unknown time-1 price in Markowitz portfolio analysis.)

Then $\boldsymbol{X}$ is a stochastic process with index set $\mathcal{I} = \{1, 2, \ldots, N\}$.

Example

Consider an n-period Binomial tree model with parameters $(p, u, d) = (p_n, u_n, d_n)$ fixed. (E.g., via a CRR model.)

The asset prices at every period can be lumped into a vector $\boldsymbol{S} := (S_0, S_1, \ldots, S_n)^T \in \mathbb{R}^{n+1}$.

Example

Let $\boldsymbol{X} \in \mathbb{R}^N$ be a random vector. (Say the unknown time-1 price in Markowitz portfolio analysis.)

Then $\boldsymbol{X}$ is a stochastic process with index set $\mathcal{I} = \{1, 2, \ldots, N\}$.

Example

Consider an n-period Binomial tree model with parameters $(p, u, d) = (p_n, u_n, d_n)$ fixed. (E.g., via a CRR model.)

The asset prices at every period can be lumped into a vector $\boldsymbol{S} := (S_0, S_1, \ldots, S_n)^T \in \mathbb{R}^{n+1}$.

Then $\boldsymbol{S}$ is a stochastic process with index set $\mathcal{I} = \{0, 1, \ldots, n\}$.

Continuum index sets

The previous examples had *discrete* index sets.

Of course, there is nothing stopping us from considering a continuous index set.

For example, if $S(t)$, $t \in [0, T]$ was the continuous-time model that we described from our $n \uparrow \infty$ limit of the CRR model, then S is a stochastic process with index set $\mathcal{I} = [0, T]$.

Continuum index sets

The previous examples had *discrete* index sets.

Of course, there is nothing stopping us from considering a continuous index set.

For example, if $S(t)$, $t \in [0, T]$ was the continuous-time model that we described from our $n \uparrow \infty$ limit of the CRR model, then S is a stochastic process with index set $\mathcal{I} = [0, T]$.

When we have an index set that's a continuum, we can discussing some standard notions of functional analysis.

Definition

A realization of a stochastic process $X(\cdot, \omega)$ (for a fixed $\omega \in \Omega$) with a continuum index set $\mathcal{I}$ is said to have a **continuous sample path** if

$$\lim_{s \to t} X(s, \omega) = X(\cancel{x}, \omega), \qquad \text{for every } t \in \mathcal{I}.$$

t

This notion of continuity involves a single, fixed ω.

A proper extension of continuity (or any other) property to "all ω" is more delicate.

Definition

Let X be a stochastic process with a continuum index set $\mathcal{I} = [0, \infty)$. Then X is **continuous** if

$$\widetilde{\Omega} := \left\{\omega \in \Omega \;\middle|\; X \text{ has a continuous sample path at } \omega\right\},$$

satisfies $P(\widetilde{\Omega}) = 1$.

Alternative languange: X is sample-path continuous, or almost surely continuous, or continuous with probability 1.

A proper extension of continuity (or any other) property to "all ω" is more delicate.

Definition

Let X be a stochastic process with a continuum index set $\mathcal{I} = [0, \infty)$. Then X is **continuous** if

$$\widetilde{\Omega} := \left\{\omega \in \Omega \;\middle|\; X \text{ has a continuous sample path at } \omega\right\},$$

satisfies $P(\widetilde{\Omega}) = 1$.

Alternative languange: X is sample-path continuous, or almost surely continuous, or continuous with probability 1.

Why not "for all ω"?

The mathematics of stochastic processes makes statements "with probability 1" the more natural statements to consider.

Asking for properties "for every ω" is too strong in the sense that asserting this limits our flexibility for analysis.

And since we really only care about probabilities of things, asking for statements in terms of probabilities is conceptually natural.

Our next major goal will be to identify and investigate a particular stochastic process that forms the foundation of mathematical finance.

To describe what kind of stochastic process we want, let's consider the log-return for our n-period discrete-time binomial tree model:

$$S_j = S(t_j), \qquad L_j = \log \frac{S_{j+1}}{S_j}, \qquad t_j = \cancel{jh_n}, \qquad h_n = \frac{T}{n}.$$

jh_j

Recall that we model L_j through a coin flip. More precisely,

$$L_j \sim \log d_n + \log \frac{u_n}{d_n} X, \qquad X \sim \text{Bernoulli}(p_n).$$

A stochastic process for securities

Our next major goal will be to identify and investigate a particular stochastic process that forms the foundation of mathematical finance.

To describe what kind of stochastic process we want, let's consider the log-return for our n-period discrete-time binomial tree model:

$$S_j = S(t_j), \qquad L_j = \log \frac{S_{j+1}}{S_j}, \qquad t_j = j h_n, \qquad h_n = \frac{T}{n}.$$

Recall that we model L_j through a coin flip. More precisely,

$$L_j \sim \log d_n + \log \frac{u_n}{d_n} X, \qquad X \sim \mathrm{Bernoulli}(p_n).$$

Suppose we take $n \uparrow \infty$, and we define $L(t)$ as the *cumulative* log-return from time 0:

$$L(T) \text{``}=\text{''} \log S_0 + \lim_{n\uparrow\infty} \sum_{j=1}^{n} L_j.$$

where $n \uparrow \infty$ affects the values of (p_n, u_n, d_n).

A stochastic process for securities

Our next major goal will be to identify and investigate a particular stochastic process that forms the foundation of mathematical finance.

To describe what kind of stochastic process we want, let's consider the log-return for our n-period discrete-time binomial tree model:

$$S_j = S(t_j), \qquad L_j = \log \frac{S_{j+1}}{S_j}, \qquad t_j = j h_n, \qquad h_n = \frac{T}{n}.$$

Recall that we model L_j through a coin flip. More precisely,

$$L_j \sim \log d_n + \log \frac{u_n}{d_n} X, \qquad X \sim \text{Bernoulli}(p_n).$$

Suppose we take $n \uparrow \infty$, and we define $L(t)$ as the *cumulative* log-return from time 0:

$$L(T) \text{``}=\text{''} \log S_0 + \lim_{n \uparrow \infty} \sum_{j=1}^{n} L_j.$$

where $n \uparrow \infty$ affects the values of (p_n, u_n, d_n).

This in principle defines a stochastic process $L(t)$ with index set $\mathcal{I} = [0, T]$. What properties should this process have?

A more explicit construction

Consider our real-world CRR tree, which assigns (p_n, u_n, d_n) as,

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

for given drift and volatility (μ, σ). These numbers affect the distribution of L_j.

A more explicit construction

Consider our real-world CRR tree, which assigns (p_n, u_n, d_n) as,

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

for given drift and volatility (μ, σ). These numbers affect the distribution of L_j.

Let's simplify things a bit: let's consider a particular real-world model with $(\mu, \sigma) = (0, 1)$.

Under this assumption, then

$$L_j = \begin{cases} \sqrt{h_n}, & \text{with probability } \frac{1}{2}, \\ -\sqrt{h_n}, & \text{with probability } \frac{1}{2}, \end{cases}$$

I.e., the cumulative sum of the L_j corresponds to a symmetric random walk.

A more explicit construction

Consider our real-world CRR tree, which assigns (p_n, u_n, d_n) as,

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

for given drift and volatility (μ, σ). These numbers affect the distribution of L_j.

Let's simplify things a bit: let's consider a particular real-world model with $(\mu, \sigma) = (0, 1)$.

Under this assumption, then

$$L_j = \begin{cases} \sqrt{h_n}, & \text{with probability } \frac{1}{2}, \\ -\sqrt{h_n}, & \text{with probability } \frac{1}{2}, \end{cases}$$

I.e., the cumulative sum of the L_j corresponds to a symmetric random walk.

The goal now is to identify/construct a stochastic process $L(t)$ that:

- is consistent with what happens to $\log(S_n/S_0)$ when we take $n \uparrow \infty$
- has the properties we desire from the finance perspective

We want the following things in a putative continuous-time stochastic process $L(t)$:

- We want $L(t)$ to be distributed like $\mathcal{N}(0, t)$ for all t.

Recall: $S(t) \sim \text{lognormal}\left(\mu t + \frac{\sigma^2}{2} t, \sigma^2 t\right)$ $\quad (S_0 = 1)$

$(\mu, \sigma) = (0, 1) \Rightarrow S(t) \sim \text{lognormal}(0, t)$

$\log\left(\frac{S(t)}{S_0}\right) = L(t) \sim \mathcal{N}(0, t)$

We want the following things in a putative continuous-time stochastic process $L(t)$:

- We want $L(t)$ to be distributed like $\mathcal{N}(0, t)$ for all t.
- We want non-overlapping *increments*, like $L(3) - L(2)$ and $L(2) - L(1)$, to have independent distributions.

Coin flips are all independent: coin flips are discrete-time increments.

So the "coin flips" determining $L(t_2) - L(t_1)$ should be independent of the "coin flips" determining $L(t_4) - L(t_3)$ so long as (t_1, t_2) and (t_3, t_4) are disjoint.

We want the following things in a putative continuous-time stochastic process $L(t)$:

- We want $L(t)$ to be distributed like $\mathcal{N}(0, t)$ for all t.
- We want non-overlapping *increments*, like $L(3) - L(2)$ and $L(2) - L(1)$, to have independent distributions.
- If we know that $L(t) = L_0$, then we want the *increments*, $L(t + \epsilon) - L(t)$, to be distributed like $\mathcal{N}(0, \epsilon)$.

I.e., starting at time t and "flipping coins" until time $t+\varepsilon$ should correspond to a (log) variance of $(t+\varepsilon)-t = \varepsilon$.

Desiderata

We want the following things in a putative continuous-time stochastic process $L(t)$:

- We want $L(t)$ to be distributed like $\mathcal{N}(0, t)$ for all t.
- We want non-overlapping *increments*, like $L(3) - L(2)$ and $L(2) - L(1)$, to have independent distributions.
- If we know that $L(t) = L_0$, then we want the *increments*, $L(t + \epsilon) - L(t)$, to be distributed like $\mathcal{N}(0, \epsilon)$.
- We want Markovian structure: If we assume the knowledge $L(t) = L_0$, then our understanding of the properties and distribution of $L(s)$ for $s > t$ is the same as if we had knowledge of $L(r)$ for all $r \leqslant t$.

Desiderata

We want the following things in a putative continuous-time stochastic process $L(t)$:

- We want $L(t)$ to be distributed like $\mathcal{N}(0, t)$ for all t.
- We want non-overlapping *increments*, like $L(3) - L(2)$ and $L(2) - L(1)$, to have independent distributions.
- If we know that $L(t) = L_0$, then we want the *increments*, $L(t + \epsilon) - L(t)$, to be distributed like $\mathcal{N}(0, \epsilon)$.
- We want Markovian structure: If we assume the knowledge $L(t) = L_0$, then our understanding of the properties and distribution of $L(s)$ for $s > t$ is the same as if we had knowledge of $L(r)$ for all $r \leqslant t$.

There is essentially just one process satisfying these conditions: Brownian motion.

Definition

With index set $\mathcal{I} = [0, \infty)$, a **standard Brownian motion/Wiener process** $B = B_t = B(t)$ is a stochastic process satisfying

1. $P(B(0) = 0) = 1$
2. B is continuous with probability 1
3. The n sequential increments formed by any choice of $n + 1$ ordered time points $t_1, \ldots, t_{n+1}$ are mutually independent
4. For any $0 \leqslant s \leqslant t < \infty$, then $B(t) - B(s) \sim \mathcal{N}(0, t - s)$.

I.e., $B(t_2) - B(t_1)$ is independent of $B(t_3) - B(t_2)$, which is ind. of $B(t_4) - B(t_3)$ .. etc.

$(t_1 < t_2 < \cdots < t_{n+1})$

Definition

With index set $\mathcal{I} = [0, \infty)$, a **standard Brownian motion/Wiener process** $B = B_t = B(t)$ is a stochastic process satisfying

1. $P(B(0) = 0) = 1$
2. B is continuous with probability 1
3. The n sequential increments formed by any choice of $n+1$ ordered time points $t_1, \ldots, t_{n+1}$ are mutually independent
4. For any $0 \leqslant s \leqslant t < \infty$, then $B(t) - B(s) \sim \mathcal{N}(0, t-s)$.

Some initial remarks:

– The reason we use indefinite articles (a Brownian motion) is the same reason we use indefinite articles to describe random variables (X is a normal random variable).

Definition

With index set $\mathcal{I} = [0, \infty)$, a **standard Brownian motion/Wiener process** $B = B_t = B(t)$ is a stochastic process satisfying

1. $P(B(0) = 0) = 1$
2. B is continuous with probability 1
3. The n sequential increments formed by any choice of $n+1$ ordered time points $t_1, \ldots, t_{n+1}$ are mutually independent
4. For any $0 \leqslant s \leqslant t < \infty$, then $B(t) - B(s) \sim \mathcal{N}(0, t-s)$.

Some initial remarks:

- The reason we use indefinite articles (a Brownian motion) is the same reason we use indefinite articles to describe random variables (X is a normal random variable).
- The last property implies that for every $s, t, h \geqslant 0$:

$$B(t+h) - B(t) \sim B(s+h) - B(s) \sim \mathcal{N}(0, h)$$

Definition

With index set $\mathcal{I} = [0, \infty)$, a **standard Brownian motion/Wiener process** $B = B_t = B(t)$ is a stochastic process satisfying

1. $P(B(0) = 0) = 1$
2. B is continuous with probability 1
3. The n sequential increments formed by any choice of $n+1$ ordered time points $t_1, \ldots, t_{n+1}$ are mutually independent
4. For any $0 \leqslant s \leqslant t < \infty$, then $B(t) - B(s) \sim \mathcal{N}(0, t-s)$.

Some initial remarks:

- The reason we use indefinite articles (a Brownian motion) is the same reason we use indefinite articles to describe random variables (X is a normal random variable).
- The last property implies that for every $s, t, h \geqslant 0$:

$$B(t+h) - B(t) \sim B(s+h) - B(s) \sim \mathcal{N}(0, h)$$

- The four properties above are typically concisely referred to, respectively, as: $B(0) = \overset{0}{\cancel{1}}$ with probability 1, B is sample-continuous with probability 1, B has independent increments, and B has (time-)stationary and normally-distributed increments.

Let $B(t)$ be a standard Brownian motion, and let $b_0 \in \mathbb{R}$, $\mu \in \mathbb{R}$, and $\sigma > 0$. Then the process

$$A(t) = b_0 + \mu t + \sigma B(t),$$

is a Brownian motion with drift and scaling:

- It has time-0 value b_0 with probability 1.
- It has deterministic drift μt
- It has scaling σ

$$\mathbb{E}[A(t) - A(0)] = \mathbb{E}[\, b_0 + \mu t + \sigma B(t) - b_0 \,]$$

$$= \mathbb{E}[\mu t + \sigma B(t)] = \mu t$$

$$\text{Var}\, A(t) = \sigma^2 \text{Var}\, B(t) = \sigma^2 t$$

Let $B(t)$ be a standard Brownian motion, and let $b_0 \in \mathbb{R}$, $\mu \in \mathbb{R}$, and $\sigma > 0$. Then the process

$$A(t) = b_0 + \mu t + \sigma B(t),$$

is a Brownian motion with drift and scaling:

- It has time-0 value b_0 with probability 1.
- It has deterministic drift μt
- It has scaling σ

Example

What distribution do the increments of $A(t)$ have?

$\mathbb{E}[A(t) - A(0)] = \mu t$ $\quad$ $\mathrm{Var}[A(t)] = \sigma^2 t$

$A(t)$ is normally distributed $\rightarrow$ $A(t) - A(0) \sim N(\mu t, \sigma^2 t)$

Other properties

Brownian motion is a *fascinating* object:

- With probability 1, the sample path $B(\cdot, \omega)$ is continuous in time.

One trajectory:

Brownian motion is a *fascinating* object:

- With probability 1, the sample path $B(\cdot, \omega)$ is continuous in time.
- With probability 1, the sample path $B(\cdot, \omega)$ is differentiable <u>nowhere</u>

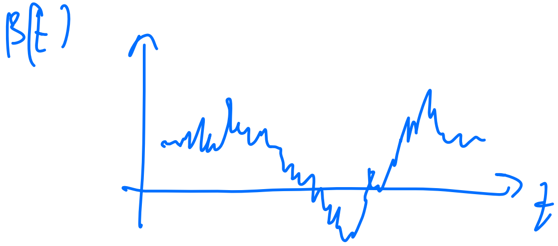

Other properties

Brownian motion is a *fascinating* object:

- With probability 1, the sample path $B(\cdot, \omega)$ is continuous in time.
- With probability 1, the sample path $B(\cdot, \omega)$ is differentiable nowhere
- For any interval $I \subset [0, \infty)$ of finite length, with probability 1, the sample path $B(\cdot, \omega)$ has infinite "variation" on I.

Brownian motion is a *fascinating* object:

- With probability 1, the sample path $B(\cdot, \omega)$ is continuous in time.
- With probability 1, the sample path $B(\cdot, \omega)$ is differentiable nowhere
- For any interval $I \subset [0, \infty)$ of finite length, with probability 1, the sample path $B(\cdot, \omega)$ has infinite "variation" on I.
- B has Markovian structure: Fix $s > 0$ and define $A(t) := B(t + s) - B(s)$. Then $A(t)$ and $B(t)$ have the same distribution, and in particular A is a standard Brownian motion.

Other properties

Brownian motion is a *fascinating* object:

- With probability 1, the sample path $B(\cdot, \omega)$ is continuous in time.
- With probability 1, the sample path $B(\cdot, \omega)$ is differentiable nowhere
- For any interval $I \subset [0, \infty)$ of finite length, with probability 1, the sample path $B(\cdot, \omega)$ has infinite "variation" on I.
- B has Markovian structure: Fix $s > 0$ and define $A(t) := B(t + s) - B(s)$. Then $A(t)$ and $B(t)$ have the same distribution, and in particular A is a standard Brownian motion.
- Sample paths of Brownian motion are self-similar/fractals. In particular, for any $c > 0$, the process $\frac{1}{c} B(c^2 t)$ is a standard Brownian motion.

Self-similarity of Brownian paths

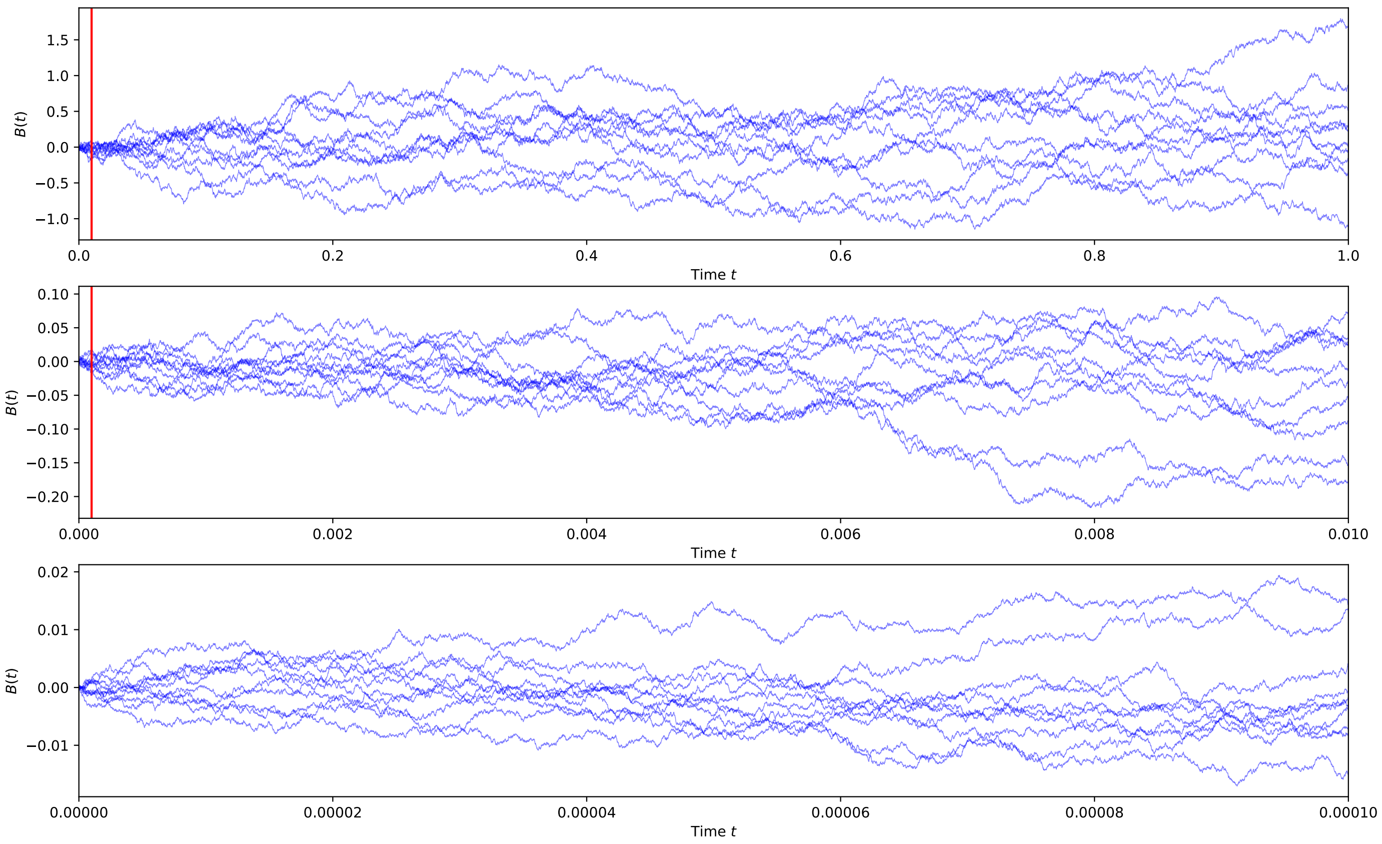

In relation to securities....

To close the loop: a standard Brownian motion $B(t)$ is exactly a continuous-time process whose properties line up with $L(t)$, our continuous-time limit of the binomial tree model.

In particular, if $(\mu, \sigma, S_0) = (0, 1, 1)$, then we will identify

$$L(t) \text{``} = \text{''} B(t).$$

In relation to securities....

To close the loop: a standard Brownian motion $B(t)$ is exactly a continuous-time process whose properties line up with $L(t)$, our continuous-time limit of the binomial tree model.

In particular, if $(\mu, \sigma, S_0) = (0, 1, 1)$, then we will identify

$$L(t) \text{``}=\text{''} B(t).$$

As one might expect, for $\mu \neq 0$ and $\sigma \neq 1$, then the appropriate identification is,

$$L(t) \text{``}=\text{''} \mu t + \sigma B(t).$$

To close the loop: a standard Brownian motion $B(t)$ is exactly a continuous-time process whose properties line up with $L(t)$, our continuous-time limit of the binomial tree model.

In particular, if $(\mu, \sigma, S_0) = (0, 1, 1)$, then we will identify

$$L(t) \text{``}=\text{''} B(t).$$

As one might expect, for $\mu \neq 0$ and $\sigma \neq 1$, then the appropriate identification is,

$$L(t) \text{``}=\text{''} \mu t + \sigma B(t).$$

While in principle we are done in terms of identifying a mathematical model for $L(t)$, we have actually just begun to reap benefits from this model.

In particular, the fact that $B(t)$ has sample paths or *trajectories* suggests that there is an underlying time-evolution law.

Stochastic calculus is the appropriate language we'll use to explore such concepts.

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.