Lorenzo Naranjo October 15, 2026

Modeling Stock Prices in Continuous-Time

The binomial model of Multiple Periods moves the stock price up by a factor $u = e^{\sigma\sqrt{\Delta t}}$ or down by $d = 1/u$ each period, and its prices approach the Black-Scholes prices as the periods get shorter. This chapter takes that limit. We describe the random part of the stock price with a *Brownian motion*, the continuous-time limit of a random walk, and model the stock price as a *geometric Brownian motion*. We then introduce *Itô's lemma*, the chain rule for functions of a Brownian motion, which is the one tool from stochastic calculus that the Black-Scholes model needs. With it, we find the distribution of the future stock price, which is what the Black-Scholes formula needs.

The chapter assumes no background in stochastic processes, and it aims for intuition rather than proofs. The notes on Stochastic Foundations for Finance develop the same material with full arguments.

From Binomial Trees to Brownian Motion

The Log Price in a Binomial Tree

In a binomial tree with $u = e^{\sigma\sqrt{\Delta t}}$ and $d = e^{-\sigma\sqrt{\Delta t}}$, each period adds $\sigma\sqrt{\Delta t}$ to the logarithm of the stock price or subtracts it. After m periods, at time $t = m\Delta t$,

$$\ln(S_t) = \ln(S_0) + \sigma\sqrt{\Delta t}\,(X_1 + X_2 + \cdots + X_m),$$

where each X_i equals $+1$ for an up move and -1 for a down move. If up and down moves are equally likely and independent across periods, each X_i has mean zero and variance one, and

$$\mathrm{V}\,(\ln(S_t)) = \sigma^2\Delta t \times m = \sigma^2 t.$$

The variance of the log price grows in proportion to time, at the rate σ^2, however short the periods are. This is why the tree uses steps of size $\sigma\sqrt{\Delta t}$: with steps proportional to

Δt, the variance would be $\sigma^2 t\Delta t$ and the stock price would stop moving as $\Delta t \to 0$, while with steps that do not shrink, the variance would explode.

The Scaled Random Walk

Setting σ aside, the random part of the log price is the *scaled random walk*

$$B_t^{(n)} = \sqrt{\Delta t}\,(X_1 + X_2 + \cdots + X_m), \qquad t = m\Delta t,$$

where $n = 1/\Delta t$ is the number of steps per unit of time. Its mean is zero and its variance is $m\Delta t = t$. Figure 1 shows one path of the walk for three values of n. With ten steps the path is a coarse zigzag. With ten thousand it looks like a continuous but very rough curve, and its range barely changes, since the variance at $t = 1$ is one for every n.

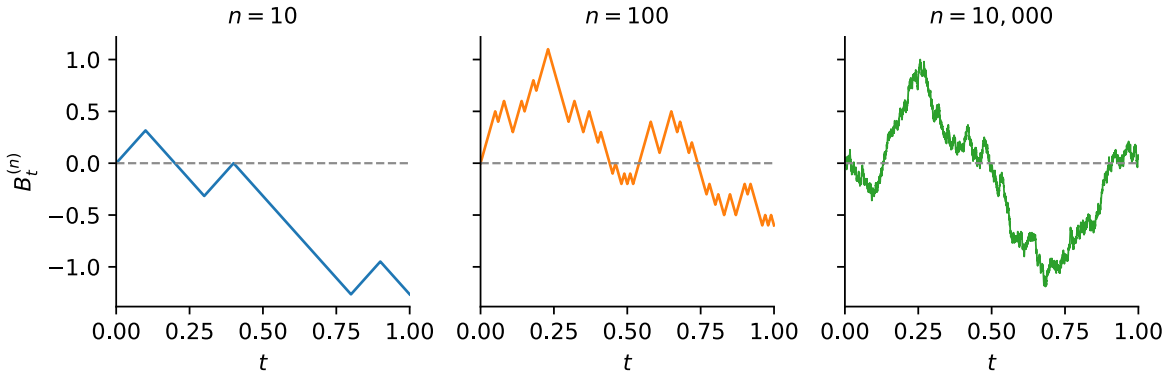


Figure 1: The figure plots one path of the scaled random walk over $0 \le t \le 1$ with $n = 10, n = 100,$ and $n = 10{,}000$ steps per unit of time. Each step moves up or down by $\sqrt{\Delta t} = 1/\sqrt{n}$.

Brownian Motion

As $\Delta t \to 0$, the value of the walk at time t is a sum of more and more independent steps, each smaller than the last. By the central limit theorem, its distribution approaches a normal distribution with mean zero and variance t. The same is true of the change of the walk over any interval $(s, t]$, which is built from the steps taken during that interval and is independent of the steps taken before. The limit of the scaled random walk is called a *Brownian motion*, or a *Wiener process*.

Brownian Motion

A *Brownian motion* $\{B_t : t \geq 0\}$ is a stochastic process with the following properties:

i. $B_0 = 0$.
ii. For $s < t$, the increment $B_t - B_s$ is normally distributed with mean zero and variance $t - s$.
iii. Increments over non-overlapping intervals are independent.
iv. The path $t \mapsto B_t$ is continuous.

In particular, $B_t \sim \mathcal{N}(0, t)$, so that $B_t = \sqrt{t}\, Z$ for a standard normal Z. The increments are normal because of the central limit theorem, not because the coin flips were normal: any steps with mean zero and variance one would lead to the same limit.

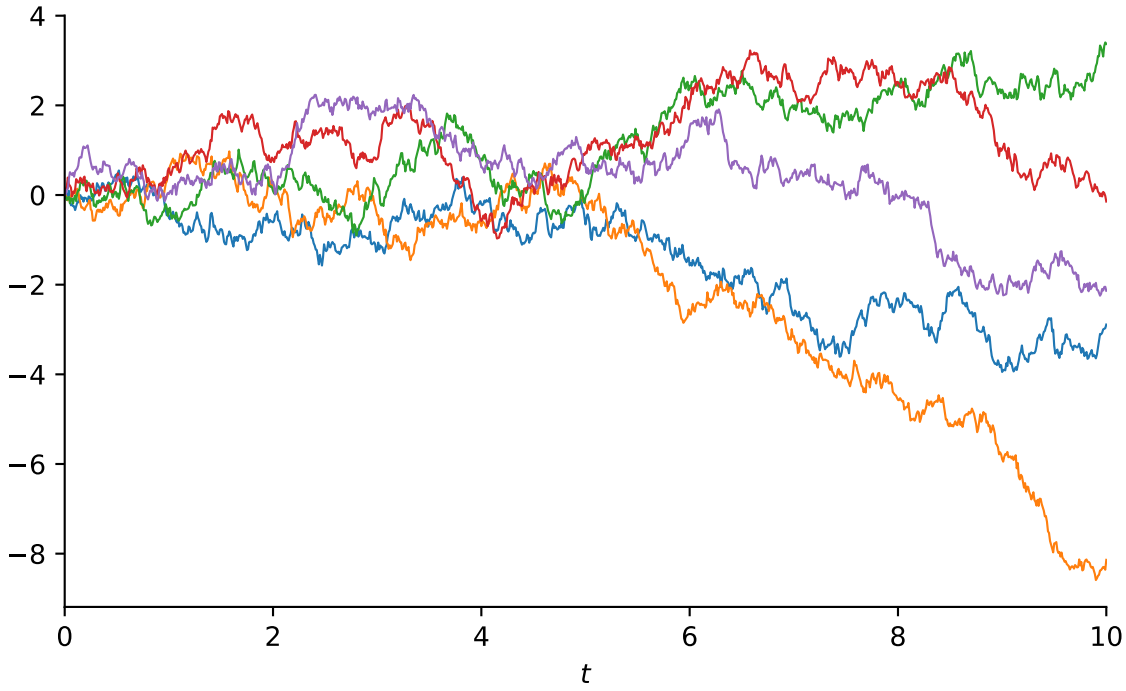


Figure 2: The figure plots simulated paths of a Brownian motion $\{B_t\}$ for $t \in [0, 10]$.

The paths in Figure 2 are continuous but nowhere smooth. Over a short interval of length dt, the increment dB has mean zero and standard deviation $\sqrt{dt}$, which is much larger than dt when dt is small. A Brownian path therefore moves a lot over every short interval,

in one direction or the other. This single fact explains why calculus with Brownian motion differs from ordinary calculus, as we show below.

Geometric Brownian Motion

A Brownian motion is not a good model for the price itself. It can become negative, and its changes do not depend on its level, while a stock trading at $200 typically moves by about twice as many dollars in a day as the same stock trading at $100. What is roughly stable over time is the distribution of *returns*, the percentage changes in price. We therefore let returns, not price changes, be driven by a Brownian motion. Over a short interval of length Δt, the return on the stock is

$$\frac{\Delta S_t}{S_t} = \mu \Delta t + \sigma \Delta B_t,$$

which is normally distributed with mean $\mu \Delta t$ and variance $\sigma^2 \Delta t$. Letting $\Delta t \to 0$ gives

$$dS = \mu S\, dt + \sigma S\, dB. \tag{1}$$

The process $\{S_t\}$ is called a *geometric Brownian motion* (GBM). The constant μ is the *expected return* of the stock per unit of time, and σ is the *volatility* of its returns, the same σ used to build binomial trees. Both are usually quoted per year. Equation (1) is shorthand for the limit of the discrete returns above, and it says that over each short interval the stock earns a deterministic return $\mu\, dt$ plus a random return $\sigma\, dB$.

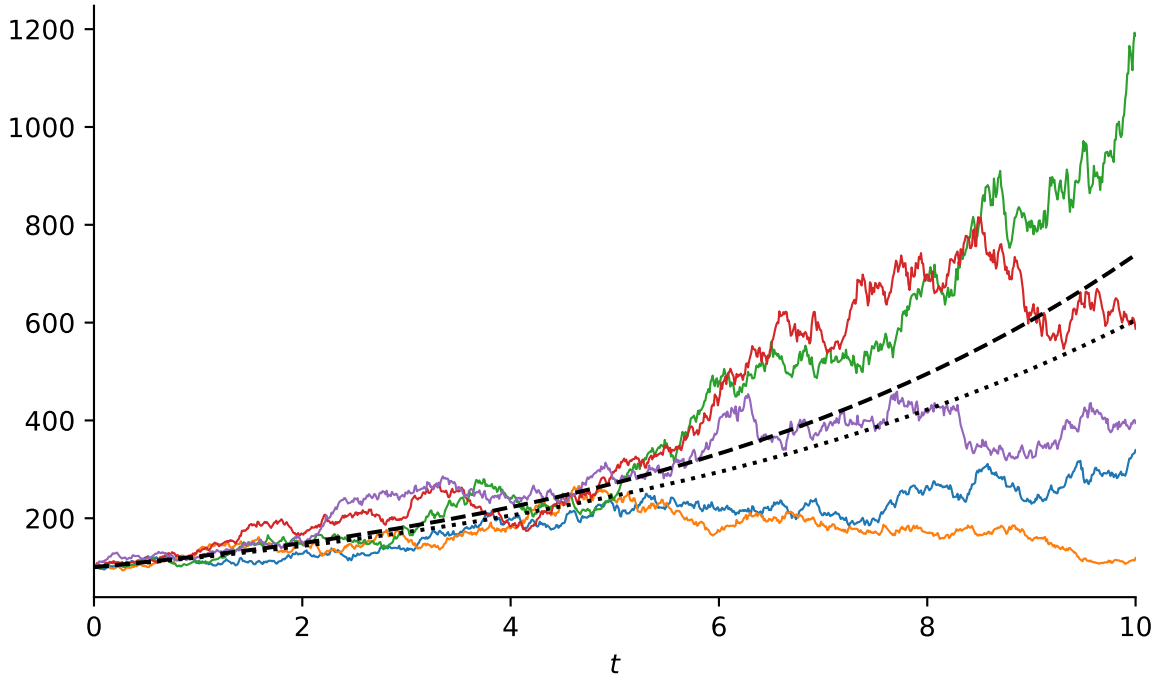


Figure 3: The figure plots simulated paths of a geometric Brownian motion $\{S_t\}$ for $t \in [0, 10]$, with $S_0 = 100$, $\mu = 0.20$, and $\sigma = 0.20$. The dashed line is the expected price $\mathbb{E}(S_t) = S_0 e^{\mu t}$ and the dotted line is the median price $S_0 e^{(\mu - \frac{1}{2}\sigma^2)t}$.

Figure 3 shows that the paths never fall below zero, that their fluctuations grow with the price level, and that they fan out unevenly, with a few ending far above the rest. To make these statements precise, we need the distribution of S_T, and to find it we need a rule for how functions of S_t, such as $\ln(S_t)$, change over time.

Itô's Lemma

Why Ordinary Calculus Fails

In ordinary calculus, if x changes by a small amount Δx, a smooth function $F(x)$ changes by

$$\Delta F \approx F'(x)\,\Delta x + \frac{1}{2}F''(x)(\Delta x)^2,$$

and the second-order term is negligible because $(\Delta x)^2$ is tiny compared with Δx. For a Brownian motion, ΔB is of order $\sqrt{\Delta t}$, so $(\Delta B)^2$ is of order Δt, the same order as the time step itself. The second-order term no longer disappears.

More precisely, $(\Delta B)^2 = \Delta t\, Z^2$ for a standard normal Z, so it has mean Δt and variance $2(\Delta t)^2$. Its standard deviation, $\sqrt{2}\,\Delta t$, is as large as its mean, so a single squared increment is far from Δt. But split an interval $[0, t]$ into n steps of length $\Delta t = t/n$ and add up the squared increments of a Brownian path. The squared increments are independent, so their sum has mean $n\Delta t = t$ and variance $n \times 2(\Delta t)^2 = 2t\,\Delta t$, which goes to zero as $\Delta t \to 0$. The sum of squared increments converges to t, on every path. Figure 4 shows this for one simulated path on $[0, 1]$.

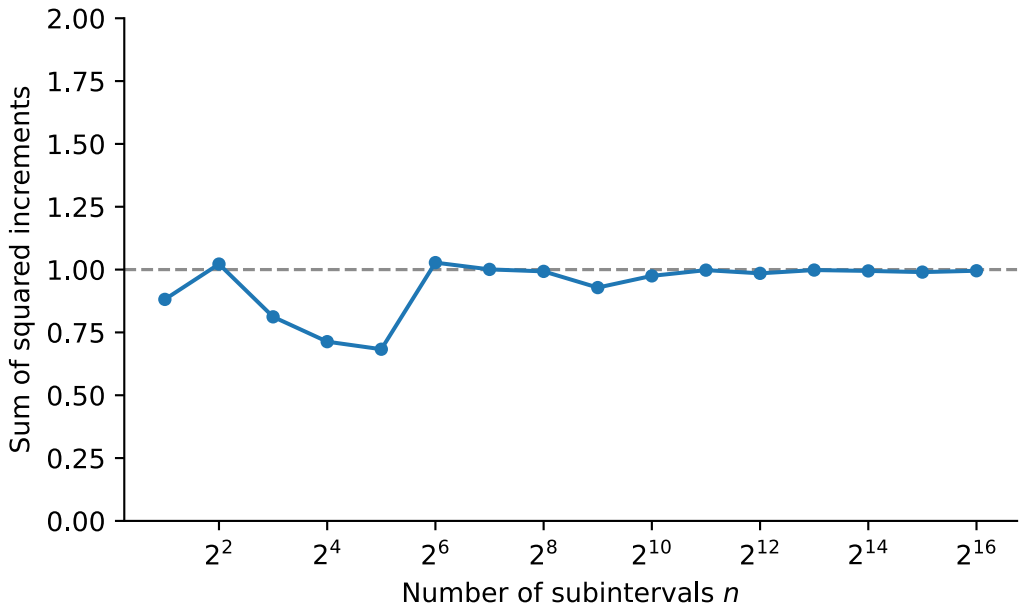


Figure 4: The figure adds up the squared increments of one simulated Brownian path on $[0, 1]$, using $n = 2, 4, \ldots, 2^{16}$ equal subintervals. As the subintervals shrink, the sum settles on $t = 1$, the dashed line.

In differential notation, this result reads $(dB)^2 = dt$. The squared increments of a Brownian motion add up like increments of time, so in any calculation that adds up many small changes, $(dB)^2$ can be replaced by dt. The products $(dt)^2$ and $(dt)(dB)$ are of order $(dt)^2$ and $(dt)^{3/2}$, smaller than dt, and they add up to zero. Together, these rules are called the *multiplication table* or *box calculus*.

Multiplication Table

The products of differentials of t and of a Brownian motion B are

$$(dB)^2 = dt, \qquad (dt)(dB) = 0, \qquad (dt)^2 = 0.$$

For a GBM, the multiplication table gives

$$(dS)^2 = (\mu S\, dt + \sigma S\, dB)^2 = \mu^2 S^2 (dt)^2 + 2\mu\sigma S^2 (dt)(dB) + \sigma^2 S^2 (dB)^2 = \sigma^2 S^2\, dt.$$

The drift does not contribute to $(dS)^2$. Only the random term, which moves the price by amounts of order $\sqrt{dt}$, does.

The Lemma

Expanding a function $F(S,t)$ to second order in dS and first order in dt, and keeping the term $(dS)^2 = \sigma^2 S^2\, dt$ that ordinary calculus would drop, gives Itô's lemma.

Itô's Lemma

Let $\{S_t\}$ follow the GBM in (1), and let $F(S,t)$ be a function that is twice differentiable in S and differentiable in t. Then

$$dF = \frac{\partial F}{\partial S}\, dS + \frac{1}{2}\frac{\partial^2 F}{\partial S^2}\,(dS)^2 + \frac{\partial F}{\partial t}\, dt,$$

where $(dS)^2 = \sigma^2 S^2\, dt$. Substituting dS,

$$dF = \left(\mu S \frac{\partial F}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 F}{\partial S^2} + \frac{\partial F}{\partial t}\right) dt + \sigma S \frac{\partial F}{\partial S}\, dB. \qquad (2)$$

Sketch of proof

A second-order Taylor expansion of F gives

$$dF = \frac{\partial F}{\partial S} dS + \frac{\partial F}{\partial t} dt + \frac{1}{2}\frac{\partial^2 F}{\partial S^2}(dS)^2 + \frac{\partial^2 F}{\partial S \partial t}(dS)(dt) + \frac{1}{2}\frac{\partial^2 F}{\partial t^2}(dt)^2.$$

By the multiplication table, $(dS)^2 = \sigma^2 S^2 dt$, while $(dS)(dt) = \mu S(dt)^2 + \sigma S(dB)(dt) = 0$ and $(dt)^2 = 0$. Higher-order terms involve $(dS)^3$ and higher powers, which are of order $(dt)^{3/2}$ or smaller and also vanish. What is left is Itô's lemma. □

The term $\frac{1}{2}\sigma^2 S^2 \frac{\partial^2 F}{\partial S^2}$ is the *Itô correction*. It has no counterpart in ordinary calculus, and it vanishes when there is no randomness ($\sigma = 0$) or when F is linear in S. When F depends only on S, the last term drops out. Equation (2) also shows that F is driven by the same Brownian motion as the stock, with a random term $\sigma S \frac{\partial F}{\partial S} dB$ that is proportional to the sensitivity of F to the stock price. This common source of risk is what allows an option to be replicated with the stock in the Black-Scholes model.

Example 1. Let $F(S,t) = e^{-rt}S$ be the stock price discounted at the risk-free rate r. Then $\frac{\partial F}{\partial S} = e^{-rt}$, $\frac{\partial^2 F}{\partial S^2} = 0$, and $\frac{\partial F}{\partial t} = -re^{-rt}S$, so that

$$dF = \left(\mu e^{-rt}S - re^{-rt}S\right)dt + \sigma e^{-rt}S\, dB = (\mu - r)F\, dt + \sigma F\, dB.$$

The discounted stock price is again a GBM, with expected return $\mu - r$, the risk premium of the stock, and the same volatility. There is no Itô correction because F is linear in S.□

Solving for the Stock Price

The logarithm turns multiplicative changes into additive ones, so we apply Itô's lemma to $F(S) = \ln(S)$. With $\frac{\partial F}{\partial S} = 1/S$ and $\frac{\partial^2 F}{\partial S^2} = -1/S^2$,

$$\begin{aligned} d\ln(S) &= \frac{1}{S}\, dS + \frac{1}{2}\left(-\frac{1}{S^2}\right)(dS)^2 \\ &= \frac{1}{S}(\mu S\, dt + \sigma S\, dB) - \frac{1}{2}\frac{1}{S^2}\sigma^2 S^2\, dt \\ &= \left(\mu - \frac{1}{2}\sigma^2\right)dt + \sigma\, dB. \end{aligned}$$

The right-hand side has constant coefficients, so the changes of $\ln(S)$ add up to

$$\ln(S_T) - \ln(S_0) = \left(\mu - \frac{1}{2}\sigma^2\right)T + \sigma B_T,$$

and taking exponentials,

$$S_T = S_0 \exp\left(\left(\mu - \frac{1}{2}\sigma^2\right)T + \sigma B_T\right). \tag{3}$$

The log price is a Brownian motion with drift $\mu - \frac{1}{2}\sigma^2$ and volatility σ, the continuous-time version of the log price in a binomial tree. The term $-\frac{1}{2}\sigma^2$ is the Itô correction. Ordinary calculus would give $d\ln(S) = dS/S$, and a log price that grows at the rate μ.

The Distribution of the Stock Price

Since $B_T \sim \mathcal{N}(0, T)$, equation (3) implies that the log price is normally distributed,

$$\ln(S_T) \sim \mathcal{N}(m, s^2), \qquad m = \ln(S_0) + \left(\mu - \frac{1}{2}\sigma^2\right)T, \qquad s = \sigma\sqrt{T}.$$

In other words, S_T has a lognormal distribution, as defined in The Lognormal Distribution. The continuously compounded return $\ln(S_T/S_0)$ over a horizon T is normal, with mean $\left(\mu - \frac{1}{2}\sigma^2\right)T$ and standard deviation $\sigma\sqrt{T}$. Probabilities about the stock price are computed by translating them into statements about this normal return.

Example 2. A stock that follows a GBM has an expected return of 12% per year and a volatility of 25% per year, and its current price is \$25. Over $T = 0.5$ years, the log price $X_T = \ln(S_T)$ has

$$\mathrm{E}(X_T) = \ln(25) + \left(0.12 - 0.5(0.25)^2\right)(0.5) = 3.2633,$$
$$\mathrm{SD}(X_T) = 0.25\sqrt{0.5} = 0.1768.$$

The 95% probability interval for X_T is $3.2633 \pm 1.96(0.1768)$, and exponentiating its endpoints gives

$$\left[e^{3.2633-1.96(0.1768)}, e^{3.2633+1.96(0.1768)}\right] = [18.48, 36.96].$$

There is a 95% probability that the stock price in six months will lie between \$18.48 and \$36.96. The interval is not symmetric around the current price, because the lognormal distribution is skewed to the right. □

Moments of the Stock Price

The moments of a lognormal random variable from The Lognormal Distribution give the mean and standard deviation of S_T.

Moments of the Stock Price

If $\{S_t\}$ follows the GBM in (1), then

$$\mathrm{E}(S_T) = S_0 e^{\mu T}, \qquad \mathrm{SD}(S_T) = \mathrm{E}(S_T)\sqrt{e^{\sigma^2 T} - 1}.$$

Proof

Since $\ln(S_T) \sim \mathcal{N}(m, s^2)$,

$$\mathrm{E}(S_T) = e^{m+\frac{1}{2}s^2} = e^{\ln(S_0)+\left(\mu-\frac{1}{2}\sigma^2\right)T+\frac{1}{2}\sigma^2 T} = S_0 e^{\mu T},$$

and $\mathrm{SD}(S_T) = \mathrm{E}(S_T)\sqrt{e^{s^2} - 1}$ with $s^2 = \sigma^2 T$. □

The expected stock price grows at the rate μ, which is why μ is called the expected return. The standard deviation is proportional to the expected price and grows even faster, so the uncertainty about the future price grows much faster than the expected price.

Example 3. For the stock in Example 2, the expected price and the standard deviation of the price six months from now are

$$\mathrm{E}(S_T) = 25e^{0.12(0.5)} = \$26.55,$$

$$\mathrm{SD}(S_T) = 26.55\sqrt{e^{0.25^2(0.5)} - 1} = \$4.73.$$

□

Expected Return and Growth Rate

The log price grows at the rate $\mu - \frac{1}{2}\sigma^2$, while the expected price grows at the faster rate μ. Both statements are correct, and the gap between them comes from the skewness of the lognormal distribution. The median of S_T is the price at which the log price is at its mean,

$$\text{Median}(S_T) = S_0 e^{\left(\mu - \frac{1}{2}\sigma^2\right)T} < S_0 e^{\mu T} = \text{E}(S_T).$$

Half of the paths end below the median, and more than half end below the mean. A few paths that end very high pull the mean above what a typical path achieves, as the dashed and dotted lines in Figure 3 show. The rate $\mu - \frac{1}{2}\sigma^2$ is the growth rate that an investor holding the stock for a long time will typically realize, while μ is the rate at which wealth grows on average across all paths.

Example 4. A stock has an expected return of $\mu = 0.10$ and a volatility of $\sigma = 0.20$. Its log price grows at the rate $0.10 - 0.5(0.20)^2 = 0.08$ per year. Over $T = 10$ years, the expected price is $e^{0.10\times 10} = 2.72$ times the current price, but the median price is only $e^{0.08\times 10} = 2.23$ times the current price. The stock ends below its expected price when $\sigma B_T < \frac{1}{2}\sigma^2 T$, which has probability

$$\text{P}\left(Z < \frac{1}{2}\sigma\sqrt{T}\right) = \Phi(0.3162) = 0.6241,$$

where Z is a standard normal random variable. □

The Black-Scholes model prices options with this lognormal distribution, with one change: the expected return μ is replaced by the risk-free rate r. The Black-Scholes Model explains why.

Practice Problems

Solutions to all problems can be found at lorenzonaranjo.com/class-materials/derivatives/.

Problem 1. Consider a stock whose price at time t is given by S_t and that follows a geometric Brownian motion (GBM). The expected return is 18% per year and the volatility is 32% per year. The current spot price is $60.

a. Compute the expected price 9 months from now.
b. Compute the mean and standard deviation of the log-spot price 9 months from now.
c. Compute the 95% probability interval of $\ln(S_T)$ 9 months from now, and report the corresponding values for S_T.

Problem 2. Consider a stock whose price at time t is given by S_t and that follows a GBM. The expected return is 11% per year and the volatility is 27% per year. The current spot price is \$60.

a. Compute the expected price of S_t 1 year from now.
b. Compute the expected price of $1/S_t$ 1 year from now.

Problem 3. Consider a stock whose price at time t is given by S_t and that follows a GBM. The expected return is 12% per year and the volatility is 35% per year. The current spot price is \$55. Let $T = 18$ months.

a. Compute $\mathrm{E}(S_T)$.
b. Compute the mean and standard deviation of the log-spot price at T.
c. Find C such that $\mathrm{P}(S_T \leq C) = 0.01$.

Problem 4. Consider a stock whose price at time T is given by S_T and that follows a GBM, i.e.,

$$\ln(S_T) \sim \mathcal{N}(\ln(S_0) + (\mu - 0.5\sigma^2)T, \sigma^2 T).$$

The expected return is 12% per year and the volatility is 35% per year. The current spot price is \$100.

a. Compute the expected price in 2 years from now.
b. Compute the mean and standard deviation of the log-spot price in 2 years from now.
c. Compute the probability that the spot price is less than \$100 in 2 years from now.
d. Compute the probability that the spot price is greater than \$120 in 2 years from now.

Problem 5. Suppose that the stock price follows a geometric Brownian motion (GBM) with drift μ and instantaneous volatility σ, i.e.,

$$dS = \mu S dt + \sigma S dB.$$

Show that $Y = S^\alpha$ also follows a GBM and determine the drift and volatility as a function of μ, σ, and α.

Problem 6. Let S be the price of TESLA stock that follows a geometric Brownian motion such that

$$dS = \mu S dt + \sigma S dB.$$

Your sales team would like to launch a new product called TESLA Quadro that tracks the price of TESLA to the power 4. In other words, the value of this instrument is given by $Y = S^4$. What is the process followed by Y?

Problem 7. GoingUp Corp. has been gaining a lot of attention in the media for its upside potential. Financial experts agree that the stock price follows a geometric Brownian motion with drift (μ) equal to 20% per year and volatility of price returns (σ) of 73% per year. The current stock price is $220. Compute the probability that the stock price is greater than $233 in 10 months from now.

Problem 8. You would like to invest in ZigZag Inc. but you are concerned that the stock price might go down. You have been studying the dynamics of the stock price and concluded that the stock follows a geometric Brownian motion with drift (μ) equal to 13% per year and volatility of price returns (σ) of 58% per year. The current stock price is $118. Compute the probability that the stock price is less than $98 in 12 months from now.

Problem 9. You are analyzing BMX stock. You believe that it is accurate to model the price evolution of the stock as a geometric Brownian motion. Using historical data, you estimate that the drift (μ) is 12.0% per year and the volatility of stock returns (σ) is 39% per year. The stock price just closed at $331. Compute the expected stock price in 9 months from now.