Lorenzo Naranjo October 15, 2026

The Lognormal Distribution

The Black-Scholes model assumes that the price of a stock at a future date has a *lognormal* distribution: its logarithm is normally distributed. This chapter collects the facts about the normal and lognormal distributions that the next chapters use. We show how to compute probabilities, the mean and standard deviation of a lognormal random variable, and the *partial expectations* that turn into the two terms of the Black-Scholes formula.

Why the Lognormal Distribution

Suppose a stock trades at S_0 today and at S_T at a future date T. Its continuously compounded return over the period is $R = \ln(S_T/S_0)$, so that $S_T = S_0 e^R$. Split the period into many short subperiods, and let $R_1, R_2, \ldots, R_n$ be the continuously compounded returns over each of them. Continuously compounded returns add up over time,

$$R = \ln\left(\frac{S_T}{S_0}\right) = \ln\left(\frac{S_{t_1}}{S_0} \times \frac{S_{t_2}}{S_{t_1}} \times \cdots \times \frac{S_T}{S_{t_{n-1}}}\right) = R_1 + R_2 + \cdots + R_n.$$

If the returns over different subperiods are independent and none of them dominates the sum, the central limit theorem says that R is approximately normal, whatever the distribution of each R_i. The price $S_T = S_0 e^R$ is then the exponential of a normal random variable, which is what we call a lognormal random variable.

The lognormal distribution has two features that suit stock prices. A lognormal random variable is always positive, consistent with the limited liability of a stock, and its distribution is skewed to the right: a stock can at most lose all of its value, but it can gain several times its value. The binomial trees of the previous chapters already have this structure, since every period multiplies the stock price by u or d, and the logarithm of the price is a sum of independent steps.

The Normal Distribution

A random variable X is *normally distributed* with mean μ and standard deviation σ, written $X \sim \mathcal{N}(\mu, \sigma^2)$, if its *probability density function* (PDF) is

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}.$$

The density is the familiar bell curve, symmetric around μ, with inflection points at $\mu - \sigma$ and $\mu + \sigma$.

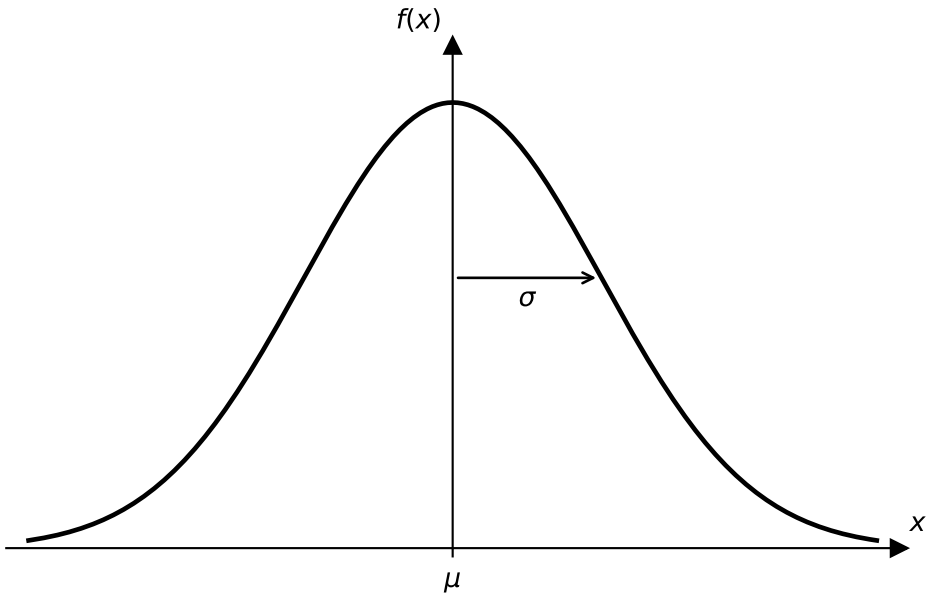


Figure 1: The figure shows the density function of a normally distributed random variable with mean μ and standard deviation σ.

Moments of the Normal Distribution

If $X \sim \mathcal{N}(\mu, \sigma^2)$, then $\mathrm{E}(X) = \mu$ and $\mathrm{V}(X) = \sigma^2$.

Proof

Let $Z = (X - \mu)/\sigma$, which has the standard normal density $\varphi(z) = e^{-z^2/2}/\sqrt{2\pi}$. The function $z\varphi(z)$ is odd, so its integral over the real line is zero and $\mathrm{E}(Z) = 0$. Since $\varphi'(z) = -z\varphi(z)$, integrating by parts gives

$$\mathrm{E}(Z^2) = \int_{-\infty}^{\infty} z \cdot z\varphi(z)\,dz = \Big[-z\varphi(z)\Big]_{-\infty}^{\infty} + \int_{-\infty}^{\infty} \varphi(z)\,dz = 0 + 1 = 1,$$

so $\mathrm{V}(Z) = 1$. Therefore, $\mathrm{E}(X) = \mu + \sigma\,\mathrm{E}(Z) = \mu$ and $\mathrm{V}(X) = \sigma^2\,\mathrm{V}(Z) = \sigma^2$. □

The Standard Normal Distribution

When $\mu = 0$ and $\sigma = 1$, we say that $Z \sim \mathcal{N}(0, 1)$ has the *standard normal distribution*. Its *cumulative distribution function* (CDF) is denoted by the capital Greek letter Φ (phi),

$$\Phi(z) = \mathrm{P}(Z \leq z) = \int_{-\infty}^{z} \frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}}\,du,$$

the area under the density to the left of z.

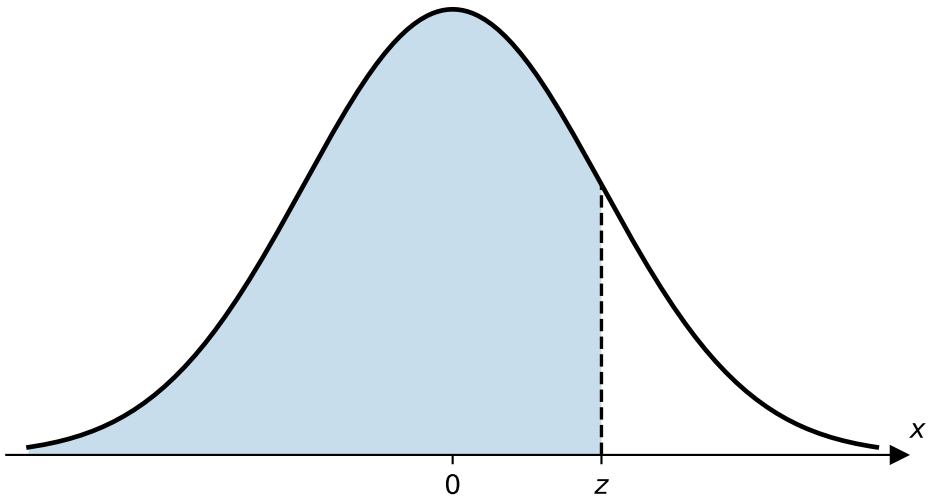


Figure 2: The blue shaded area represents $\Phi(z)$.

The integral has no closed form, so $\Phi(z)$ is obtained from a table or a computer. In Python, SciPy computes $\Phi(-0.4)$ as follows:

```
from scipy.stats import norm

print(norm.cdf(-0.4))
```

```
0.3445782583896758
```

In Excel, `=norm.s.dist(-0.4,TRUE)` gives the same answer.

Every normal random variable is a standard normal in disguise. If $X \sim \mathcal{N}(\mu, \sigma^2)$, then its *z-score* $Z = (X - \mu)/\sigma$ has the standard normal distribution, so that $X = \mu + \sigma Z$. We will use this representation throughout.

Computing Probabilities

To compute the probability that X is below a value x, we translate the statement into one about Z:

$$\mathrm{P}(X \leq x) = \mathrm{P}\left(\frac{X-\mu}{\sigma} \leq \frac{x-\mu}{\sigma}\right) = \Phi\left(\frac{x-\mu}{\sigma}\right).$$

Example 1. Suppose that $X \sim \mathcal{N}(\mu, \sigma^2)$ with $\mu = 10$ and $\sigma = 25$. The probability that $X \leq 0$ is

$$\mathrm{P}(X \leq 0) = \Phi\left(\frac{0-10}{25}\right) = \Phi(-0.40) = 0.3446.$$

□

The probability of the whole distribution is one, so the *right-tail* probability is

$$\mathrm{P}(X > x) = 1 - \mathrm{P}(X \leq x).$$

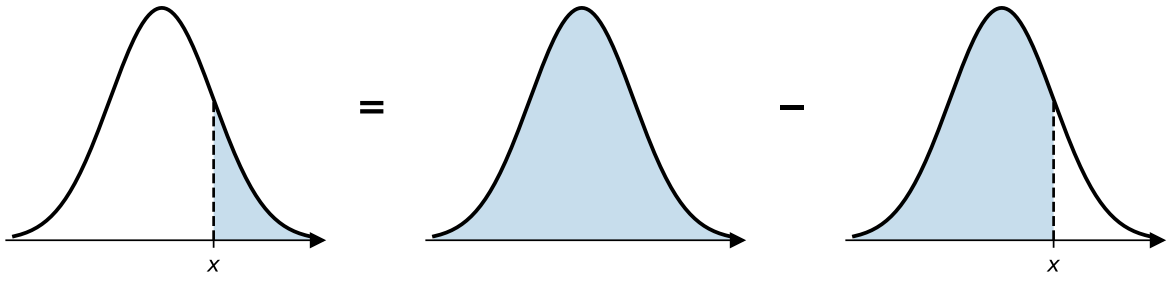


Figure 3: The right-tail probability is the probability of the whole distribution, which is one, minus the left-tail probability.

Since the standard normal density is symmetric around zero, the area to the right of z equals the area to the left of $-z$, so that $1 - \Phi(z) = \Phi(-z)$. We will use this identity to write right-tail probabilities as values of Φ.

Example 2. Suppose that $X \sim \mathcal{N}(\mu, \sigma^2)$ with $\mu = 10$ and $\sigma = 25$. The probability that $X > 12$ is

$$\mathrm{P}(X > 12) = 1 - \Phi\left(\frac{12 - 10}{25}\right) = 1 - \Phi(0.08) = 1 - 0.5319 = 0.4681.$$

□

The probability that X falls in an interval $(x_1, x_2]$ is the area to the left of x_2 minus the area to the left of x_1,

$$\mathrm{P}(x_1 < X \leq x_2) = \mathrm{P}(X \leq x_2) - \mathrm{P}(X \leq x_1).$$

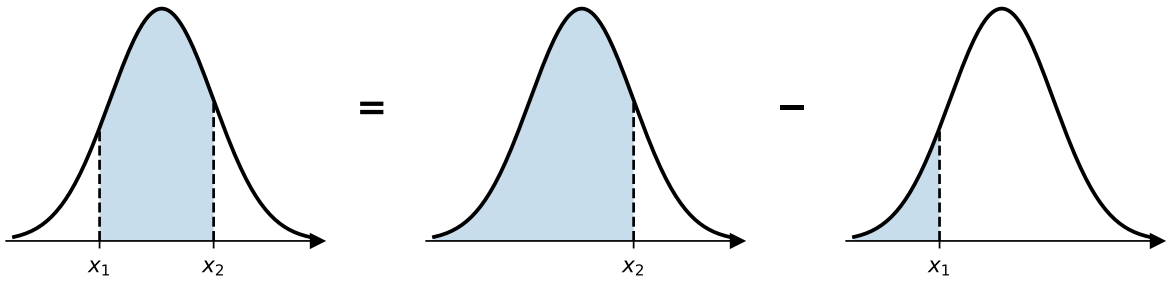


Figure 4: Subtracting the area to the left of x_1 from the area to the left of x_2 gives the probability that $x_1 < X \leq x_2$.

Example 3. Suppose that $X \sim \mathcal{N}(\mu, \sigma^2)$ with $\mu = 10$ and $\sigma = 25$. The probability that $2 < X \leq 14$ is

$$\begin{aligned}\mathrm{P}(2 < X \leq 14) &= \Phi\left(\frac{14 - 10}{25}\right) - \Phi\left(\frac{2 - 10}{25}\right) \\ &= \Phi(0.16) - \Phi(-0.32) \\ &= 0.5636 - 0.3745 = 0.1891.\end{aligned}$$

□

Percentiles and Probability Intervals

We often need the reverse calculation: the value that leaves a given probability in the tail. For a standard normal Z, the *right-tail percentile* z_α is the value above which the probability is α,

$$\mathrm{P}(Z > z_\alpha) = \alpha, \qquad \text{or equivalently} \qquad z_\alpha = \Phi^{-1}(1-\alpha),$$

where Φ^{-1} is the inverse of Φ.

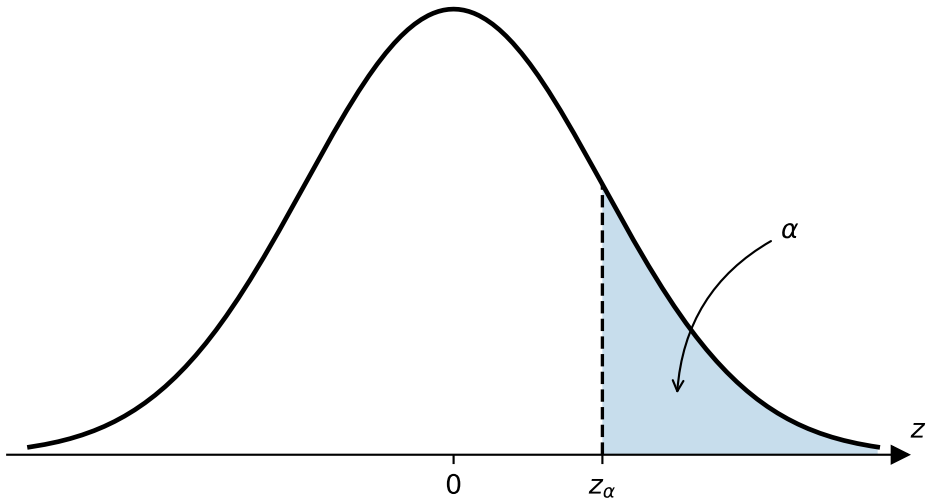


Figure 5: The right-tail percentile is the value z_α that gives an area to the right equal to α.

Like Φ, its inverse has no closed form. In Python, `norm.ppf` computes it, so that $z_{0.025} = \Phi^{-1}(0.975)$ is

```
print(norm.ppf(0.975))
```

```
1.959963984540054
```

In Excel, `=norm.s.inv(0.975)` gives the same result. The table below shows the values we use most often.

α	z_α
0.050	1.645
0.025	1.960
0.010	2.326
0.005	2.576

By symmetry, Z falls between $-z_{\alpha/2}$ and $z_{\alpha/2}$ with probability $1-\alpha$, leaving $\alpha/2$ in each tail. For instance, Z lies in $[-1.96, 1.96]$ with probability 95%.

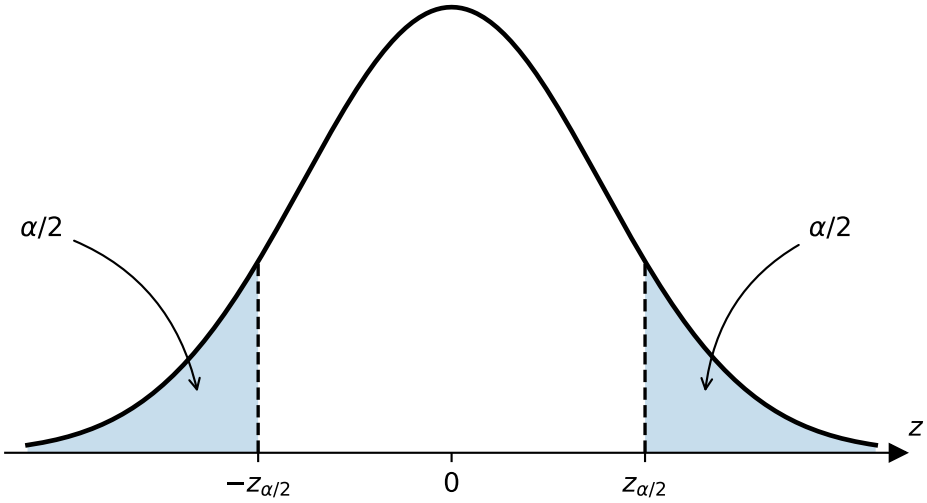


Figure 6: The areas on each side are both equal to $\alpha/2$.

Since $X = \mu + \sigma Z$, the same statement for $X \sim \mathcal{N}(\mu, \sigma^2)$ is that X lies in

$$[\mu - z_{\alpha/2}\sigma, \mu + z_{\alpha/2}\sigma]$$

with probability $1-\alpha$. We call this interval a $(1-\alpha)$ *probability interval* for X. It is often called a confidence interval, although in statistics that term refers to an interval estimated from data.

Example 4. Suppose that $X \sim \mathcal{N}(\mu, \sigma^2)$ with $\mu = 10$ and $\sigma = 25$. Since $z_{0.025} = 1.96$, the 95% probability interval for X is

$$[10 - 1.96(25), 10 + 1.96(25)] = [-39, 59].$$

□

Exponentials of a Normal Random Variable

One formula does most of the work in the rest of this chapter. It gives the expected value of e^{aX} for any constant a, the *moment generating function* of the normal distribution.

Expected Exponential of a Normal Random Variable

If $X \sim \mathcal{N}(\mu, \sigma^2)$, then for any constant a,

$$\mathrm{E}\left(e^{aX}\right) = e^{a\mu + \frac{1}{2}a^2\sigma^2}. \tag{1}$$

Proof

Let $\varphi(z) = e^{-z^2/2}/\sqrt{2\pi}$ be the standard normal density. Completing the square, $az - \frac{1}{2}z^2 = -\frac{1}{2}(z-a)^2 + \frac{1}{2}a^2$, so that

$$e^{az}\varphi(z) = e^{\frac{1}{2}a^2}\varphi(z-a). \tag{2}$$

The right-hand side is $e^{\frac{1}{2}a^2}$ times the density of a normal random variable with mean a and variance one, which integrates to one. Integrating both sides gives $\mathrm{E}(e^{aZ}) = e^{\frac{1}{2}a^2}$ for a standard normal Z. Writing $X = \mu + \sigma Z$,

$$\mathrm{E}\left(e^{aX}\right) = e^{a\mu}\,\mathrm{E}\left(e^{a\sigma Z}\right) = e^{a\mu + \frac{1}{2}a^2\sigma^2}.$$

□

The term $\frac{1}{2}a^2\sigma^2$ in (1) is worth noticing. Since e^x is convex, the expected value of e^{aX} is larger than $e^{a\,\mathrm{E}(X)} = e^{a\mu}$, and the gap grows with the variance of X. The same term reappears in the next chapter as the difference between the expected return of a stock and the growth rate of its price.

The Lognormal Distribution

If $X \sim \mathcal{N}(\mu, \sigma^2)$, then $Y = e^X$ is *lognormally distributed* with parameters μ and σ^2. The parameters are the mean and variance of $\ln(Y)$, not of Y. Figure 7 compares the two densities.

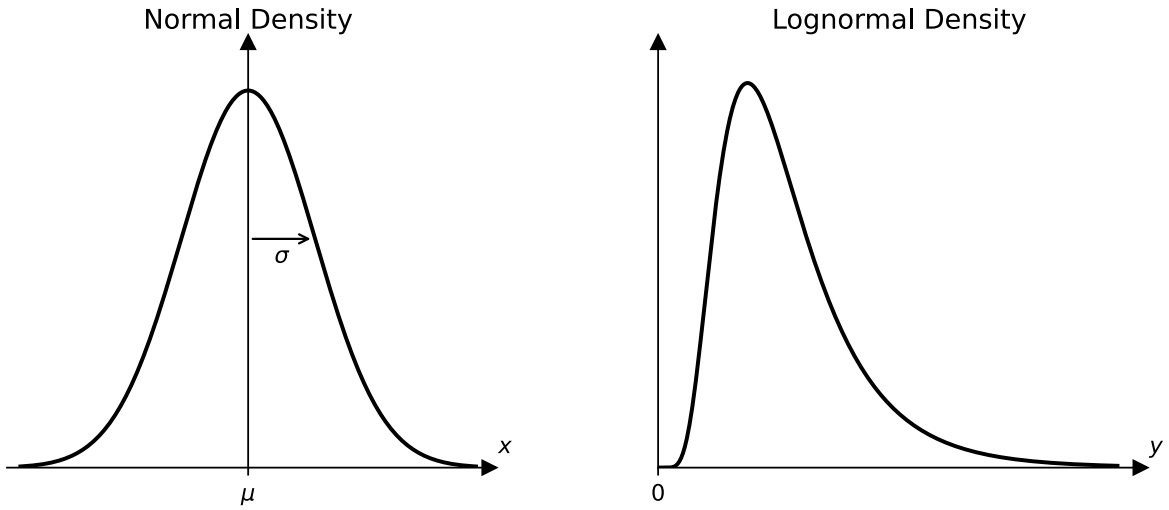


Figure 7: The figure shows the difference between a normal and a lognormal PDF with the same parameters.

Lognormal Density

If Y is lognormally distributed with parameters μ and σ^2, its PDF is

$$f(y) = \frac{1}{y\sqrt{2\pi\sigma^2}} e^{-\frac{(\ln(y)-\mu)^2}{2\sigma^2}}, \qquad y > 0.$$

Proof

For $y > 0$, the CDF of Y is $\mathrm{P}(Y \leq y) = \mathrm{P}(X \leq \ln(y))$. Differentiating with respect to y and using the chain rule, the density of Y is the normal density of X evaluated at $\ln(y)$, times the derivative $d\ln(y)/dy = 1/y$, which is the expression above. □

Unlike the normal density, the lognormal density is not symmetric. A normal random variable can take any real value, whereas a lognormal random variable is always positive,

and its long right tail means that large values are more likely than equally distant small ones.

Computing Probabilities

The logarithm is an increasing function, so any statement about Y can be translated into a statement about the normal random variable $\ln(Y)$:

$$\mathrm{P}(Y \le y) = \mathrm{P}(\ln(Y) \le \ln(y)) = \Phi\left(\frac{\ln(y) - \mu}{\sigma}\right).$$

Example 5. Let $Y = e^{4+1.5Z}$ where $Z \sim \mathcal{N}(0,1)$. The probability that $Y \le 100$ is

$$\mathrm{P}(Y \le 100) = \mathrm{P}\left(Z \le \frac{\ln(100) - 4}{1.5}\right) = \Phi(0.4034) = 0.6567.$$

□

Probability intervals translate in the same way. If $\ln(Y)$ lies in $[\mu - z_{\alpha/2}\sigma, \mu + z_{\alpha/2}\sigma]$ with probability $1 - \alpha$, then Y lies in $[e^{\mu - z_{\alpha/2}\sigma}, e^{\mu + z_{\alpha/2}\sigma}]$ with the same probability.

Example 6. Let $Y = e^{4+1.5Z}$ where $Z \sim \mathcal{N}(0,1)$. The 95% probability interval for Y is

$$[e^{4-1.96(1.5)}, e^{4+1.96(1.5)}] = [2.89, 1032.77].$$

The interval is far from symmetric around $e^4 = 54.60$, the value of Y when $\ln(Y)$ is at its mean. □

Moments

Since $Y^a = e^{aX}$, the moments of Y follow directly from (1).

Moments of a Lognormal Distribution

Let $Y = e^X$ where $X \sim \mathcal{N}(\mu, \sigma^2)$. For any constant a,

$$\mathrm{E}(Y^a) = e^{a\mu + \frac{1}{2}a^2\sigma^2}.$$

In particular,

$$\mathrm{E}(Y) = e^{\mu + \frac{1}{2}\sigma^2},$$
$$\mathrm{V}(Y) = e^{2\mu+\sigma^2}\left(e^{\sigma^2} - 1\right),$$
$$\mathrm{SD}(Y) = \mathrm{E}(Y)\sqrt{e^{\sigma^2} - 1}.$$

Proof

The first formula is (1). Taking $a = 1$ gives $\mathrm{E}(Y)$, and taking $a = 2$ gives $\mathrm{E}(Y^2) = e^{2\mu+2\sigma^2}$, so that

$$\mathrm{V}(Y) = \mathrm{E}(Y^2) - \mathrm{E}(Y)^2 = e^{2\mu+2\sigma^2} - e^{2\mu+\sigma^2} = e^{2\mu+\sigma^2}\left(e^{\sigma^2} - 1\right),$$

and the standard deviation is its square root. □

Example 7. Let $Y = e^{4+1.5Z}$ where $Z \sim \mathcal{N}(0, 1)$. The expectation and standard deviation of Y are

$$\mathrm{E}(Y) = e^{4+0.5(1.5^2)} = 168.17,$$
$$\mathrm{SD}(Y) = 168.17\sqrt{e^{1.5^2} - 1} = 489.95.$$

□

The mean of Y is not e^{μ}. Half of the distribution of $\ln(Y)$ lies below μ, so the *median* of Y is e^{μ}, while its mean $e^{\mu+\frac{1}{2}\sigma^2}$ is larger. A few very large values pull the mean above the typical value. In Example 7 the median is $e^4 = 54.60$ and the mean is 168.17, and Y falls below its mean with probability $\Phi(\frac{1}{2}\sigma) = \Phi(0.75) = 0.7734$.

Partial Expectations

A call option with strike price K pays $Y - K$ if $Y > K$ and nothing otherwise. To write such payoffs compactly, we use the *indicator function*

$$\mathbb{1}_{\{Y>K\}} = \begin{cases} 1 \text{ if } Y > K, \\ 0 \text{ if } Y \le K, \end{cases}$$

so that the payoff of the call is $\max(Y - K, 0) = Y\mathbb{1}_{\{Y>K\}} - K\mathbb{1}_{\{Y>K\}}$. Its expected value requires two *partial expectations*, which average Y and K only over the outcomes where the option is exercised.

Partial Expectations

Let $Y = e^X$ where $X \sim \mathcal{N}(\mu, \sigma^2)$, let $K > 0$, and define

$$d = \frac{\mu - \ln(K)}{\sigma}.$$

Then

$$\mathrm{P}(Y > K) = \Phi(d), \qquad \mathrm{E}\left(Y\mathbb{1}_{\{Y>K\}}\right) = \mathrm{E}(Y)\Phi(d + \sigma),$$

where $\mathrm{E}(Y) = e^{\mu+\frac{1}{2}\sigma^2}$. The second partial expectation is $\mathrm{E}(K\mathbb{1}_{\{Y>K\}}) = K\,\mathrm{P}(Y > K) = K\,\Phi(d)$.

Proof

Write $Y = e^{\mu+\sigma Z}$ with Z standard normal. Then $Y > K$ exactly when $Z > -d$, so that $\mathrm{P}(Y > K) = 1 - \Phi(-d) = \Phi(d)$. For the partial expectation, (2) with $a = \sigma$ gives

$$\mathrm{E}\left(Y\mathbb{1}_{\{Y>K\}}\right) = e^{\mu}\int_{-d}^{\infty} e^{\sigma z}\varphi(z)\,dz = e^{\mu+\frac{1}{2}\sigma^2}\int_{-d}^{\infty}\varphi(z-\sigma)\,dz = \mathrm{E}(Y)\mathrm{P}(Z + \sigma > -d),$$

and $\mathrm{P}(Z + \sigma > -d) = \mathrm{P}(Z > -(d + \sigma)) = \Phi(d + \sigma)$. □

The first result is the probability that the option is exercised. The second is the full mean of Y times a probability whose argument is shifted from d to $d + \sigma$. Weighting each

outcome by $Y = e^{\mu+\sigma Z}$ gives more weight to high values of Z, and the effect is the same as shifting Z to the right by σ.

Example 8. Let $Y = e^{4+1.5Z}$ where $Z \sim \mathcal{N}(0,1)$, and let $K = 100$. Then $d = (4 - \ln(100))/1.5 = -0.4034$ and $d + \sigma = 1.0966$, so that

$$\mathrm{P}(Y > 100) = \Phi(-0.4034) = 0.3433,$$
$$\mathrm{E}\left(Y \mathbb{1}_{\{Y>100\}}\right) = 168.17\,\Phi(1.0966) = 145.23.$$

The expected payoff of a call on Y with strike price 100 is $\mathrm{E}(\max(Y - 100, 0)) = 145.23 - 100 \times 0.3433 = 110.90$. □

The Black-Scholes chapter computes the price of a call in exactly this way: it takes the expected payoff of the option under a particular lognormal distribution and discounts it at the risk-free rate.

Practice Problems

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

Problem 1. Suppose that X is a normally distributed random variable with mean $\mu = 12$ and standard deviation $\sigma = 20$.

a. What is the probability that $X \leq 0$?
b. What is the probability that $X \leq -4$?
c. What is the probability that $X > 8$?
d. What is the probability that $4 < X \leq 10$?

Problem 2. Suppose that X is a normally distributed random variable with mean $\mu = 10$ and standard deviation $\sigma = 20$. Compute the 90%, 95%, and 99% probability intervals for X.

Problem 3. Suppose that $X = \ln(Y)$ is a normally distributed random variable with mean $\mu = 3.9$ and standard deviation $\sigma = 15$.

a. What is the probability that $Y \leq 6$?

b. What is the probability that $Y > 4$?
c. What is the probability that $3 < Y \leq 12$?
d. What is the probability that $Y \leq 0$?

Problem 4. Suppose that X is a normally distributed variable with mean $\mu = 3.70$ and standard deviation $\sigma = 0.80$. If $Y = e^X$, what is the probability that Y is greater than 45?

Problem 5. Suppose that $X = \ln(Y)$ is a normally distributed random variable with mean $\mu = 2.7$ and standard deviation $\sigma = 1$. Compute the 90%, 95%, and 99% probability intervals for X and report the corresponding values for Y.

Problem 6. Let $Y = e^{\mu+\sigma Z}$ where $\mu = 1, \sigma = 2$ and $Z \sim \mathcal{N}(0,1)$. Compute:

a. $\mathrm{E}(Y)$
b. $\mathrm{SD}(Y) = \sqrt{\mathrm{E}(Y^2) - \mathrm{E}(Y)^2}$
c. $\mathrm{E}(Y^{0.3})$
d. $\mathrm{E}(Y^{-1})$

Problem 7. Let $Y = e^X$ where $X \sim \mathcal{N}(\mu, \sigma^2)$ with $\mu = 4.75$ and $\sigma = 0.35$. Compute $\mathrm{E}(Y\mathbb{1}_{\{Y>95\}})$, $\mathrm{P}(Y > 95)$, and the expected payoff $\mathrm{E}(\max(Y - 95, 0))$ of a call on Y with strike price 95.