

The Black-Scholes Model

A Mathematical Treatment of Options Pricing

From Geometric Brownian Motion to the Greeks

Covering: Stochastic Calculus Foundations, Itô's Lemma, the Black-Scholes PDE, Closed-Form Solutions, Risk-Neutral Pricing, the Greeks, and Model Extensions

A Reference for Quantitative Finance

Prerequisites: Single-variable calculus (derivatives, integrals, chain rule)

© 2026 **AlphaPilot**. All rights reserved.

No part of this document may be reproduced, distributed, or transmitted in any form without prior written permission from AlphaPilot.

Contents

1	Introduction to Options	3
1.1	Payoff Functions	3
1.2	European vs. American Options	4
1.3	The Central Problem	4
2	Mathematical Preliminaries	4
2.1	Random Variables and Expectation	4
2.2	The Normal Distribution	5
2.3	The Log-Normal Distribution	5
2.4	Brownian Motion (Wiener Process)	6
2.5	Stochastic Differential Equations	7
3	Geometric Brownian Motion	8
3.1	Solution of the GBM SDE	8
3.2	Key Assumptions of the GBM Model	9
4	Itô's Lemma	9
4.1	Why Ordinary Calculus Fails	10
4.2	Statement of Itô's Lemma	10
4.3	Derivation	10
4.4	Application: Solving the GBM Equation	11
5	Derivation of the Black-Scholes PDE	11
5.1	Setup	12
5.2	Constructing the Hedging Portfolio	12
5.3	The No-Arbitrage Condition	13
5.4	Key Observations	13
6	Solving the Black-Scholes Equation	13
6.1	Boundary Conditions	13
6.2	Transformation to the Heat Equation	14
6.3	Solution via the Heat Kernel	14
7	The Black-Scholes Pricing Formulas	14
7.1	European Call Option	15
7.2	European Put Option	15
7.3	Interpretation of the Terms	15
7.4	The Five Inputs	16
8	Put-Call Parity	17
8.1	Derivation	17
8.2	Verification with Black-Scholes	18

9	The Greeks: Sensitivity Analysis	18
9.1	Delta (Δ)	18
9.2	Gamma (Γ)	19
9.3	Theta (Θ)	20
9.4	Vega (\mathcal{V})	21
9.5	Rho (ρ)	21
9.6	Summary Table	22
10	Risk-Neutral Pricing Framework	22
10.1	The Risk-Neutral Measure	22
10.2	Deriving Black-Scholes via Risk-Neutral Pricing	23
10.3	Girsanov's Theorem (Sketch)	23
11	Limitations and Extensions	23
11.1	Non-Constant Volatility	23
11.2	Jump Risk	24
11.3	Transaction Costs and Discrete Hedging	24
11.4	Dividends	25
11.5	Non-Normal Returns	25
11.6	Machine Learning Approaches	25
12	References	25

Introduction to Options

A financial option is a type of **derivative**—a contract whose value is derived from an underlying asset (typically a stock). Options give the holder the *right, but not the obligation*, to buy or sell the underlying asset at a predetermined price.

Before diving in, let us establish the key terminology:

- **Underlying asset:** The stock (or other asset) the option is based on, with price denoted S .
- **Strike price (K):** The predetermined price at which the holder can buy or sell the stock.
- **Expiration date (T):** The date by which the option must be exercised.
- **Premium:** The price paid to purchase the option.

There are two fundamental types of options:

- A **call option** gives the holder the right to *buy* the stock at price K .
- A **put option** gives the holder the right to *sell* the stock at price K .

Payoff Functions

At expiration time T , the value of an option depends entirely on the stock price S_T relative to the strike K . If you hold a call option and the stock is trading at \$150 with a strike of \$100, you can buy at \$100 and immediately sell at \$150, netting \$50. If the stock is below \$100, the option is worthless—you would not exercise it.

Mathematically, the **call option payoff** is:

$$C(S_T, T) = \max(S_T - K, 0) \quad (1)$$

The **put option payoff** is:

$$P(S_T, T) = \max(K - S_T, 0) \quad (2)$$

The max function ensures you never have a negative payoff—you simply do not exercise the option if it would result in a loss.

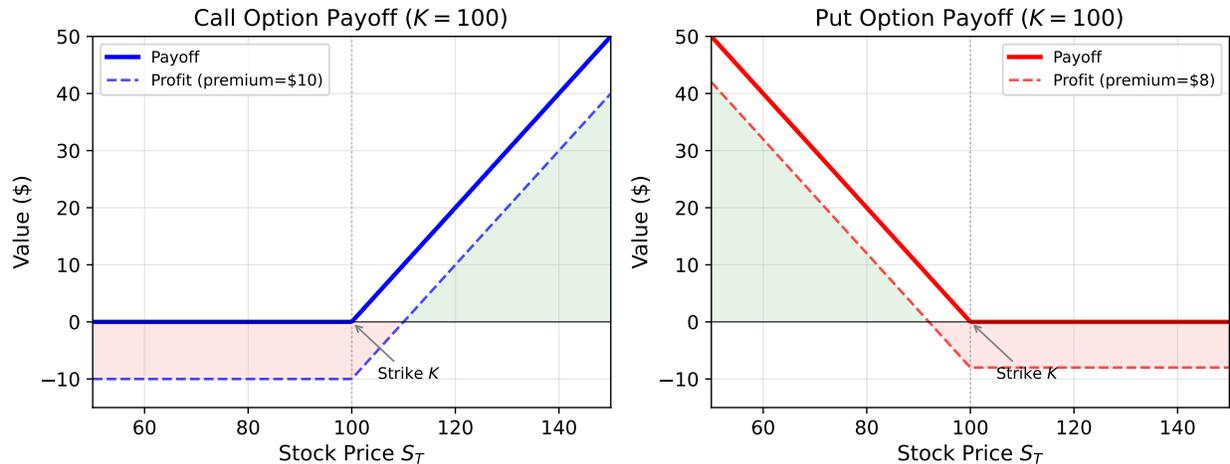


Figure 1: Option payoff and profit diagrams at expiration ($K = 100$). The solid line shows the payoff; the dashed line shows profit after subtracting the premium paid. Green shading indicates profit; red indicates loss. Notice the asymmetry: the call has unlimited upside, while the put's maximum payoff is K (if the stock drops to zero).

European vs. American Options

A **European option** can only be exercised at the expiration date T . An **American option** can be exercised at any time t where $0 \leq t \leq T$. The Black-Scholes model prices European options. American options require more complex numerical methods (binomial trees, finite differences).

The Central Problem

The fundamental question is: *What is the fair price of an option today, at time $t < T$?* We know the payoff at expiration—but the future stock price S_T is uncertain. The Black-Scholes model gives a precise, closed-form answer under specific assumptions about how stock prices move.

Further Reading

- [Investopedia: Options Basics](#) — Gentle introduction to what options are and how they work.
- [Khan Academy: Derivative Securities](#) — Free video lectures covering options from scratch.

Mathematical Preliminaries

Before we can derive the Black-Scholes equation, we need tools from probability theory and stochastic calculus. If you are comfortable with basic calculus (derivatives, integrals, the chain rule), the material below will extend those concepts to handle *randomness*.

Random Variables and Expectation

A **random variable** X is a quantity whose value is determined by a random outcome. The **expected value** (or mean) of X , denoted $\mathbb{E}[X]$, is the probability-weighted average of all possible

values:

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} x f(x) dx$$

where $f(x)$ is the **probability density function** (PDF) of X . The **variance** measures the spread:

$$\text{Var}(X) = \mathbb{E}[(X - \mathbb{E}[X])^2] = \mathbb{E}[X^2] - (\mathbb{E}[X])^2$$

The Normal Distribution

The **normal (Gaussian) distribution** with mean μ and variance σ^2 has PDF:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right)$$

We write $X \sim \mathcal{N}(\mu, \sigma^2)$. The **standard normal** has $\mu = 0$ and $\sigma = 1$.

The **cumulative distribution function** (CDF) of the standard normal, denoted $N(x)$, gives the probability that a standard normal variable is less than or equal to x :

$$N(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-u^2/2} du \quad (3)$$

Key properties: $N(0) = 0.5$, $N(\infty) = 1$, $N(-\infty) = 0$, and the symmetry relation $N(-x) = 1 - N(x)$. This function appears throughout the Black-Scholes formula.

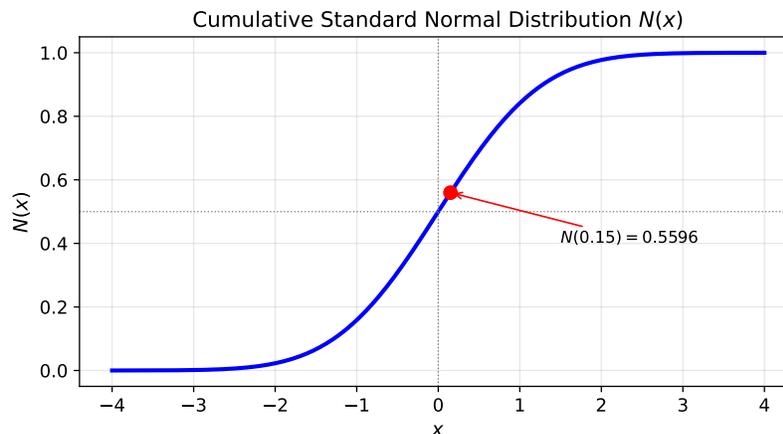


Figure 2: The cumulative standard normal distribution $N(x)$. This S-shaped curve maps any real number to a probability between 0 and 1. In the Black-Scholes formula, $N(d_2)$ gives the risk-neutral probability of the option expiring in the money.

The Log-Normal Distribution

If $\ln(X)$ is normally distributed, then X follows a **log-normal distribution**. This is important because the Black-Scholes model implies stock prices are log-normally distributed—ensuring they remain positive (a stock price can never be negative).

If $\ln(S_T) \sim \mathcal{N}(m, v^2)$, then:

$$\mathbb{E}[S_T] = e^{m+v^2/2} \quad (4)$$

$$\text{Var}(S_T) = e^{2m+v^2} (e^{v^2} - 1) \quad (5)$$

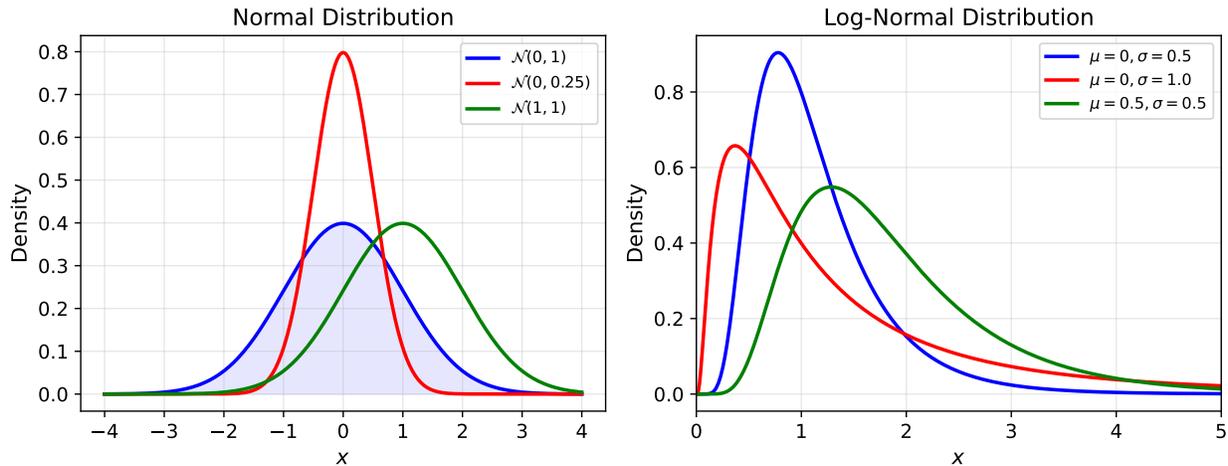


Figure 3: Left: Normal distributions with different means and variances. Right: Log-normal distributions. Notice that the log-normal is always positive and right-skewed—matching the behavior of stock prices, which cannot go negative and occasionally make large upward moves.

Why log-normal?

If stock returns are normally distributed (which is a reasonable approximation for short time intervals), then stock *prices* are log-normally distributed. This follows because the price is the exponential of the cumulative returns: $S_T = S_0 e^{\text{cumulative return}}$.

Brownian Motion (Wiener Process)

A **standard Brownian motion** (or Wiener process) $W(t)$ is a continuous-time random process that models “pure randomness.” Think of it as the limiting case of a random walk as the step size goes to zero.

Definition 2.1 (Standard Brownian Motion). *A process $W(t)$ is a standard Brownian motion if:*

1. $W(0) = 0$.
2. **Independent increments:** For $0 \leq s < t$, the change $W(t) - W(s)$ is independent of everything that happened before time s .
3. **Normal increments:** $W(t) - W(s) \sim \mathcal{N}(0, t - s)$.
4. **Continuous paths:** $W(t)$ is a continuous function of t (no jumps).

Over a small time interval dt , the increment dW satisfies:

$$dW \sim \mathcal{N}(0, dt), \quad \mathbb{E}[dW] = 0, \quad \mathbb{E}[dW^2] = dt$$

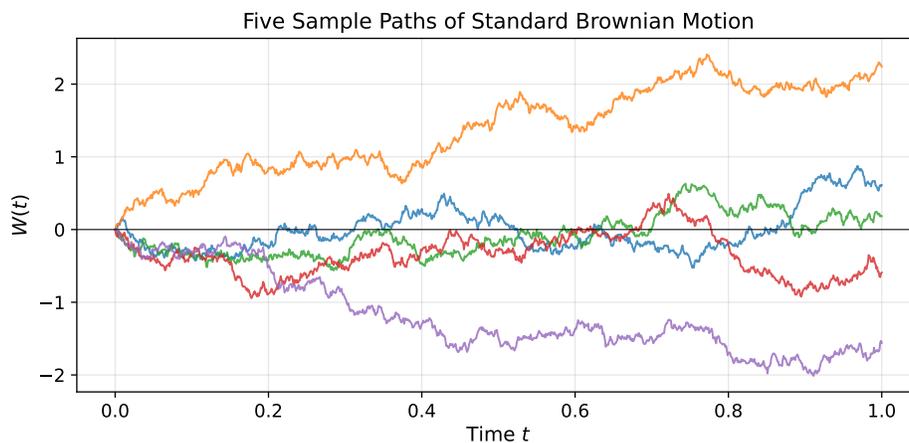


Figure 4: Five sample paths of standard Brownian motion $W(t)$ over one year. Each path starts at 0 and wanders randomly. Notice how the paths are continuous but extremely jagged—this is why ordinary calculus does not apply and we need stochastic calculus.

Brownian motion has a remarkable property: although its paths are continuous, they are *nowhere differentiable*. The paths are infinitely jagged at every scale. This is why we need a new kind of calculus (stochastic calculus) to work with it.

Stochastic Differential Equations

A **stochastic differential equation** (SDE) describes how a random process evolves:

$$dX(t) = a(X, t) dt + b(X, t) dW(t) \tag{6}$$

Here:

- $a(X, t)$ is the **drift** — the deterministic trend (“where the process tends to go”).
- $b(X, t)$ is the **diffusion** — the magnitude of the random fluctuations.
- $dW(t)$ is the Brownian motion increment (the source of randomness).

The dt term pulls the process in a particular direction; the dW term adds random noise. The relative sizes of a and b determine how “noisy” the process is.

Further Reading

- [Introduction to Probability \(probabilitycourse.com\)](http://probabilitycourse.com) — Free textbook covering random variables, distributions, and expectation.
- [Wikipedia: Wiener Process](#) — Good overview with visualizations of Brownian motion paths.
- [MIT OCW: Stochastic Calculus](#) — Lecture notes for a deeper treatment of SDEs and Itô calculus.

Geometric Brownian Motion

The Black-Scholes model assumes the stock price $S(t)$ follows a **Geometric Brownian Motion** (GBM):

$$dS = \mu S dt + \sigma S dW \quad (7)$$

where:

- S = stock price at time t
- μ = **drift rate** (expected return per unit time, e.g., $\mu = 0.10$ means 10% expected annual return)
- σ = **volatility** (standard deviation of returns per unit time, e.g., $\sigma = 0.20$ means 20% annual volatility)
- dW = increment of a standard Brownian motion

Interpreting the SDE

Dividing both sides of Equation (7) by S gives the **return process**:

$$\frac{dS}{S} = \mu dt + \sigma dW$$

This says: the instantaneous percentage return on the stock equals a predictable part (μdt) plus a random part (σdW). The randomness is *proportional to the price itself*—a \$100 stock fluctuates twice as much in dollar terms as a \$50 stock, but the same in percentage terms.

Solution of the GBM SDE

We will derive the solution in the next chapter using Itô's Lemma. The result is:

$$S(T) = S(0) \exp \left[\left(\mu - \frac{\sigma^2}{2} \right) T + \sigma W(T) \right] \quad (8)$$

This tells us the logarithm of the stock price is normally distributed:

$$\ln S(T) \sim \mathcal{N} \left(\ln S(0) + \left(\mu - \frac{\sigma^2}{2} \right) T, \sigma^2 T \right)$$

The term $\sigma^2/2$ is called the **Itô correction** (or convexity adjustment). It appears because for a random variable X , $\mathbb{E}[e^X] \neq e^{\mathbb{E}[X]}$ (Jensen's inequality). Without this correction, the expected stock price would not match what the drift μ implies. With it, we get $\mathbb{E}[S(T)] = S(0) e^{\mu T}$, as desired.

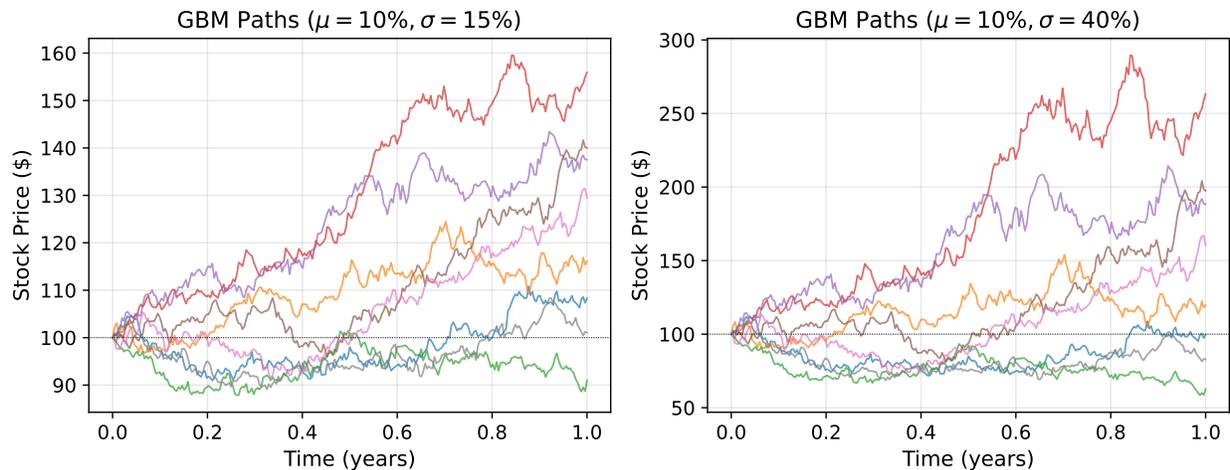


Figure 5: Simulated stock price paths under GBM with the same drift ($\mu = 10\%$) but different volatilities. Left: $\sigma = 15\%$ produces relatively smooth paths. Right: $\sigma = 40\%$ produces wild swings. Both start at \$100. Higher volatility means wider dispersion of outcomes, which is why options become more valuable when volatility increases.

Key Assumptions of the GBM Model

1. **Continuous trading:** The stock can be traded at any instant.
2. **Constant parameters:** μ and σ do not change over the life of the option.
3. **No jumps:** The price path is continuous (no overnight gaps or crash discontinuities).
4. **Log-normal returns:** Returns over any interval are normally distributed.
5. **No dividends:** The stock pays no dividends (this assumption can be relaxed later).
6. **No transaction costs:** Trading is frictionless.
7. **No arbitrage:** There are no risk-free profit opportunities.

Further Reading

- [Wikipedia: Geometric Brownian Motion](#) — Includes simulated sample paths and statistical properties.
- [Investopedia: Monte Carlo Simulation](#) — How GBM is used in practice to simulate stock prices.

Itô's Lemma

In ordinary calculus, if you have a function $f(x)$ and x changes, the **chain rule** tells you how f changes:

$$df = f'(x) dx$$

But when x is a random process driven by Brownian motion, this rule needs modification. The reason is subtle but fundamental: Brownian motion is so “rough” that second-order terms, which we normally neglect in calculus, become significant.

Why Ordinary Calculus Fails

Recall from calculus that a Taylor expansion to second order gives:

$$df = f'(x) dx + \frac{1}{2} f''(x) (dx)^2 + \dots$$

In ordinary calculus, $(dx)^2$ is infinitesimally smaller than dx , so we drop it. But for Brownian motion:

$$(dW)^2 = dt \neq 0$$

This is the **quadratic variation** of Brownian motion—it accumulates at a finite rate. Intuitively, Brownian paths wiggle so violently that the squared increments sum to something non-negligible. This single fact is what makes stochastic calculus different from ordinary calculus.

Statement of Itô’s Lemma

Theorem 4.1 (Itô’s Lemma). *Let $X(t)$ satisfy the SDE $dX = a dt + b dW$. Let $f(X, t)$ be a twice continuously differentiable function. Then:*

$$df = \left(\frac{\partial f}{\partial t} + a \frac{\partial f}{\partial X} + \frac{1}{2} b^2 \frac{\partial^2 f}{\partial X^2} \right) dt + b \frac{\partial f}{\partial X} dW \quad (9)$$

Compare this with the ordinary chain rule: Itô’s Lemma has an extra term $\frac{1}{2} b^2 \frac{\partial^2 f}{\partial X^2}$. This is the correction arising from the quadratic variation of Brownian motion.

Derivation

Start with the Taylor expansion of $f(X, t)$ to second order:

$$df = \frac{\partial f}{\partial t} dt + \frac{\partial f}{\partial X} dX + \frac{1}{2} \frac{\partial^2 f}{\partial X^2} (dX)^2 + \frac{\partial^2 f}{\partial t \partial X} (dt)(dX) + \dots$$

Substitute $dX = a dt + b dW$:

$$(dX)^2 = (a dt + b dW)^2 = a^2 (dt)^2 + 2ab(dt)(dW) + b^2 (dW)^2$$

Now apply the **multiplication rules for stochastic infinitesimals**:

$$(dt)^2 = 0, \quad (dt)(dW) = 0, \quad (dW)^2 = dt \quad (10)$$

The first two rules are the same as ordinary calculus (higher-order infinitesimals vanish). The third rule, $(dW)^2 = dt$, is the game-changer. Therefore:

$$(dX)^2 = b^2 dt$$

Substituting back into the Taylor expansion and keeping terms up to order dt :

$$df = \frac{\partial f}{\partial t} dt + \frac{\partial f}{\partial X}(a dt + b dW) + \frac{1}{2} \frac{\partial^2 f}{\partial X^2} (b^2 dt)$$

Collecting the dt and dW terms:

$$df = \left(\frac{\partial f}{\partial t} + a \frac{\partial f}{\partial X} + \frac{1}{2} b^2 \frac{\partial^2 f}{\partial X^2} \right) dt + b \frac{\partial f}{\partial X} dW$$

This is Itô's Lemma. \square

Application: Solving the GBM Equation

Let us apply Itô's Lemma to find the distribution of stock prices. Set $f(S) = \ln S$ where $dS = \mu S dt + \sigma S dW$.

We need the partial derivatives:

$$\frac{\partial f}{\partial S} = \frac{1}{S}, \quad \frac{\partial^2 f}{\partial S^2} = -\frac{1}{S^2}, \quad \frac{\partial f}{\partial t} = 0$$

Applying Itô's Lemma with $a = \mu S$ and $b = \sigma S$:

$$\begin{aligned} d(\ln S) &= \left(0 + \mu S \cdot \frac{1}{S} + \frac{1}{2} (\sigma S)^2 \cdot \left(-\frac{1}{S^2} \right) \right) dt + \sigma S \cdot \frac{1}{S} dW \\ &= \left(\mu - \frac{\sigma^2}{2} \right) dt + \sigma dW \end{aligned}$$

This is a simple SDE with constant coefficients! Integrating from 0 to T :

$$\ln S(T) - \ln S(0) = \left(\mu - \frac{\sigma^2}{2} \right) T + \sigma W(T) \quad (11)$$

$$S(T) = S(0) \exp \left[\left(\mu - \frac{\sigma^2}{2} \right) T + \sigma W(T) \right] \quad (12)$$

This confirms Equation (8) and shows where the $\sigma^2/2$ correction comes from: it arises naturally from the $\frac{1}{2}b^2 f''$ term in Itô's Lemma.

Further Reading

- [Wikipedia: Itô's Lemma](#) — Multiple examples and the multidimensional version.
- [MIT OCW: Itô Calculus \(YouTube\)](#) — Video lecture explaining the intuition behind stochastic calculus.

Derivation of the Black-Scholes PDE

This is the intellectual heart of the Black-Scholes theory. The key idea is to construct a portfolio that eliminates all randomness, then use the fact that a risk-free portfolio must earn the risk-free rate.

Setup

Let $V(S, t)$ denote the price of a derivative (call or put) as a function of the stock price S and time t . Since S follows GBM, we apply Itô's Lemma to V :

$$dV = \left(\frac{\partial V}{\partial t} + \mu S \frac{\partial V}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} \right) dt + \sigma S \frac{\partial V}{\partial S} dW \quad (13)$$

Notice that dV has a random component (the dW term). Our goal is to eliminate it.

Constructing the Hedging Portfolio

The brilliant insight of Black and Scholes: construct a portfolio Π consisting of:

- One derivative V (long position)
- $-\frac{\partial V}{\partial S}$ shares of stock (short position)

The portfolio value is:

$$\Pi = V - \frac{\partial V}{\partial S} S$$

The change in the portfolio over a small time dt is:

$$d\Pi = dV - \frac{\partial V}{\partial S} dS$$

Now substitute Equation (13) for dV and $dS = \mu S dt + \sigma S dW$:

$$\begin{aligned} d\Pi &= \left(\frac{\partial V}{\partial t} + \mu S \frac{\partial V}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} \right) dt + \sigma S \frac{\partial V}{\partial S} dW \\ &\quad - \frac{\partial V}{\partial S} (\mu S dt + \sigma S dW) \end{aligned}$$

Now watch what happens—the dW terms cancel exactly:

$$\sigma S \frac{\partial V}{\partial S} dW - \frac{\partial V}{\partial S} \cdot \sigma S dW = 0$$

And the $\mu S \frac{\partial V}{\partial S} dt$ terms also cancel. We are left with:

$$d\Pi = \left(\frac{\partial V}{\partial t} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} \right) dt \quad (14)$$

This is remarkable!

The portfolio change $d\Pi$ has **no dW term**. All randomness has been eliminated! The portfolio is instantaneously risk-free. This “delta hedging” strategy neutralizes the random fluctuations of the stock price by holding exactly the right number of shares.

The No-Arbitrage Condition

A risk-free portfolio must earn the risk-free interest rate r (if it earned more or less, there would be an arbitrage opportunity—a risk-free profit). Therefore:

$$d\Pi = r \Pi dt$$

Substituting $\Pi = V - \frac{\partial V}{\partial S} S$ and Equation (14):

$$\left(\frac{\partial V}{\partial t} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} \right) dt = r \left(V - \frac{\partial V}{\partial S} S \right) dt$$

Dividing by dt and rearranging:

$$\frac{\partial V}{\partial t} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS \frac{\partial V}{\partial S} - rV = 0 \quad (15)$$

The Black-Scholes Partial Differential Equation

Key Observations

The drift μ has disappeared. This is one of the most profound results in financial mathematics. The option price does not depend on the expected return of the stock! Two investors who completely disagree about where the stock is headed—but agree on its volatility—will agree on the option's fair price. This is because the hedging strategy eliminates exposure to the stock's direction; only its volatility matters.

The equation is a PDE, not an SDE. By constructing the hedging portfolio, we transformed a stochastic problem into a deterministic one. The Black-Scholes PDE is a second-order linear parabolic PDE, structurally similar to the *heat equation* from physics.

Further Reading

- [Wikipedia: Black-Scholes Equation](#) — Includes the connection to the heat equation.
- [Investopedia: Delta Hedging](#) — Practical explanation of how the hedging portfolio works.
- [Wikipedia: Heat Equation](#) — The physics analogue, for those interested in the mathematical structure.

Solving the Black-Scholes Equation

The Black-Scholes PDE (Equation (15)) must be solved subject to boundary conditions determined by the option type.

Boundary Conditions

For a **European call**:

$$\text{At expiration: } V(S, T) = \max(S - K, 0)$$

$$\text{As } S \rightarrow 0 : V(0, t) = 0$$

$$\text{As } S \rightarrow \infty : V(S, t) \sim S - Ke^{-r(T-t)}$$

For a **European put**:

$$\begin{aligned} \text{At expiration: } V(S, T) &= \max(K - S, 0) \\ \text{As } S \rightarrow 0: V(0, t) &= Ke^{-r(T-t)} \\ \text{As } S \rightarrow \infty: V(S, t) &= 0 \end{aligned}$$

Transformation to the Heat Equation

The standard method is to transform the Black-Scholes PDE into the classical **heat equation**, whose solution is well-known from physics and mathematics.

Introduce the substitutions:

$$x = \ln\left(\frac{S}{K}\right), \quad \tau = \frac{1}{2}\sigma^2(T-t), \quad V(S, t) = K v(x, \tau)$$

Note that x measures how far the stock is from the strike (in log terms), and τ is a rescaled time-to-expiration that runs *forward* (from 0 at expiration to $\frac{1}{2}\sigma^2T$ at time 0).

After substitution (the algebra is tedious but mechanical), we seek v in the form:

$$v(x, \tau) = e^{\alpha x + \beta \tau} u(x, \tau)$$

where α and β are chosen to eliminate lower-order terms. Setting:

$$\alpha = -\frac{1}{2}(k-1), \quad \beta = -\frac{1}{4}(k+1)^2, \quad k = \frac{2r}{\sigma^2}$$

the PDE reduces to the **heat equation**:

$$\frac{\partial u}{\partial \tau} = \frac{\partial^2 u}{\partial x^2} \tag{16}$$

Solution via the Heat Kernel

The heat equation has the well-known fundamental solution (Green's function):

$$u(x, \tau) = \frac{1}{\sqrt{4\pi\tau}} \int_{-\infty}^{\infty} u(s, 0) \exp\left(-\frac{(x-s)^2}{4\tau}\right) ds$$

The initial condition $u(s, 0)$ comes from transforming the option's payoff at expiration. Evaluating this integral involves splitting it at the exercise boundary, completing the square in each exponential, and recognizing the resulting integrals as cumulative normal probabilities $N(\cdot)$. Transforming back to the original variables yields the Black-Scholes formula.

The Black-Scholes Pricing Formulas

European Call Option

$$C(S, t) = S N(d_1) - K e^{-r(T-t)} N(d_2) \quad (17)$$

where:

$$d_1 = \frac{\ln(S/K) + \left(r + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \quad (18)$$

$$d_2 = d_1 - \sigma\sqrt{T-t} = \frac{\ln(S/K) + \left(r - \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \quad (19)$$

European Put Option

$$P(S, t) = K e^{-r(T-t)} N(-d_2) - S N(-d_1) \quad (20)$$

Interpretation of the Terms

Understanding what each piece “means” helps build intuition:

- $N(d_2)$ is approximately the probability (under the risk-neutral measure) that the option expires in the money, i.e., $\Pr(S_T > K)$. So $K e^{-r(T-t)} N(d_2)$ is the present value of paying K at expiration, weighted by the probability you actually have to pay it.
- $N(d_1)$ is the **delta** of the option—the expected fraction of a share you receive, conditional on the option being exercised, adjusted for the log-normal distribution. So $S N(d_1)$ is the present value of receiving the stock, weighted by the probability of exercise.
- $K e^{-r(T-t)}$ is the **present value** of the strike price, discounted from expiration back to today at the risk-free rate.
- The call price is the difference: what you expect to receive ($S N(d_1)$) minus what you expect to pay ($K e^{-r(T-t)} N(d_2)$).

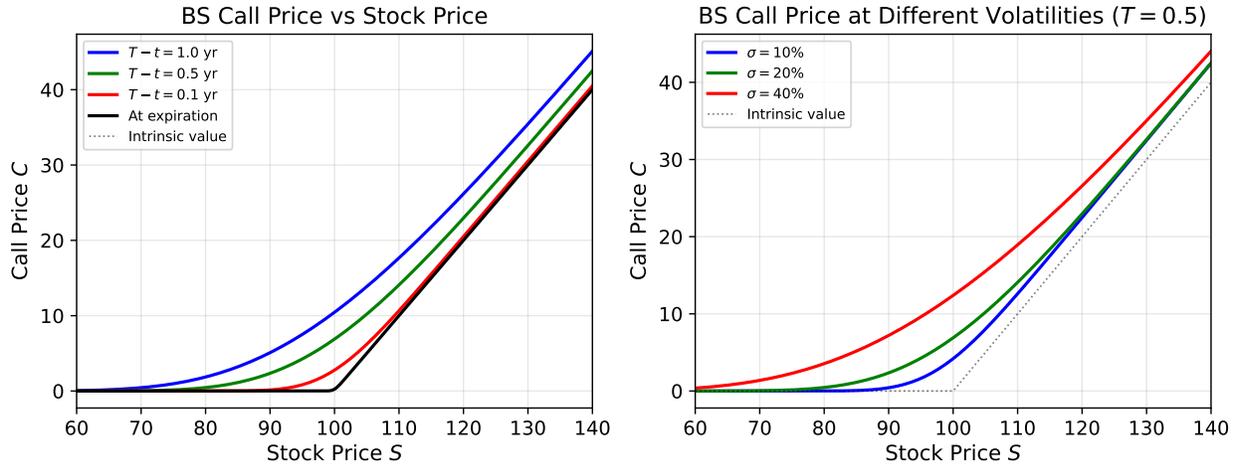


Figure 6: Black-Scholes call price as a function of stock price S ($K = 100$, $r = 5\%$). Left: Different times to expiration—as expiration approaches, the curve converges to the kinked payoff function. Right: Different volatilities—higher σ increases the option value, especially near the strike. The dotted line shows intrinsic value $\max(S - K, 0)$; the gap above it is the **time value**.

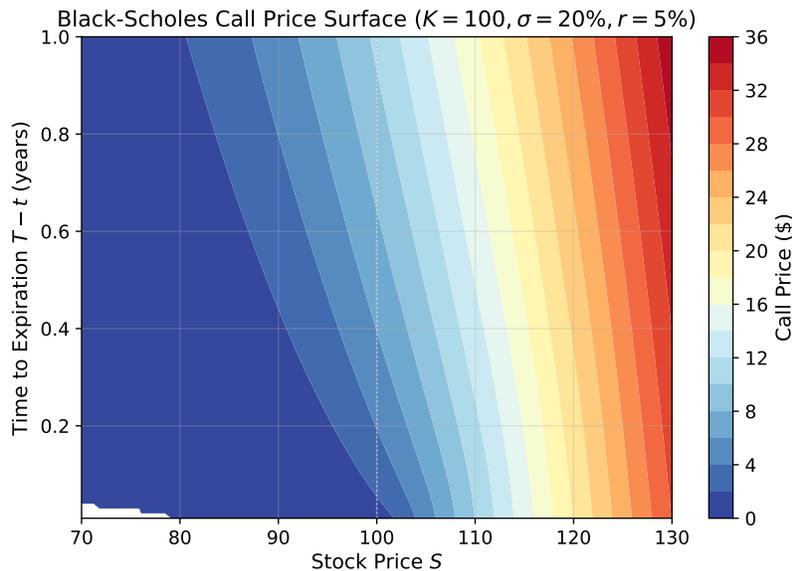


Figure 7: Heat map of the Black-Scholes call price as a function of both stock price S and time to expiration $T - t$. Deep in-the-money options (high S) with long time remaining are most valuable. The price surface smoothly transitions to the kinked payoff as $T - t \rightarrow 0$.

The Five Inputs

The formula requires exactly five inputs:

Symbol	Input	Observable?
S	Current stock price	Yes (market data)
K	Strike price	Yes (contract terms)
$T - t$	Time to expiration	Yes (contract terms)
r	Risk-free interest rate	Approximately (Treasury yields)
σ	Volatility	No — the only unobservable input

Since σ is the only unobservable input, it plays a special role. Traders often *invert* the formula: given the market price of an option, they solve for the σ that makes Black-Scholes match the market price. This is called the **implied volatility**.

Numerical Example

Suppose: $S = \$100$, $K = \$100$ (at the money), $T - t = 1$ year, $r = 5\%$, $\sigma = 20\%$.

$$d_1 = \frac{\ln(100/100) + (0.05 + 0.02)(1)}{0.20 \cdot 1} = \frac{0 + 0.07}{0.20} = 0.35$$

$$d_2 = 0.35 - 0.20 = 0.15$$

$$N(0.35) \approx 0.6368, \quad N(0.15) \approx 0.5596$$

$$\begin{aligned} C &= 100 \times 0.6368 - 100 \times e^{-0.05} \times 0.5596 \\ &= 63.68 - 95.12 \times 0.5596 \approx 63.68 - 53.23 = \$10.45 \end{aligned}$$

The fair price of this call option is approximately **\$10.45**.

Put-Call Parity

Put-call parity is a fundamental no-arbitrage relationship between European call and put prices. It holds regardless of which pricing model you use.

Derivation

Consider two portfolios at time t :

Portfolio A	One call option C + cash of $Ke^{-r(T-t)}$
Portfolio B	One put option P + one share of stock S

At expiration, if $S_T > K$:

$$\text{Portfolio A: } (S_T - K) + K = S_T$$

$$\text{Portfolio B: } 0 + S_T = S_T$$

If $S_T \leq K$:

$$\text{Portfolio A: } 0 + K = K$$

$$\text{Portfolio B: } (K - S_T) + S_T = K$$

Both portfolios give identical payoffs in *every* scenario. By the **law of one price** (no-arbitrage), they must cost the same today:

$$C + Ke^{-r(T-t)} = P + S \quad \Leftrightarrow \quad C - P = S - Ke^{-r(T-t)} \quad (21)$$

Put-Call Parity

Verification with Black-Scholes

We can verify that the Black-Scholes formulas satisfy put-call parity:

$$\begin{aligned} C - P &= [S N(d_1) - Ke^{-r(T-t)} N(d_2)] - [Ke^{-r(T-t)} N(-d_2) - S N(-d_1)] \\ &= S [N(d_1) + N(-d_1)] - Ke^{-r(T-t)} [N(d_2) + N(-d_2)] \end{aligned}$$

Since $N(x) + N(-x) = 1$ for all x :

$$C - P = S \cdot 1 - Ke^{-r(T-t)} \cdot 1 = S - Ke^{-r(T-t)} \quad \checkmark$$

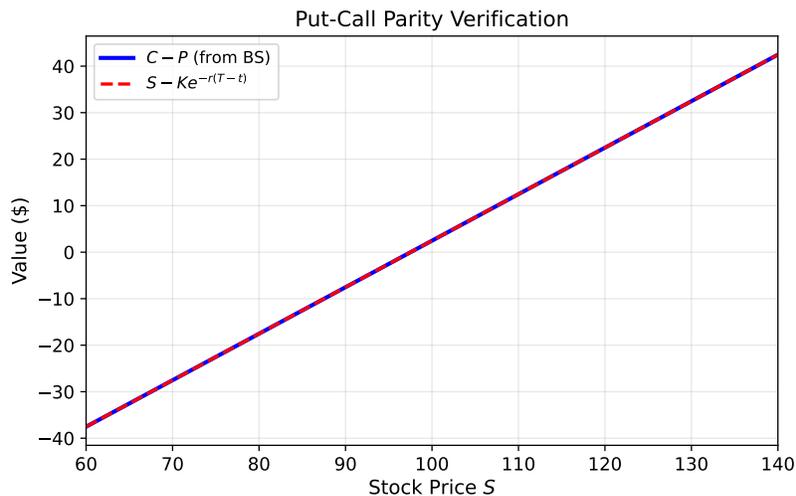


Figure 8: Verification of put-call parity: the blue line ($C - P$ computed from Black-Scholes) lies exactly on the red dashed line ($S - Ke^{-r(T-t)}$). If these lines ever diverged in the real market, there would be an arbitrage opportunity.

The Greeks: Sensitivity Analysis

The “Greeks” measure how the option price responds to changes in each input. They are essential tools for risk management, hedging, and understanding option behavior.

Delta (Δ)

Delta measures how much the option price changes per \$1 change in the stock price:

$$\Delta_{\text{call}} = \frac{\partial C}{\partial S} = N(d_1) \quad \Delta_{\text{put}} = \frac{\partial P}{\partial S} = N(d_1) - 1 \quad (22)$$

Derivation sketch: Differentiating $C = S N(d_1) - K e^{-r(T-t)} N(d_2)$ with respect to S requires the product rule and chain rule:

$$\frac{\partial C}{\partial S} = N(d_1) + S N'(d_1) \frac{\partial d_1}{\partial S} - K e^{-r(T-t)} N'(d_2) \frac{\partial d_2}{\partial S}$$

Using the identity $S N'(d_1) = K e^{-r(T-t)} N'(d_2)$ (verifiable by direct computation) and $\frac{\partial d_1}{\partial S} = \frac{\partial d_2}{\partial S} = \frac{1}{S\sigma\sqrt{T-t}}$, the last two terms cancel, leaving $\Delta = N(d_1)$.

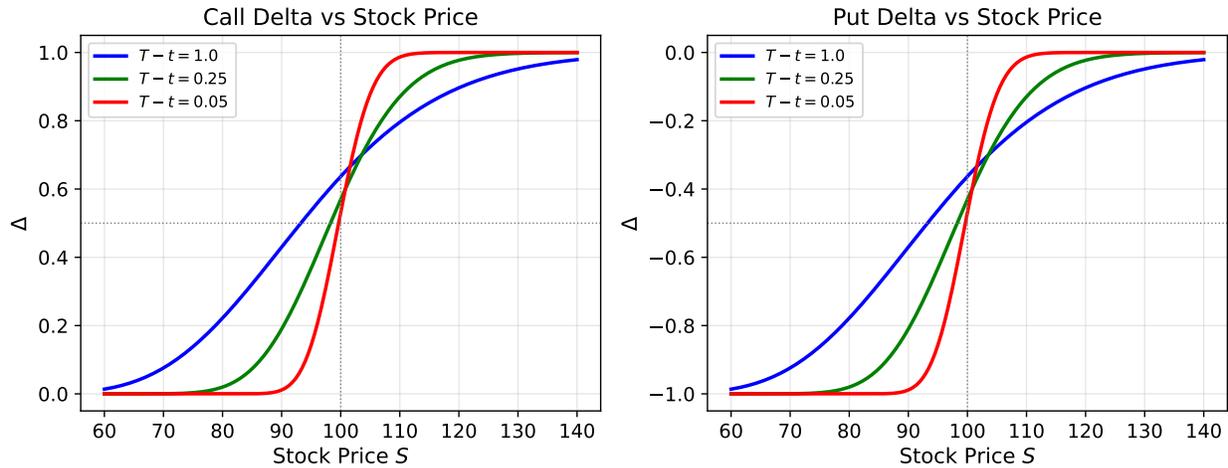


Figure 9: Delta as a function of stock price for calls (left) and puts (right) at different times to expiration. Near expiration (red), delta becomes a step step function—the option is either worthless or fully in the money. With more time remaining (blue), the transition is gradual. At-the-money options have $|\Delta| \approx 0.5$.

Interpretation: Call delta ranges from 0 (deep out of the money) to 1 (deep in the money). An at-the-money option has delta ≈ 0.5 . Delta also tells you the hedge ratio: to delta-hedge one call, sell Δ shares of stock.

Gamma (Γ)

Gamma measures how fast delta changes as the stock price moves:

$$\Gamma = \frac{\partial^2 C}{\partial S^2} = \frac{\partial^2 P}{\partial S^2} = \frac{N'(d_1)}{S\sigma\sqrt{T-t}} \quad (23)$$

where $N'(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$ is the standard normal PDF.

Gamma is always positive for long options and is highest for at-the-money options near expiration. High gamma means the hedge ratio is changing rapidly—you need to rebalance more frequently. Gamma is identical for calls and puts with the same strike and expiration.

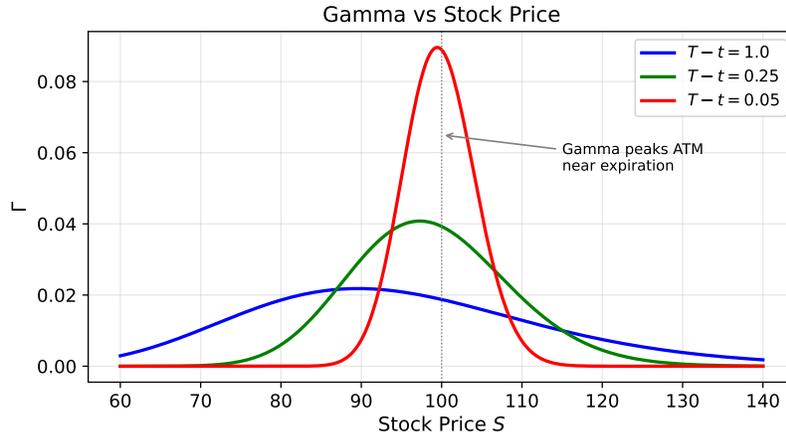


Figure 10: Gamma peaks sharply at-the-money near expiration (red curve). This is when delta is changing most rapidly—a small move in the stock price causes a large change in delta, requiring frequent hedge rebalancing. Far from the strike, gamma is near zero.

Theta (Θ)

Theta measures the rate of time decay—how much the option loses per day just from time passing:

$$\Theta_{\text{call}} = -\frac{\partial C}{\partial(T-t)} = -\frac{S N'(d_1) \sigma}{2\sqrt{T-t}} - rK e^{-r(T-t)} N(d_2) \quad (24)$$

$$\Theta_{\text{put}} = -\frac{\partial P}{\partial(T-t)} = -\frac{S N'(d_1) \sigma}{2\sqrt{T-t}} + rK e^{-r(T-t)} N(-d_2) \quad (25)$$

Theta is almost always *negative* for long options—options lose value as time passes, all else equal. Time decay accelerates near expiration (notice the $1/\sqrt{T-t}$ term, which blows up as $T-t \rightarrow 0$).

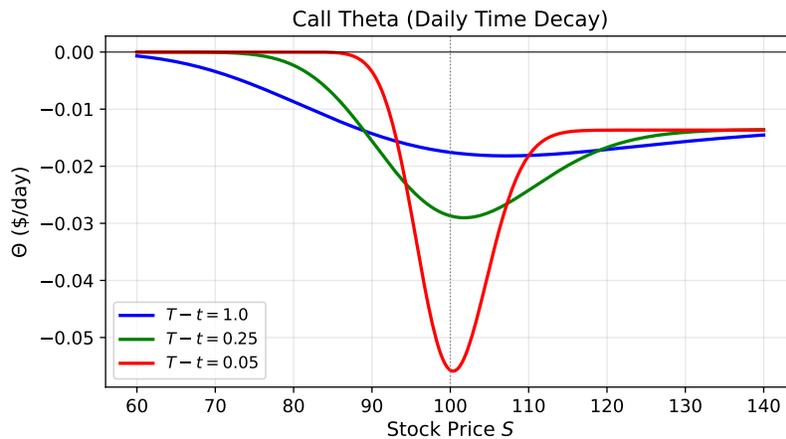


Figure 11: Daily time decay (theta) for a call option. The red curve ($T-t = 0.05$, about 2.5 weeks) shows devastating time decay near the strike—an ATM option might lose several cents per day. With more time remaining (blue), decay is gentler and more evenly spread.

The Theta-Gamma Tradeoff

The Black-Scholes PDE implies a fundamental relationship for a delta-hedged portfolio:

$$\Theta + \frac{1}{2}\sigma^2 S^2 \Gamma = rV$$

Since gamma is positive and theta is negative, this captures a deep truth: you cannot have convexity (gamma) for free. The “price” of gamma is theta—time decay.

Vega (\mathcal{V})

Vega measures sensitivity to volatility changes:

$$\mathcal{V} = \frac{\partial C}{\partial \sigma} = \frac{\partial P}{\partial \sigma} = S N'(d_1) \sqrt{T-t} \quad (26)$$

Vega is always positive: higher volatility increases the value of both calls and puts. This is intuitive—options have asymmetric payoffs (unlimited upside, limited downside), so more uncertainty is beneficial to the option holder. Vega is identical for calls and puts and is maximized for at-the-money options with long time to expiration.

In practice, vega is often the most important Greek. A trader who buys an option is implicitly betting on volatility, and vega quantifies that exposure.

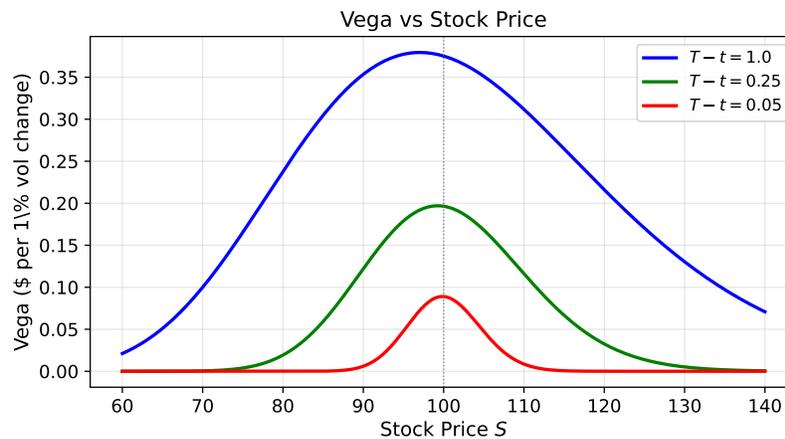


Figure 12: Vega (per 1% volatility change) as a function of stock price. Unlike gamma, vega is *highest* for longer-dated options (blue)—they have more time for volatility changes to affect the price. Vega always peaks at-the-money.

Rho (ρ)

Rho measures sensitivity to the risk-free interest rate:

$$\rho_{\text{call}} = \frac{\partial C}{\partial r} = K(T-t)e^{-r(T-t)}N(d_2) \quad (27)$$

$$\rho_{\text{put}} = \frac{\partial P}{\partial r} = -K(T-t)e^{-r(T-t)}N(-d_2) \quad (28)$$

Rho is typically the least significant Greek for short-dated options. It becomes more important for long-dated options (LEAPs) where discounting effects are more pronounced.

Summary Table

Greek	Measures	Call Formula	Sign (Long)
Δ	$\partial V/\partial S$	$N(d_1)$	+ (call), - (put)
Γ	$\partial^2 V/\partial S^2$	$\frac{N'(d_1)}{S\sigma\sqrt{T-t}}$	Always +
Θ	$\partial V/\partial t$	$-\frac{SN'(d_1)\sigma}{2\sqrt{T-t}} - rKe^{-r\tau}N(d_2)$	Usually -
\mathcal{V}	$\partial V/\partial \sigma$	$SN'(d_1)\sqrt{T-t}$	Always +
ρ	$\partial V/\partial r$	$K\tau e^{-r\tau}N(d_2)$	+ (call), - (put)

Further Reading

- [Investopedia: The Greeks](#) — Practical explanation with examples of each Greek.
- [Options Education: Greeks Tutorial](#) — Interactive tools from the Options Industry Council.

Risk-Neutral Pricing Framework

The fact that μ vanishes from the Black-Scholes PDE hints at something deeper. This leads to the **risk-neutral pricing framework**—one of the most powerful ideas in modern finance.

The Risk-Neutral Measure

In the real world, stocks have an expected return μ that compensates for risk. Under **risk-neutral pricing**, we replace the real-world probability measure \mathbb{P} with a *risk-neutral* measure \mathbb{Q} under which the stock grows at the risk-free rate r :

$$dS = rS dt + \sigma S dW^{\mathbb{Q}}$$

The fundamental result is:

$$V(S, t) = e^{-r(T-t)} \mathbb{E}^{\mathbb{Q}}[\text{Payoff}(S_T)] \quad (29)$$

Risk-Neutral Pricing Formula: The price of any derivative equals the expected payoff under \mathbb{Q} , discounted at the risk-free rate.

This is remarkable: we can price options *without knowing the stock's expected return or investors' risk preferences*.

Deriving Black-Scholes via Risk-Neutral Pricing

For a call option under \mathbb{Q} :

$$C = e^{-r(T-t)} \mathbb{E}^{\mathbb{Q}}[\max(S_T - K, 0)]$$

Under \mathbb{Q} , we have:

$$\ln S_T \sim \mathcal{N}\left(\ln S + \left(r - \frac{\sigma^2}{2}\right)(T-t), \sigma^2(T-t)\right)$$

Evaluating the expectation by splitting the integral at $S_T = K$, completing the square, and identifying cumulative normal probabilities yields the same Black-Scholes formula (Equations (17)–(19)).

Girsanov's Theorem (Sketch)

The mathematical justification for changing measures is **Girsanov's theorem**. Define the *market price of risk*:

$$\lambda = \frac{\mu - r}{\sigma}$$

This measures the excess return per unit of volatility (the Sharpe ratio). Girsanov's theorem states that:

$$dW^{\mathbb{Q}} = dW + \lambda dt$$

is a Brownian motion under the new measure \mathbb{Q} . The Radon-Nikodým derivative relating the two measures is:

$$\frac{d\mathbb{Q}}{d\mathbb{P}} = \exp\left(-\lambda W(T) - \frac{1}{2}\lambda^2 T\right)$$

This change of measure simultaneously removes the stock's risk premium and makes pricing independent of risk preferences.

Further Reading

- [Wikipedia: Risk-Neutral Measure](#) — Explains the concept with examples.
- [Wikipedia: Girsanov's Theorem](#) — The mathematical foundation for changing probability measures.
- Shreve, S.E. (2004). *Stochastic Calculus for Finance II*, Chapters 5–6 — The authoritative textbook treatment.

Limitations and Extensions

The Black-Scholes model is elegant and useful, but its assumptions are violated in real markets. Understanding these violations motivates more sophisticated models.

Non-Constant Volatility

The most significant limitation. Empirical evidence shows volatility changes over time and varies across strike prices (the **volatility smile/skew**).

Local Volatility (Dupire, 1994): Let $\sigma = \sigma(S, t)$ be a deterministic function of stock price and time, calibrated to match observed market option prices.

Stochastic Volatility (Heston, 1993): Volatility itself is random:

$$dv = \kappa(\theta - v) dt + \xi\sqrt{v} dW_v$$

where $v = \sigma^2$ is the variance, κ is the speed of mean reversion, θ is the long-run variance, and ξ is the “vol-of-vol.” The Heston model captures the empirical observation that volatility tends to cluster (high-vol periods followed by more high-vol) and mean-revert.

SABR Model (Hagan et al., 2002): Popular in interest rate markets, provides approximate closed-form solutions for the implied volatility smile.

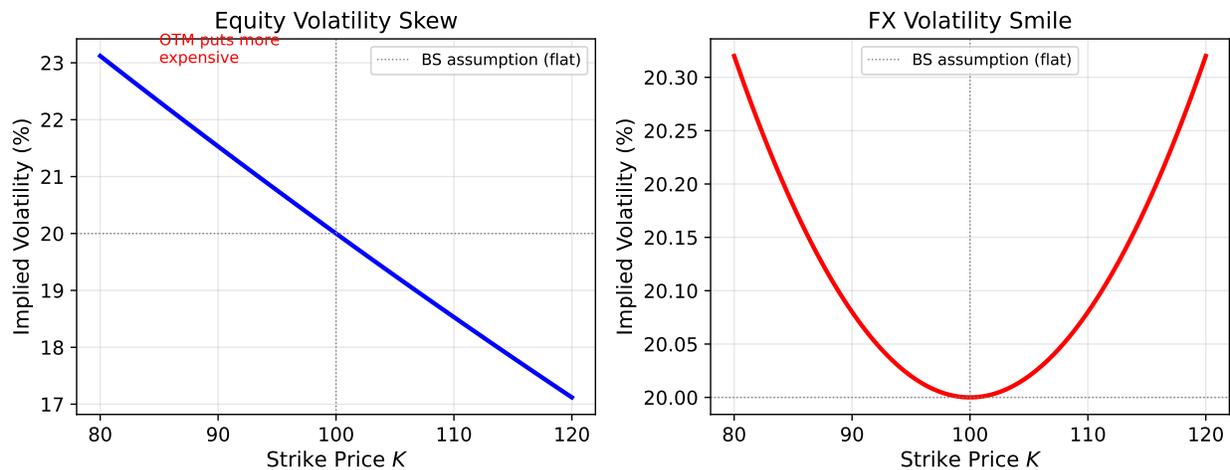


Figure 13: Left: The equity volatility **skew**—implied volatility is higher for low strikes (OTM puts) than high strikes, reflecting the market’s fear of crashes. Right: The FX volatility **smile**—symmetric, with higher vol for both deep OTM and deep ITM options. The gray dashed line shows the Black-Scholes assumption of constant volatility, which clearly fails in practice.

Jump Risk

Stock prices exhibit sudden jumps (earnings announcements, crashes) that GBM cannot capture. **Merton’s Jump-Diffusion Model** (1976) adds a Poisson jump process:

$$\frac{dS}{S} = (\mu - \lambda \bar{k}) dt + \sigma dW + J dN$$

where N is a Poisson process with intensity λ (average number of jumps per year), J is the random jump size, and $\bar{k} = \mathbb{E}[J]$. This produces heavier tails (more extreme moves) than the log-normal distribution.

Transaction Costs and Discrete Hedging

Black-Scholes assumes continuous, costless hedging. In reality, every trade incurs costs, and you can only rebalance at discrete times. Leland (1985) showed this effectively increases the volatility:

$$\sigma_{\text{adj}}^2 = \sigma^2 \left(1 + \sqrt{\frac{2}{\pi}} \cdot \frac{c}{\sigma\sqrt{\Delta t}} \right)$$

where c is the proportional transaction cost and Δt is the rebalancing interval.

Dividends

For stocks paying a continuous dividend yield q , replace S with $Se^{-q(T-t)}$ in the formulas:

$$C = Se^{-q(T-t)}N(d_1) - Ke^{-r(T-t)}N(d_2)$$

where d_1 and d_2 use $r - q$ instead of r in the numerator.

Non-Normal Returns

Empirical returns have heavier tails (more extreme events) and negative skewness (crashes are more common than equivalent rallies) compared to the normal distribution. Extensions include variance-gamma processes, normal-inverse Gaussian models, and generalized hyperbolic distributions.

Machine Learning Approaches

Recent research uses neural networks and deep learning to learn pricing functions directly from market data, potentially capturing dynamics that parametric models miss. Active areas include neural approximation of pricing functions, reinforcement learning for optimal hedging, and generative models for simulating realistic price paths.

Further Reading

- [Wikipedia: Heston Model](#) — The most widely-used stochastic volatility model.
- [Wikipedia: Volatility Smile](#) — Why Black-Scholes implied vol varies by strike.
- [Deep Hedging \(arXiv\)](#) — Using neural networks for hedging in incomplete markets.

References

1. Black, F. and Scholes, M. (1973). “The Pricing of Options and Corporate Liabilities.” *Journal of Political Economy*, 81(3), 637–654.
2. Merton, R.C. (1973). “Theory of Rational Option Pricing.” *Bell Journal of Economics and Management Science*, 4(1), 141–183.
3. Cox, J.C. and Rubinstein, M. (1985). *Options Markets*. Prentice-Hall.
4. Hull, J.C. (2021). *Options, Futures, and Other Derivatives*. 11th Edition. Pearson.
5. Shreve, S.E. (2004). *Stochastic Calculus for Finance II: Continuous-Time Models*. Springer.
6. Natenberg, S. (2014). *Option Volatility and Pricing*. 2nd Edition. McGraw-Hill.
7. Heston, S.L. (1993). “A Closed-Form Solution for Options with Stochastic Volatility.” *Review of Financial Studies*, 6(2), 327–343.
8. Dupire, B. (1994). “Pricing with a Smile.” *Risk*, 7(1), 18–20.
9. Merton, R.C. (1976). “Option Pricing When Underlying Stock Returns Are Discontinuous.” *Journal of Financial Economics*, 3(1–2), 125–144.

10. Hagan, P.S., Kumar, D., Lesniewski, A.S. and Woodward, D.E. (2002). “Managing Smile Risk.” *Wilmott Magazine*, 84–108.
11. Leland, H.E. (1985). “Option Pricing and Replication with Transaction Costs.” *Journal of Finance*, 40(5), 1283–1301.
12. Øksendal, B. (2003). *Stochastic Differential Equations*. 6th Edition. Springer.