

Pricing a real call option on the S&P 500.

$$C = S_0 \Phi(d_1) - K e^{-rT} \Phi(d_2)$$

CONTRACT **SPX261120C07900000**

SPX European call · strike 7,900 · expires Nov 20, 2026

$S_0 = 7,414$

$K = 7,900$

$T = 0.523$ yr

$r = 0.036$

$\sigma = 0.1946$

$C_{\text{mkt}} = \$180.30$

Begin with Part 1 →

Nathaniel Freed · May 2026 · [download PDF](#)



FTC and Taylor approximation

 $S_0 = 7,414$ SPX index

 $K = 7,900$ strike

 $T = 0.523$ yr time to expiry

 $r = 0.036$ 13-wk T-bill

 $C_{\text{mkt}} = \$180.30$ market mid

§1.1

Estimating σ from history

Eight-year window, daily log-returns, annualized.

Lookback_days = 2018 (≈ 8 calendar years). Long enough for statistical stability ($n \approx 2,017$ daily returns) and to span the calm of 2019, the COVID crash, the 2022 bear market, and the 2023–2026 rally. Shorter windows are dominated by the current regime; much longer windows contaminate the estimate with structural conditions no longer in force.

$$r_t = \ln(S_t/S_{t-1}), \quad \sigma = s_{\text{daily}} \cdot \sqrt{252}.$$

	VALUE
Calendar days requested	2,018
Daily returns computed	2,017

	VALUE
Daily std. dev. of log-returns	0.012257
Annualized volatility σ	0.1946 (19.46%)

ANNUALIZED VOLATILITY

$$\sigma = 0.1946$$

§1.2

Taylor series for Φ

Build the series for e^x , substitute, integrate term by term, then evaluate at d_1 and d_2 .

(I) Taylor series for e^x at $x = 0$

$$e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \frac{x^4}{4!} + \dots = \sum_{n=0}^{\infty} \frac{x^n}{n!}.$$

(II) Substitute $x = -t^2/2$

$$e^{-t^2/2} = \sum_{n=0}^{\infty} \frac{(-1)^n t^{2n}}{2^n n!} = 1 - \frac{t^2}{2} + \frac{t^4}{8} - \frac{t^6}{48} + \frac{t^8}{384} - \dots$$

(III) Integrate term-by-term to get $\Phi(z)$

$$\Phi(z) = \frac{1}{2} + \frac{1}{\sqrt{2\pi}} \int_0^z e^{-t^2/2} dt = \frac{1}{2} + \frac{1}{\sqrt{2\pi}} \sum_{n=0}^{\infty} \frac{(-1)^n z^{2n+1}}{2^n n! (2n+1)}.$$

$$\Phi(z) = \frac{1}{2} + \frac{1}{\sqrt{2\pi}} \left[z - \frac{z^3}{6} + \frac{z^5}{40} - \frac{z^7}{336} + \frac{z^9}{3456} - \dots \right].$$

(IV) Compute d_1 , d_2 , and evaluate Φ

$$\begin{aligned}\sigma^2 &= (0.1946)^2 = 0.0379, \\ \sigma^2/2 &= 0.0190, \\ r + \sigma^2/2 &= 0.0550, \\ (r + \sigma^2/2)T &= 0.0550 \times 0.523 = 0.0288, \\ S_0/K &= 7414/7900 = 0.9385, \\ \ln(S_0/K) &= -0.0635, \\ \sqrt{T} &= 0.7232, \\ \sigma\sqrt{T} &= 0.1407.\end{aligned}$$

$$d_1 = \frac{\ln(S_0/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}} = \frac{-0.0635 + 0.0288}{0.1407} = -0.2466.$$

$$d_2 = d_1 - \sigma\sqrt{T} = -0.2466 - 0.1407 = -0.3873.$$

$$T_0 = -0.2466, \quad |T_0| = 0.2466$$

$$T_1 = -z^3/6 = +0.0025, \quad |T_1| = 0.0025$$

$$T_2 = +z^5/40 = -0.0000, \quad |T_2| = 0.000023 < 0.0001 \Rightarrow \text{stop (2 terms)}$$

$$\Phi(d_1) = \frac{1}{2} + 0.3989 \times (-0.2441) = 0.4026.$$

$$T_0 = -0.3873, \quad |T_0| = 0.3873$$

$$T_1 = +0.0097, \quad |T_1| = 0.0097$$

$$T_2 = -0.0002, \quad |T_2| = 0.0002$$

$$T_3 = +0.0000, \quad |T_3| = 3.9 \times 10^{-6} < 0.0001 \Rightarrow \text{stop (3 terms)}$$

$$\Phi(d_2) = \frac{1}{2} + 0.3989 \times (-0.3778) = 0.3493.$$

$\Phi(D_1)$

0.4026

$\Phi(D_2)$

0.3493

§1.3

Hand-computed call price

$$C = S_0 \Phi(d_1) - K e^{-rT} \Phi(d_2).$$

DISCOUNT FACTOR e^{-rT} via Taylor series

$$\begin{aligned} rT &= 0.036 \times 0.523 = 0.0188, \\ x &= 0.0188, \quad x^2 = 0.0004, \quad x^2/2 = 0.0002, \\ e^{-0.0188} &\approx 1 - 0.0188 + 0.0002 - 0.0000 = 0.9814. \end{aligned}$$

TERM 1 $S_0 \cdot \Phi(d_1)$

$$S_0 \Phi(d_1) = 7414 \times 0.4026 = 2984.8764.$$

TERM 2 $K \cdot e^{-rT} \cdot \Phi(d_2)$

$$\begin{aligned} K e^{-rT} &= 7900 \times 0.9814 = 7753.0600, \\ K e^{-rT} \Phi(d_2) &= 7753.0600 \times 0.3493 = 2708.1439. \end{aligned}$$

$$C = 2984.8764 - 2708.1439 = 276.7325.$$

HAND-COMPUTED CALL PRICE

$$C = \$276.73$$

Python verification

scipy.stats.norm.cdf for Φ , full precision.

```
import numpy as np
from scipy.stats import norm

def black_scholes_call(S0, K, r, sigma, T):
    d1 = (np.log(S0/K) + (r + sigma**2/2) * T) / (sigma * np.sqrt(T))
    d2 = d1 - sigma * np.sqrt(T)
    return S0 * norm.cdf(d1) - K * np.exp(-r*T) * norm.cdf(d2)

# Sanity check (expect ~10.45)
print(black_scholes_call(100, 100, 0.05, 0.20, 1)) # → 10.4506

# Project parameters
S0, K, r, sigma, T = 7414, 7900, 0.036, 0.1946, 0.523
print(black_scholes_call(S0, K, r, sigma, T)) # → 277.29
```

	TAYLOR (HAND)	NORM.CDF (PYTHON)	Δ
$\Phi(d_1)$	0.4026	0.4025	0.0001
$\Phi(d_2)$	0.3493	0.3491	0.0002

	HAND	PYTHON
C	\$276.7325	\$277.2896
difference		\$0.5571

Market gap – model vs \$180.30

	VALUE
Model C (hand)	\$276.7325
Market C	\$180.30
Gap (model – market)	+\$96.4325

The market price sits **\$96.43 below** the model — a discount of roughly 35%. Three factors plausibly account for the gap. First, $\sigma = 19.5\%$ was estimated from eight years of realized daily returns, which include the COVID crash and the 2022 bear market; the option market is forward-looking and is currently pricing closer to 13–14% implied vol for OTM SPX calls at this maturity. Second, the formula assumes the index pays no dividends, but SPX carries a ~1.3% effective dividend yield from its constituents; including dividends shaves several dollars off the model price. Third, the volatility skew/smile: out-of-the-money calls (the \$7,900 strike is \$486 above spot) trade at lower implied vols than at-the-money options, because hedging demand concentrates in puts and ATM calls.

Greeks, parity, portfolio protection

CARRYING FORWARD FROM PART 1

$$\sigma = 0.1946$$

$$d_1 = -0.2466$$

$$d_2 = -0.3873$$

$$\Phi(d_1) = 0.4026$$

$$\Phi(d_2) = 0.3493$$

$$C = \$276.7325$$

$$e^{-rT} = 0.9814$$

§2.1

Deriving Δ

(I) Product rule on C with respect to S_0

$$\frac{\partial C}{\partial S_0} = \Phi(d_1) + S_0 \varphi(d_1) \frac{\partial d_1}{\partial S_0} - K e^{-rT} \varphi(d_2) \frac{\partial d_2}{\partial S_0}.$$

(II) $\partial d_1 / \partial S_0 = \partial d_2 / \partial S_0$

$$d_1 = \frac{\ln(S_0/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}} \implies \frac{\partial d_1}{\partial S_0} = \frac{1}{S_0 \sigma\sqrt{T}}.$$

$$d_2 = d_1 - \sigma\sqrt{T} \implies \frac{\partial d_2}{\partial S_0} = \frac{\partial d_1}{\partial S_0} = \frac{1}{S_0 \sigma\sqrt{T}}.$$

(III) **Substitute and factor**

$$\frac{\partial C}{\partial S_0} = \Phi(d_1) + \frac{S_0 \varphi(d_1) - K e^{-rT} \varphi(d_2)}{S_0 \sigma \sqrt{T}}.$$

(IV) **Identity: So $\varphi(d_1) = K e^{-rT} \varphi(d_2)$**

$$\ln S_0 - \frac{d_1^2}{2} \stackrel{?}{=} \ln K - rT - \frac{d_2^2}{2} \iff \ln(S_0/K) + rT \stackrel{?}{=} \frac{d_1^2 - d_2^2}{2}. \quad (\star)$$

$$d_1^2 - d_2^2 = (d_1 - d_2)(d_1 + d_2) = \sigma \sqrt{T} (2d_1 - \sigma \sqrt{T}) = 2\sigma \sqrt{T} d_1 - \sigma^2 T.$$

$$\sigma \sqrt{T} d_1 = \ln(S_0/K) + (r + \frac{\sigma^2}{2})T \implies d_1^2 - d_2^2 = 2 \ln(S_0/K) + 2rT.$$

$$\frac{d_1^2 - d_2^2}{2} = \ln(S_0/K) + rT. \quad \checkmark (\star).$$

(V) **Conclude**

$$\Delta = \frac{\partial C}{\partial S_0} = \Phi(d_1) = 0.4026.$$

INTERPRETATION

For a 1-index-point rise in S_0 , the call price rises by ~\$0.4026. Holding one call gives ~40.3% of the dollar move of one unit of underlying.

Vega and Theta

SETUP $\varphi(d_1)$ via Taylor

$$d_1^2 = 0.0608, \quad d_1^2/2 = 0.0304,$$

$$e^{-0.0304} \approx 1 - 0.0304 + 0.0005 - 0.0000 = 0.9701,$$

$$\varphi(d_1) = \frac{1}{\sqrt{2\pi}} e^{-d_1^2/2} = 0.3989 \times 0.9701 = 0.3870.$$

(I) VEGA $\partial C/\partial \sigma = S_0 \sqrt{T} \cdot \varphi(d_1)$

$$S_0 \sqrt{T} = 7414 \times 0.7232 = 5361.8048,$$

$$\text{Vega (raw)} = 5361.8048 \times 0.3870 = 2075.0185 \quad (\text{per unit rise in } \sigma),$$

$$\text{Vega (per 1pp)} = 2075.0185 \times 0.01 = 20.7502.$$

(I) THETA $\partial C/\partial t = -S_0 \sigma \varphi(d_1) / (2\sqrt{T}) - r K e^{-rT} \Phi(d_2)$

$$S_0 \sigma = 7414 \times 0.1946 = 1442.7644,$$

$$S_0 \sigma \varphi(d_1) = 558.3498,$$

$$2\sqrt{T} = 1.4464,$$

$$\text{first term} = -558.3498/1.4464 = -386.0272,$$

$$rK = 0.036 \times 7900 = 284.4000,$$

$$rK e^{-rT} = 279.1102,$$

$$rK e^{-rT} \Phi(d_2) = 279.1102 \times 0.3493 = 97.4932,$$

$$\text{second term} = -97.4932,$$

$$\Theta_{\text{raw}} = -386.0272 - 97.4932 = -483.5204,$$

$$\Theta_{\text{day}} = -483.5204/365 = -1.3247.$$

VEGA (PER 1PP RISE IN Σ)

20.7502

Θ (PER CALENDAR DAY)

-1.3247

(II) SIGN INTERPRETATION

Vega > 0. A rise in σ increases the call's value: more uncertainty in S_T makes the payoff $\max(S_T - K, 0)$ more valuable, since upside is unbounded and downside is floored at zero.

$\Theta < 0$. The call loses value as calendar time passes — shrinking time-to-expiry leaves fewer opportunities for S to drift above $K = 7900$.

(iii) **Magnitude vs C = \$276.73:** $|\text{Vega}_{1\text{pp}}|/C = 20.7502/276.7325 \approx 7.5\%$ per 1pp shift in σ — materially sensitive to volatility. $|\Theta_{\text{day}}|/C = 1.3247/276.7325 \approx 0.48\%$ per calendar day; over the ~191 days to expiry the cumulative decay is a sizable fraction of the premium.

§2.3

Put-call parity

A no-arbitrage identity that prices the put from the call.

$$P = C - S_0 + Ke^{-rT}.$$

(I) Compute P

$$C - S_0 = 276.7325 - 7414.0000 = -7137.2675,$$

$$Ke^{-rT} = 7900 \times 0.9814 = 7753.0600,$$

$$P = -7137.2675 + 7753.0600 = 615.7925 > 0. \checkmark$$

PUT PRICE (PARITY-IMPLIED)

$$P = \$615.79$$

(ii)(a) Correct hedging instrument for a long-equity investor

A **long put** on the index. Its payoff $\max(K - S_T, 0)$ activates exactly when the long-equity position is losing money, capping the downside while leaving the upside intact. Calls move the wrong way for a holder who fears a decline.

(ii)(b) Non-insurance buyers of OTM SPX calls

- **Directional speculators.** A trader expecting SPX to rally above 7,900 by November can buy this call for \$180.30 of premium and obtain convex exposure to ~7,900 index points of notional — capped downside (the premium), unbounded upside.
- **Volatility traders / dealer desks running long- Γ or long-Vega books.** A delta-hedged long call has near-zero directional exposure but profits when realized vol exceeds implied (gamma scalping) and from a rise in implied vol (positive Vega). The underlying delta is hedged away with a short position in SPX futures.

§2.4

Hedged portfolio Δ

200 units of SPX exposure + 200 of the $K = 7,900$ puts.

(I) Original position Δ

$$\Delta_{\text{unit}} = \partial S_0 / \partial S_0 = 1 \implies \Delta_{\text{position}} = 200 \times 1 = 200.$$

(II) $\Delta_{\text{put}} = \Phi(d_1) - 1$

$$\Delta_{\text{put}} = 0.4026 - 1 = -0.5974.$$

(III) Apply linearity

$$\Delta_{\text{portfolio}} = \sum_i n_i \Delta_i = 200(1) + 200(-0.5974) = 200 - 119.4800 = 80.4800.$$

$$\frac{200 - 80.4800}{200} = \frac{119.4800}{200} = 0.5976 = 59.76\%.$$

Δ HEDGED PORTFOLIO

80.4800

REDUCTION VS UNHEDGED

59.76%

§2.5

Solving N for delta-neutral

Choose N puts so $\Delta_{\text{portfolio}} = 0$.

$$200(1) + N(\Phi(d_1) - 1) = 0 \implies 200 + N(-0.5974) = 0 \implies N = \frac{200}{0.5974} = 334.7841.$$

DELTA-NEUTRAL N (WITH $\Phi(D_i) = 0.4026$)

$$N = 334.7841$$

round to 335 puts

$$\Delta N = 334.7841 - 200 = 134.7841, \quad \frac{N}{200} = \frac{334.7841}{200} = 1.6739.$$

Where the model bends

Black-Scholes is a model, not the world. Three short sections: which assumptions routinely fail in real markets, what came before BS in 1973, and what came after to patch those failures.

§3.1

Two violated assumptions

1. Constant volatility

The Black-Scholes model uses one σ for every option on a stock, assuming volatility is the same for all strike and expiration date combinations. However, in reality, this is not entirely accurate. For one, volatility will differ for short-term vs long-term options. For another, volatility differs for deep OTM options. Puts at those extremes are underpriced by BS compared to the market which hurts the seller in the case of a crash (which BS assumes is essentially impossible).^[1]

2. Continuous price paths (no jumps)

The second assumption is that price teleportation is impossible. For a stock to drop from \$100 to \$90, it must pass through \$99, \$98, \$97... and every price in between. At each step, σ predicts how big the next move is — such that a large drop requires many such moves. Therefore, under BS, massive near-instant crashes are essentially impossible. In reality, these crashes do happen. Prices can teleport when news arrives between trades or when prices change overnight between market close and open.^[2]

A predecessor

Louis Bachelier's doctoral thesis at the Sorbonne derived the first equation for option prices. Stock prices were modeled as Brownian motion and option value was computed as the expected payoff under that distribution. His model differed from Black-Scholes in several ways, but the most crucial was the price process itself. Bachelier used *arithmetic* Brownian motion, where the price change ($S_t - S_0$) is Gaussian, which allowed the price to drift below zero. Black-Scholes switched to *geometric* Brownian motion and treated the log return as Gaussian, which kept the price strictly positive and matched the empirical lognormal distribution of stock returns.^[3]

A successor

The Heston model (1993) addresses the aforementioned Black-Scholes assumption of constant volatility. Heston treats variance itself as a random, mean-reverting process which drifts back toward a long-run level, thus giving volatility its own "volatility" while preventing it from running off to extremes. This better matches real markets, where vol spikes are temporary and equity indices show a persistent skew. The flexibility comes at the cost of five parameters to calibrate instead of BS's one, and harder hedging: because variance isn't directly tradeable, volatility exposure must be hedged via portfolios of other options rather than the underlying alone.^[4]

Where the model breaks, where it earns its keep

Two pages on Black-Scholes against the real world. Where it failed (LTCM, 1998), where it earns its keep (vanilla options markets quote in implied vol), and where that leaves the question of when to use it.

§4.1

Critique: LTCM, 1998



1994 **Founded**

John Meriwether launches Long-Term Capital Management. Myron Scholes and Robert Merton join the board. The fund pursues convergence trades at 25–30:1 on-balance-sheet leverage.



1995 - 1997 **Peak**

CONSTANT Σ ASSUMED

~40% annual returns. By 1997 the fund holds over \$1.25 trillion in notional derivative exposure on roughly \$5 billion of equity.



1997 **Scholes & Merton Nobel**

The pair share the Nobel Memorial Prize in Economic Sciences "for a new method to determine the value of derivatives" — the Black-Scholes-Merton model itself.

AUG 17, 1998 Russia defaults

§3.1 NO-JUMPS

Russia devalues the ruble and defaults on its domestic debt. A discrete shock the size of years of "normal" daily moves — exactly what continuous GBM cannot represent.

AUG 1998 LTCM loses 44% in a month

§3.1 CONSTANT Σ

Spreads across nearly every LTCM position diverge instead of converging. Correlations across "uncorrelated" trades spike to 1 and the firm cannot reduce risk because no counterparty will warehouse it.

SEP 23, 1998 Fed organizes the rescue

Fourteen banks invest \$3.65 billion under Federal Reserve supervision to wind LTCM down in an orderly fashion. The argument: a default would cascade through the entire derivatives complex.

2000 Wound down

Trades unwound, capital returned, LTCM dissolved. The post-mortem consensus: not that BS is wrong (it is), but that confidence in any continuous-path model underprices the tail.

Long-Term Capital Management (LTCM) was a hedge fund founded in 1994 by John Meriwether, with Nobel laureates Myron Scholes and Robert Merton on its board. The fund pursued convergence trades — long undervalued securities and short overvalued "similar" ones — with leverage of roughly 25–30:1 on-balance-sheet (\approx \$30 of borrowed assets controlled per \$1 of investor equity) and over \$1.25 trillion in notional derivative exposure (the size their swap and futures bets were calibrated to, not the cash they posted) built on just \$5 billion of equity. ^[1]

On August 17, 1998, Russia devalued the ruble and defaulted on its domestic debt, and spreads across nearly every LTCM position diverged instead of converged. The fund lost 44 percent in a single month, and by mid-September its equity was nearly wiped out. ^[2] An LTCM default would have sent shockwaves through the global banking

system, so on September 23, 1998, fourteen banks invested \$3.65 billion under Federal Reserve supervision to wind the fund down in an orderly fashion.

The failure traces directly to two Black-Scholes assumptions named in §3.1. First, the continuous-paths assumption broke: the Russian default was a jump, a single discrete event whose impact dwarfed any "normal" daily move. LTCM's models, calibrated to the calm of 1995–1997, treated returns as continuous Gaussian and assigned essentially zero probability to what actually happened. Second, the constant-volatility assumption failed: when σ spiked 5–10× simultaneously across uncorrelated markets, the risk budget built on prior-year volatility was wrong by an order of magnitude — and Jorion (2000) shows LTCM compounded this by using the same covariance matrix to measure and optimize risk, causing the firm to take its largest bets on whatever its miscalibrated model called safest.^[3]

§4.2

Defense: implied volatility as a quoting language

In liquid vanilla options markets, traders do not quote option prices in dollars. They quote in *implied volatility* — the σ that, plugged into Black-Scholes, recovers the market price. A trader observes a price, runs BS in reverse, and the resulting IV becomes the quote. BS functions as a universal translator that maps dollar prices into a single number comparable across strikes and maturities. It is not a truth claim about how markets work.

Why does this work? **First, everyone uses it.** Every trader and analyst has BS. Counterparties agree on the formula, observe market prices, and exchange IV quotes without negotiating over model choice. A more sophisticated model like Heston or jump-diffusion introduces parameter disputes that slow everything down.

Second, speed. BS prices and Greeks are closed-form, computable in microseconds. A market-maker quoting and re-hedging thousands of strikes at machine speed cannot afford the Fourier inversions or Monte Carlo runs that more realistic models require.

Third, the surface absorbs the wrongness. The skew and term structure named in §3.1 are themselves the system's continuous, market-priced correction layer. BS plus the IV surface gives traders a model they know is wrong, alongside a real-time correction that tells them exactly how it is wrong. That combination is more transparent than any heavier model whose wrongness is hidden inside its parameters. [4]

The defense, then, is that in vanilla liquid markets, a transparently wrong model with a market-updated correction layer beats an opaquely "right" one.

§4.3

Taking a position

Black-Scholes is not a pricing model. It is an imaginary equation built on imaginary tokens that everyone pretends has meaning, and so it does. Nobody on a real trading desk has believed BS genuinely represents reality since 1987 at the latest, when the post-crash volatility skew became a flashing warning sign above every BS calculator. The reasons why were shown in §3.1: constant σ and continuous paths, both known to be false for decades, and both still baked into the formula. The real model that prices SPX options is the volatility surface — the empirical correction the market layers on top of BS to undo what BS gets wrong. The defense in §4.2 reads as a rescue but actually concedes the point: it defends the convention of quoting in implied volatility, not the formula itself. Calling BS a pricing model at this point is generous. The industry uses BS coordinates because every desk has them, not because the formula is right — BS is a poor system grandfathered in, which cannot be easily removed because the entire financial system is built around it.

The natural pushback is that LTCM proves the danger of mistaking notation for reality, and I agree. §4.1 showed LTCM treating BS-calibrated risk numbers as truth claims at 25:1 leverage, sizing its \$5 billion of equity against trillion-dollar notional exposure under the assumption that the model's tails were the market's tails. They were not. The model assigned essentially zero probability to a Russian default, and the market

handed LTCM the bill five weeks later. The model was not merely wrong; it was wrong in a direction the IV surface had not yet had time to learn. Therefore, BS is appropriate exactly when you treat it as notation and refuse to treat it as reality. Use it to write the quote. Do not use it to size the position. The industry post-1998 calls this discipline "model risk" and builds whole desks to enforce it.

Research Note

Sources, one cross-checked claim, and a short reflection.

§R.1

Sources used

Sources consulted for Parts 3 and 4. Numbers in brackets match the footnote markers in the body text; click any footnote to jump here, and the back-arrows below return you to the citation site.

[1] TEXTBOOK ↩ used in §3.1

Hull, J. C. *Options, Futures, and Other Derivatives*. Pearson.

[2] PAPER ↩ used in §3.1

Rubinstein, M. (1994). "Implied Binomial Trees." *Journal of Finance* 49(3), 771–818.

[3] PAPER ↩ used in §3.1

Merton, R. C. (1976). "Option Pricing When Underlying Stock Returns Are Discontinuous." *Journal of Financial Economics* 3, 125–144.

[4] PAPER ↩ used in §3.1

Bates, D. (1996). "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options." *Review of Financial Studies* 9(1), 69–107.

[5] PAPER ↩ used in §3.2

Bachelier, L. (1900). *Théorie de la spéculation*. *Annales Scientifiques de l'É.N.S.* 17, 21–86.

[6] BOOK ↩ used in §3.2

Davis, M. & Etheridge, A. (2006). *Louis Bachelier's Theory of Speculation: The Origins of Modern Finance*. Princeton University Press.

[7] PAPER ↔ used in §3.3

Heston, S. L. (1993). "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." *Review of Financial Studies* 6(2), 327–343.

[8] PAPER ↔ used in §4.1

Jorion, P. (2000). "Risk Management Lessons from Long-Term Capital Management." *European Financial Management* 6(3), 277–300.

[9] BOOK ↔ used in §4.1

Lowenstein, R. (2000). *When Genius Failed: The Rise and Fall of Long-Term Capital Management*. Random House.

[10] BOOK ↔ used in §4.2

Gatheral, J. (2006). *The Volatility Surface: A Practitioner's Guide*. Wiley.

AI tools

Claude (Anthropic). Claude code was used for research and paper summaries as well as layout design and solution verification.

§R.2

One claim cross-checked

Popular accounts of LTCM widely cite "100:1" leverage at peak. Jorion (2000) puts on-balance-sheet leverage at 25–30:1 (\$125B in assets on roughly \$5B equity), with the higher "effective" figure coming from \$1.25 trillion in notional derivative exposure. I cross-checked Jorion against the Federal Reserve History essay on LTCM and the standard Wikipedia summary; all three agreed on the 25–30:1 on-balance-sheet figure plus the trillion-dollar notional, so I used Jorion's number in §4.1 with the notional caveat called out explicitly.

§R.3

Brief reflection

Sources mostly converged: the technical claims around Heston's parameter count, Merton's jump-diffusion mechanics, and LTCM's rescue numbers were consistent across the canonical references, so the harder work turned out to be framing rather than fact-checking. The biggest surprise was discovering Bachelier developed the mathematics of Brownian motion five years before Einstein used it for pollen-grain physics — a detail the standard textbook narrative on stochastic calculus tends to drop. The hardest section to write was §4.3, because sources cannot tell you what to think; the position emerged from synthesising the earlier sections rather than from any single citation.

Appendix

Downloads, source materials, and links.

PDF

Full report (PDF)

The complete report as a single PDF — for printing or offline reference during Friday's discussion.

NOTEBOOK

Part 1 Python notebook

The starter notebook with the σ estimation, Black-Scholes function, and sanity checks.

PDF

Part 1 hand work

Original handwritten/typeset work for Part 1 — Taylor series derivation, Φ evaluation, hand C.

PDF

Part 2.1 hand work

Original hand work for Part 2.1 — Δ derivation, the φ -identity proof, Vega/Theta, put-call parity.

PDF

Project brief

The original assignment from D. Dickmann, including parameters and grading criteria.

This site is a static report generated with Astro and deployed on Cloudflare Pages. The math is rendered with KaTeX; numerics match the values in the original Python notebook and hand-work PDFs. All theme tokens swap live via the bottom-right pill.