








Modelling Financial Markets

Seminar 12: Options and the Volatility Surface — Hands-on



Daniel Traian PELE
Drd. Antoaneta Amza, antoaneta.amza@csie.ase.ro

Bucharest University of Economic Studies
IDA Institute Digital Assets
Blockchain Research Center
AI4EFin Artificial Intelligence for Energy Finance
Romanian Academy, Institute for Economic Forecasting
MSCA Digital Finance

Today's Question and Seminar Structure

- Question of the seminar
  - **how much can we trust the numbers that option prices give us: Greeks, implied volatility, a fitted smile and the variance risk premium?**
- Deliverables
  - the answers to the A exercises
  - two figures (B1, B2) and one table per B exercise
  - the corrected C1 code
  - one interpretation sentence per exercise
- Assessment criteria
  - reproducibility: the notebook runs top to bottom and reproduces every reported number
  - correct inference: robust standard errors on overlapping windows, every estimate with its uncertainty
  - interpretation: a claim supported by the numbers, with its limitation
- Notebook for today: open the seminar notebook in Colab
  - each task names its notebook section
- Seminar: Prof. Daniel Traian Pele, danpele@ase.ro; Antoaneta Amza, PhD candidate, antoaneta.amza@csie.ase.ro

Exercise Map

Exercise	Question	Type	Model
A1	why are gamma and vega of one expiry proportional?	Solved	–
A2	which volatility reproduces a market price of 4.50?	Solved	–
B1	how fast does the hedging error shrink with rebalancing?	Solved	–
B2	is the S&P 500 variance risk premium significant and stable?	Solved	–
C1	what is wrong in an artificial intelligence (AI) answer?	Proposed	A1, A2
B3 (home)	is the VIX (Cboe Volatility Index) an efficient forecast of volatility?	Proposed	B2
B4 (after the lecture)	how precisely does SVI describe the Bitcoin smile?	Solved	–
Extension packs	A3–A8, B5–B8, C2, C3 (optional, after the seminar)	both	named on each task

- **[Solved]**: full solution in the slides and in the notebook, a model to copy; **[Proposed]**: you solve it, the solution is discussed in class
- Part A: derivations on paper; Part B: estimation and inference on real data, each ending with an interpretation question; Part C: open questions and the critique of an AI (Artificial Intelligence) answer

Data Used

Series	Symbol, field	Period	Units	Used in
S&P 500	`GSPC.INDX`, close	1990–2026	index points	B2, B3, B5, C2, C3
VIX, VIX3M	`VIX.INDX`, `VIX3M.INDX`, close	1990 / 2007–2026	annualised %	B2, B3, B5, C2
Bitcoin	`BTC-USD.CC`, close	2021–2026	USD	B7, C2
Bitcoin options	Deribit option snapshot: strike, type, bid, ask, mark price (BTC), mark IV (implied volatility), forward, index	22 Sep 2026, 07:51 UTC	BTC; % for IV	A3, A4, B4, B6, B8
DVOL	Deribit volatility index, daily	2021–2026	annualised %	B7, C2
Simulated	Geometric Brownian Motion (GBM) paths, seed 42	–	–	B1

- UTC: Coordinated Universal Time; USD: US (United States) dollar; BTC: Bitcoin. Closing prices run through 18 September 2026; Bitcoin option quotes are a single snapshot from Deribit, a crypto-options exchange.
- VIX3M: the 3-month analogue of the VIX (Cboe 3-Month Volatility Index); bid/ask: best buy/sell quotes; mark: the exchange's reference price
- Two series together: join the prices on common days first, then compute returns; C3 also needs the Kenneth French daily risk-free rate and the replication package of the case study

Setup: Run and Verify

```
# notebook section "Data" defines
# load_close(name) and deribit_chain()
px = pd.concat([load_close('sp500'),
                load_close('vix')],
               axis=1, join='inner').dropna()
print(len(px), px.index[0], px.index[-1])
print(px.iloc[[0, -1]])
c = deribit_chain()
print(len(c), c['expiry'].nunique())
```

- Expected output: 9 245 common days, 1990-01-02 to 2026-09-18

date	S&P 500	VIX
1990-01-02	359.69	17.24
2026-09-18	7 650.50	14.81

- Option snapshot: 986 contracts (calls and puts), 11 expiries
- If your counts differ, stop and check the symbols, the join and the end date

Notation Used in Every Exercise (1/2): Prices, Strikes and Time

- S: price of the underlying today; S_T: its price at expiry T; K: strike (exercise price)
- $\tau = T - t$: time to expiry in years (90 calendar days: $\tau = 90/365$; the Cboe formulas and the code write T for this duration); r: continuously compounded risk-free rate
- σ: volatility, the annualised standard deviation of log returns (a decimal in formulas, in % when quoted)
- Forward price and log-moneyness

$$F = Se^{r\tau}, \qquad k = \ln(K/F)$$

  - $k < 0$: strike below the forward (out of the money, OTM puts); $k > 0$: above it (OTM calls)
- σ_{imp}: implied volatility; total implied variance $w = \sigma_{imp}^2 \tau$
- Φ, φ: distribution function and density of the standard Normal distribution, $\varphi(x) = e^{-x^2/2}/\sqrt{2\pi}$
- "Vol points": percentage points of volatility; vega per vol point = vega per unit of σ divided by 100

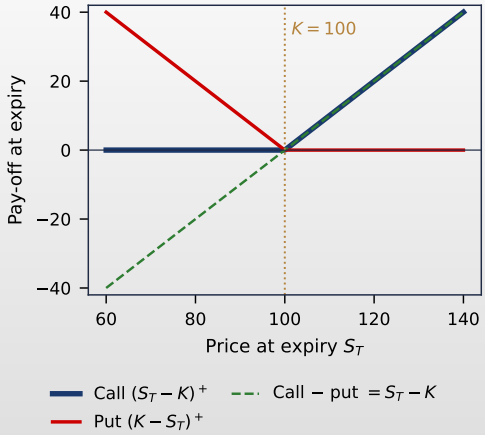
Notation Used in Every Exercise (2/2): Black–Scholes and a Worked Line

- One worked line for the Black–Scholes call price C Black & Scholes (1973) (no dividends): $S = 100$, $K = 95$, $\tau = 0.25$, $r = 2\%$, $\sigma = 20\%$
  - $\sigma\sqrt{\tau} = 0.1000$, $d_1 = 0.6129$, $d_2 = 0.5129$, $C = 7.21$
- Daily log return in %: $r_t = 100 \ln(P_t / P_{t-1})$, P_t the close of day t; A: days per year (252 trading or 365 calendar days)
- Realised variance over the next h days, annualised, in $\%^2$

$$RV_{t,t+h} = \frac{A}{h} \sum_{i=1}^{h} r_{t+i}^2$$

  - $\sqrt{RV_{t,t+h}}$: realised volatility in % per year, comparable with the VIX (the 30-day implied volatility of the S&P 500) Carr & Wu (2009)

Options and Their Pay-offs



- **European call**: the right to buy at K on the expiry date; pay-off
$$(S_T - K)^+, \qquad x^+ = \max(x, 0)$$
- **European put**: the right to sell at K; pay-off $(K - S_T)^+$
- Moneyness
  - at the money (ATM): $K \approx S$
  - OTM: a call with $K > S$, a put with $K < S$; zero pay-off if exercised today; here moneyness is relative to spot, not forward
- Green line: a long call and a short put pay $S_T - K$ in every state, like a forward bought at K

Put–Call Parity and No-Arbitrage Bounds

- **Put–call parity** (no dividends; C, P: today's prices of a call and a put with the same K and τ)

$$C - P = S - Ke^{-r\tau}$$

  - $Ke^{-r\tau}$: today's value of K paid at expiry
  - a call plus cash $Ke^{-r\tau}$ and a put plus one share both pay $\max(S_T, K)$, so they cost the same; no model is needed
  - with $D = e^{-r\tau}$ and $F = Se^{r\tau}$: $C - P = DF - DK$, a straight line in K (A4)

- **No-arbitrage bounds**

$$(S - Ke^{-r\tau})^+ \le C \le S, \qquad (Ke^{-r\tau} - S)^+ \le P \le Ke^{-r\tau}$$

  - a price outside them is a data error or an arbitrage: a portfolio that costs nothing and never loses

- Example: $S = K = 100$, $r = 4\%$, $\tau = 0.5$: $Ke^{-r\tau} = 98.02$
  - a call at 7.00 implies the put $7.00 - 100 + 98.02 = 5.02$

Geometric Brownian Motion and the Black–Scholes Price

- **The stock as a Geometric Brownian Motion** (GBM, Chapter 11)

$$dS_t = \mu S_t\, dt + \sigma S_t\, dW_t$$

  - μ: drift (expected return per year); W_t: standard Brownian motion, $W_t - W_s \sim N(0, t-s)$
  - the log return over Δ years is $N\big((\mu - \sigma^2/2)\Delta,\ \sigma^2\Delta\big)$, so $S_t > 0$

- **Black–Scholes** Black & Scholes (1973): constant σ and r, continuous trading without costs

$$C = S\Phi(d_1) - Ke^{-r\tau}\Phi(d_2), \qquad P = Ke^{-r\tau}\Phi(-d_2) - S\Phi(-d_1)$$

  - $\Phi(d_2)$: risk-neutral probability that the call ends in the money; $\Phi(d_1)$: shares held in the replicating portfolio
  - μ does not appear: Δ shares and cash, rebalanced continuously, replicate the option whatever the drift
  - European call without dividends, $r \geq 0$: the price rises with S, σ, τ and r and falls with K
  - European put: the price falls with S and r and rises with σ and K; in τ it can fall deep in the money when $r > 0$

The Greeks

- **Greeks**: sensitivities of the call price, Black–Scholes without dividends

$$\Delta = \frac{\partial C}{\partial S} = \Phi(d_1), \qquad \Gamma = \frac{\partial^2 C}{\partial S^2} = \frac{\varphi(d_1)}{S\sigma\sqrt{\tau}}, \qquad \mathcal{V} = \frac{\partial C}{\partial \sigma} = S\varphi(d_1)\sqrt{\tau}$$

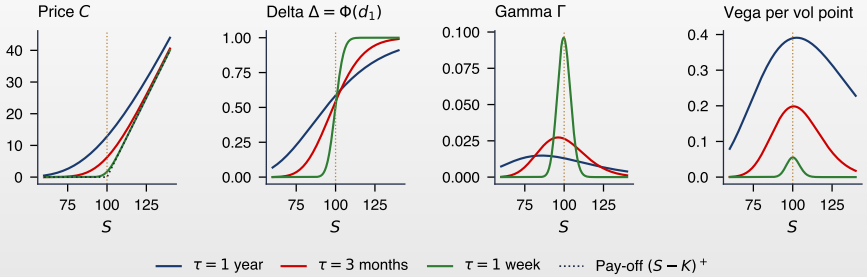
  - delta $\Delta \in (0,1)$: shares to hold per call; if S rises by 1, the call gains about Δ
  - gamma $\Gamma > 0$: how fast delta moves; vega $\mathcal{V} > 0$: price change per unit of σ (per vol point: $/100$)
  - theta $\Theta = \partial C/\partial t$, usually negative: the time value decays as expiry approaches

- **Second-order Taylor expansion** of the price change

$$\delta C \approx \Delta\,\delta S + \tfrac{1}{2}\Gamma\,(\delta S)^2 + \mathcal{V}\,\delta\sigma + \Theta\,\delta t$$

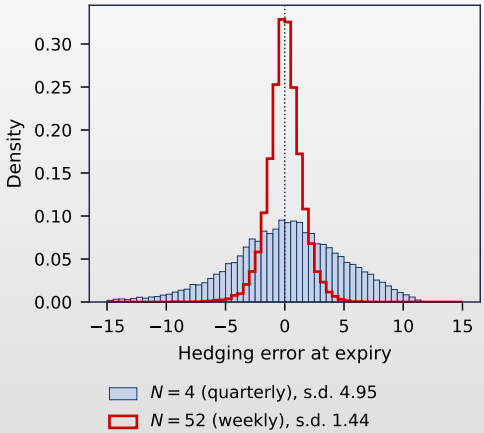
- **Portfolio Greeks add up**: n_i units of option i give $\Delta_\Pi = \sum_i n_i \Delta_i$, $\Gamma_\Pi = \sum_i n_i \Gamma_i$
  - delta–gamma hedge: a second option sets $\Gamma_\Pi = 0$, then shares ($\Delta = 1$, $\Gamma = 0$) set $\Delta_\Pi = 0$ (A7)

Price and Greeks as Functions of S



- Call with $K = 100$, $r = 2\%$, $\sigma = 30\%$, for three times to expiry τ
- Price: above the pay-off by the time value, which vanishes as $\tau \to 0$
- Delta: an S-shaped curve from 0 to 1, steeper near expiry
- Gamma and vega: largest near K; gamma grows as expiry approaches, vega shrinks

Delta Hedging and the Hedging Error



- **Delta hedging** a sold call
  - receive C_0, buy Δ_0 shares, keep $B_0 = C_0 - \Delta_0 S_0$ in cash at rate r
  - at each of N dates reset Δ, paying from cash: the portfolio is self-financing
- **Hedging error** at expiry

$$e_T = \Delta_{N-1} S_T + B_T - (S_T - K)^+$$

  - B_T: cash at T, interest included; $e_T = 0$ with continuous rebalancing in the Black–Scholes model
  - between dates gamma makes the hedge imperfect: $\text{s.d.}(e_T) \propto 1/\sqrt{N}$
- Chart (simulated): one-year ATM call, $\sigma = 30\%$; 13 times more trades cut the s.d. from 4.95 to 1.44, close to the factor $\sqrt{13} = 3.6$

Itô's Lemma, the Black–Scholes Equation, the Hedging P&L

- **Itô's lemma** (Chapter 11) for an option value $V(S,t)$ when $dS = \mu S\,dt + \sigma S\,dW$

$$dV = \left(V_t + \mu S V_S + \tfrac{1}{2}\sigma^2 S^2 V_{SS}\right)dt + \sigma S V_S\,dW_t$$

  - V_t, V_S, V_{SS}: partial derivatives; the term $\frac{1}{2}\sigma^2 S^2 V_{SS}$ comes from $(dS)^2 = \sigma^2 S^2 dt$
- **Black–Scholes equation**: holding $\Delta = V_S$ removes dW and μ; a riskless position earns r

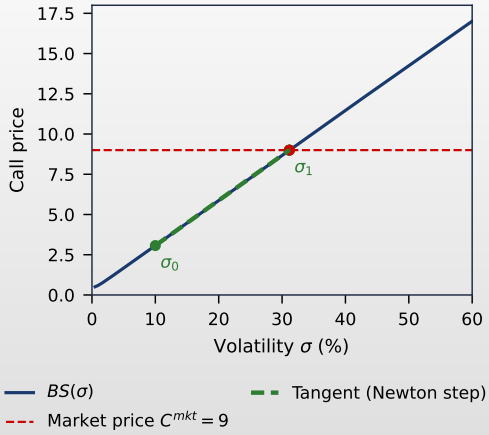
$$V_t + \tfrac{1}{2}\sigma^2 S^2 V_{SS} + rSV_S - rV = 0$$

- **Hedging profit and loss (P&L) when volatility is wrong** (A5): sold at σ_{imp}, the price moves with σ_{real}

$$d\,\mathrm{P\&L} = r\,\mathrm{P\&L}\,dt + \tfrac{1}{2}\Gamma S^2\left(\sigma_{imp}^2 - \sigma_{real}^2\right)dt$$

  - the seller gains when realised volatility is below implied; the gain accrues through ΓS^2, so it depends on the path

Implied Volatility and Newton's Method



- S, K, τ, r are observed; σ is not. **Implied volatility** σ_{imp} solves

$$BS(\sigma_{imp}) = C^{mkt}$$

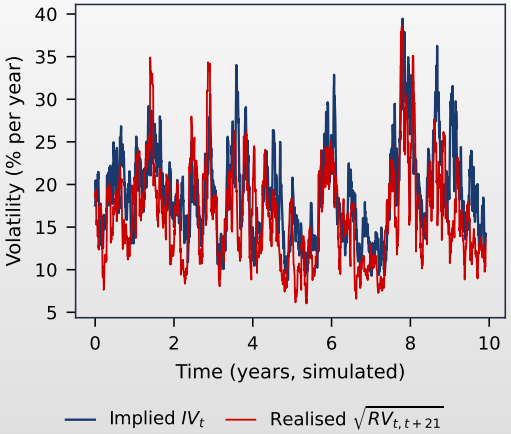
  - $BS(\sigma)$: Black–Scholes price; C^{mkt}: market price
  - unique, since vega > 0; it exists iff C^{mkt} lies strictly inside the no-arbitrage bounds

- **Newton's method**: follow the tangent from the iterate σ_k

$$\sigma_{k+1} = \sigma_k - \frac{BS(\sigma_k) - C^{mkt}}{\mathcal{V}(\sigma_k)}$$

- Chart: $S = K = 100$, $\tau = 0.5$, $r = 1\%$, $C^{mkt} = 9$: from $\sigma_0 = 10\%$ to $\sigma_1 = 31.14\%$, then 31.15%
- Question for the class: does a market price above the model price require a higher or a lower volatility?
  - Answer: higher, because vega > 0
- Traders quote options in σ_{imp}: comparable across strikes, maturities and assets

The Variance Risk Premium

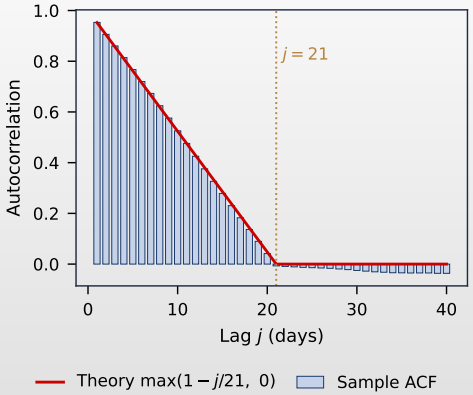


- **Variance risk premium** (VRP) Carr & Wu (2009): implied minus subsequently realised variance

$$VRP_t = IV_t^2 - RV_{t,t+h}$$

  - IV_t: implied volatility in % (VIX or DVOL); in vol points $IV_t - \sqrt{RV_{t,t+h}}$
  - the profit of a seller of variance insurance; positive on average if investors pay to insure high-volatility states
- Example: $VIX = 20$, realised volatility 15%: $VRP = 400 - 225 = 175\,\%^2$, i.e. 5 vol points
- Chart (simulated): implied variance 20% above the current instantaneous variance, used as a proxy for the expected 21-day variance; the premium is positive on 71% of days but negative in volatility spikes
- Variance swap (pays $RV - K_{var}$; VIX section) with vega notional 1: P&L $\approx (IV^2 - RV)/(2\,IV)$ vol points (C2)

Overlapping Windows and Newey–West Errors



- Consecutive $RV_{t,t+\mathbf{21}}$ share 20 of 21 days: even for independent and identically distributed (i.i.d.) returns their autocorrelation function (ACF) at lag j is $1 - j/21$ (chart)
- The naive s.e. $s/\sqrt{n}$ of a mean $\bar{x}$ is then too small; **Newey–West** Newey & West (1987) with L lags

$$\hat{\Omega} = \hat{\gamma}_{\mathbf{0}} + 2\sum_{j=\mathbf{1}}^{L}\Big(1 - \frac{j}{L+1}\Big)\hat{\gamma}_j, \qquad \mathrm{se}(\bar{x}) = \sqrt{\hat{\Omega}/n}$$

  - $\hat{\gamma}_j$: sample autocovariance at lag j; the Bartlett weights keep $\hat{\Omega} \geq 0$; take $L \geq h$
  - $t = \bar{x}/\mathrm{se}(\bar{x})$ is approximately $N(0,1)$: $|t| > 1.96$ rejects $H_{\mathbf{0}}$: mean $= 0$ at 5%
- Two periods: regress $x_t = c_{\mathbf{0}} + c_{\mathbf{1}} D_t + u_t$, $D_t = 1$ after the break; HAC (Heteroskedasticity and Autocorrelation Consistent) t-test of $c_{\mathbf{1}} = 0$
- Alternative: **block bootstrap** Künsch (1989), resampling blocks of consecutive days

Is a Forecast Unbiased? Mincer–Zarnowitz and Wald

- **Mincer–Zarnowitz regression** Mincer & Zarnowitz (1969) of the outcome on the forecast

$$RV_{t,t+\mathbf{21}} = \alpha + \beta\, VIX_t^{\mathbf{2}} + \varepsilon_t$$

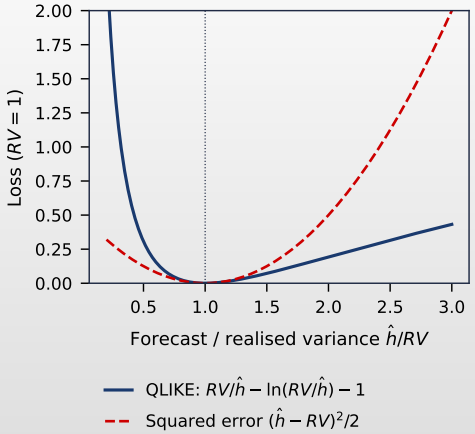
  - horizon: the VIX covers 30 calendar days, the target 21 trading days; both approximate one month, but the calendars are not equal
  - unbiased forecast: $\alpha = 0$ and $\beta = 1$ jointly; efficient: no other variable known at t adds information

- **Wald test** of q linear restrictions $H_{\mathbf{0}}$: $R\theta = c$

$$W = (R\hat{\theta} - c)^{\top} \left(R \hat{V} R^{\top}\right)^{-\mathbf{1}} (R\hat{\theta} - c) \sim \chi^{\mathbf{2}}_{q}$$

  - $\hat{\theta} = (\hat{\alpha}, \hat{\beta})$; $\hat{V}$: HAC covariance; here $R = I_{\mathbf{2}}$, $c = (0, 1)$, $q = 2$; reject at 5% if $W > 5.99$

Comparing Variance Forecasts: HAR-RV, QLIKE, Diebold–Mariano



- **HAR-RV in logs** (Heterogeneous AutoRegressive model of Realised Volatility) Corsi (2009)

$$\ln RV_{t,t+\mathbf{21}} = b_{\mathbf{0}} + \sum_{j \in \{\mathbf{1},\mathbf{5},\mathbf{22}\}} b_j \ln RV_t^{(j)} + u_t$$

  - $RV_t^{(j)}$: mean of the last j squared daily returns, annualised; level forecast $\exp(\widehat{\ln RV} + \hat{s}^{\mathbf{2}}/2)$, $\hat{s}^{\mathbf{2}}$: residual variance (log-normal correction; alternative: smearing, the mean of $e^{\hat{u}}$ on the training data)

- **QLIKE loss** (Quasi-LIKElihood) Patton (2011) of a variance forecast $\hat{h}$

$$L(RV, \hat{h}) = \frac{RV}{\hat{h}} - \ln \frac{RV}{\hat{h}} - 1 \geq 0$$

  - zero only if $\hat{h} = RV$; penalises under-prediction more (chart); the ranking is robust to noise in RV if RV is a conditionally unbiased proxy of the variance

- **Diebold–Mariano** Diebold & Mariano (1995): $d_t = L_t^A - L_t^B$, $DM = \bar{d}/\mathrm{se}_{HAC}(\bar{d}) \approx N(0,1)$

A1 [Solved]: One Identity for All the Greeks — Task

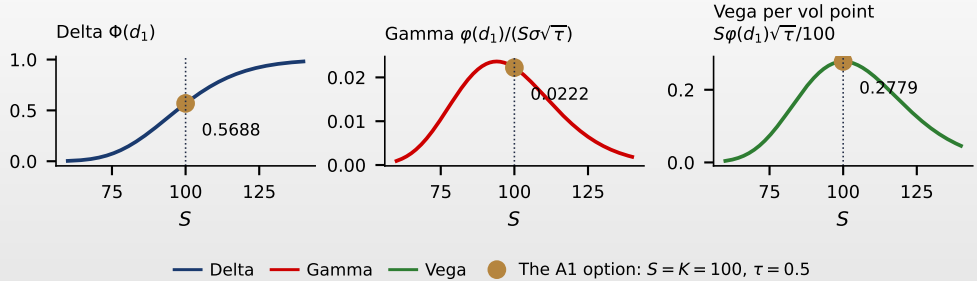
- **Questions**
  - why are the gamma and the vega of options with the same expiry proportional?
  - what are they for an at-the-money call?
- **Inputs**
  - The Black–Scholes call in "Geometric Brownian Motion and the Black–Scholes Price"; no dividends
  - Numbers: $S = K = 100$, six months ($\tau = 0.5$), $r = 3\%$, $\sigma = 25\%$; no data file
- **Tasks**
  1. Prove the identity $S\varphi(d_1) = Ke^{-r\tau}\varphi(d_2)$ by computing $d_1^2 - d_2^2$.
  2. Use the identity to show that $\Delta = \Phi(d_1)$ and vega $= S\varphi(d_1)\sqrt{\tau}$.
  3. Show that vega $= \Gamma S^2\sigma\tau$, with Γ in "The Greeks".
  4. Compute d_1, d_2, both sides of the identity, Δ, Γ, vega in two ways, the call, and the put from put–call parity.
- **Report**: the three proofs, a table of the numbers, and one sentence on what vega $= \Gamma S^2\sigma\tau$ implies for hedging gamma with options of the same expiry
- Notebook: section A1

Back Exercise list

A1 [Solved]: Solution (1/2) — Proofs

- Step 1: $\varphi(d_2)/\varphi(d_1) = e^{(d_1^2 - d_2^2)/2}$ and $d_1^2 - d_2^2 = (d_1 - d_2)(d_1 + d_2) = \sigma\sqrt{\tau} \cdot \dfrac{2\ln(S/K) + 2r\tau}{\sigma\sqrt{\tau}}$
- Hence $\varphi(d_2)/\varphi(d_1) = Se^{r\tau}/K$, i.e. $S\varphi(d_1) = Ke^{-r\tau}\varphi(d_2)$
- Step 2: $\partial_S C = \Phi(d_1) + [S\varphi(d_1) - Ke^{-r\tau}\varphi(d_2)]\,\partial_S d_1 = \Phi(d_1)$, since $\partial_S d_1 = \partial_S d_2$ and the bracket is 0 by Step 1
- Vega: $S\varphi(d_1)\partial_\sigma d_1 - Ke^{-r\tau}\varphi(d_2)\partial_\sigma d_2 = S\varphi(d_1)\,\partial_\sigma(d_1 - d_2) = S\varphi(d_1)\sqrt{\tau}$
- Step 3: $\Gamma S^2\sigma\tau = \dfrac{\varphi(d_1)}{S\sigma\sqrt{\tau}}\,S^2\sigma\tau = S\varphi(d_1)\sqrt{\tau} = \text{vega}$

A1 [Solved]: Solution (2/2) — Numbers and Chart



- Step 4: $\sigma\sqrt{\tau} = 0.1768$, $d_1 = (0.03 + 0.03125) \cdot 0.5/0.1768 = 0.1732$, $d_2 = -0.0035$; identity: $100 \cdot 0.3930 = 39.30$ and $98.51 \cdot 0.3989 = 39.30$
- $\Delta = 0.5688$, $\Gamma = 0.0222$; vega $= 100 \cdot 0.3930\sqrt{0.5} = 27.79 = \Gamma S^2 \sigma\tau = 27.79$ per unit of σ (0.2779 per vol point); $C = 7.76$, $P = C - 100 + 98.51 = 6.27$
- **Interpretation**: for one expiry vega and gamma are proportional, so a gamma-neutral book of same-expiry options is also vega-neutral, but only while the smile is flat (the same σ_{imp} for every K; extension A7)

MFM_ch12_seminar

A2 [Solved]: Implied Volatility by Newton's Method — Task

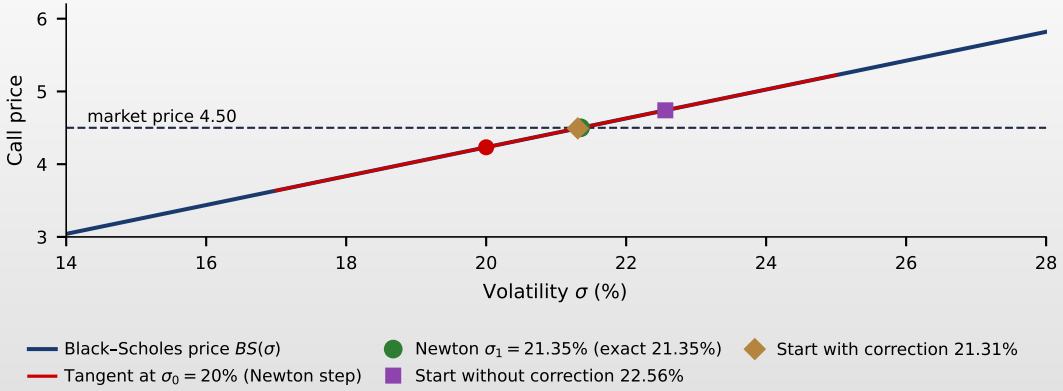
- **Questions**
  - which volatility makes the Black–Scholes price equal to a market price of 4.50?
  - Before calculating, write whether a price above the model price requires a higher or lower volatility, and explain your prediction using vega. Then check it with Newton.
- **Inputs**: a European call with $S = K = 100$, three months ($\tau = 0.25$), $r = 2\%$, market price 4.50; vega $\mathcal{V}(\sigma) = S\varphi(d_1)\sqrt{\tau}$ from A1
- **Tasks**
  1. Prove that $\sigma \mapsto BS(\sigma)$ is continuous and strictly increasing, with limits $(S - Ke^{-r\tau})^+$ as $\sigma \to 0$ and S as $\sigma \to \infty$, and state when an implied volatility exists.
  2. Derive the at-the-money-forward approximation $C \approx S\sigma\sqrt{\tau}/\sqrt{2\pi}$ Brenner & Subrahmanyam (1988).
  3. Invert it to get two starting values: from the price 4.50, and from 4.50 minus half the discounted forward intrinsic value $(S - Ke^{-r\tau})/2$.
  4. Run one Newton step $\sigma_1 = \sigma_0 - (BS(\sigma_0) - 4.50)/\mathcal{V}(\sigma_0)$ from $\sigma_0 = 20\%$ and compare σ_1 with the exact implied volatility.
- **Report**: the two starting values, σ_1 and the exact implied volatility; one sentence on why Newton converges so fast here and when it would fail
- Notebook: section A2

Back Exercise list

A2 [Solved]: Solution (1/3) — Existence and Approximation

- Step 1: $\partial_\sigma BS = S\varphi(d_1)\sqrt{\tau} > 0$ (A1): strictly increasing and continuous
- $\sigma \to 0$: $d_{1,2} \to \pm\infty$ according to the sign of $\ln(S/K) + r\tau$, so $BS \to (S - Ke^{-r\tau})^+$; $\sigma \to \infty$: $d_1 \to \infty$, $d_2 \to -\infty$, so $BS \to S$
- By the intermediate value theorem a unique σ_{imp} exists if and only if the price lies strictly inside the no-arbitrage bounds $(S - Ke^{-r\tau})^+ < C < S$
- Step 2: if $S = Ke^{-r\tau}$, $d_{1,2} = \pm\sigma\sqrt{\tau}/2$ and $C = S[\Phi(\sigma\sqrt{\tau}/2) - \Phi(-\sigma\sqrt{\tau}/2)] \approx S\,\sigma\sqrt{\tau}\,\varphi(0) = S\sigma\sqrt{\tau}/\sqrt{2\pi}$

A2 [Solved]: Solution (2/3) — Newton Step and Chart



MFM_ch12_seminar

A2 [Solved]: Solution (3/3) — Starting Values and the Newton Result

- Step 3: without correction $\sigma \approx 4.50\sqrt{2\pi}/(100 \cdot 0.5) = 22.56\%$
- With correction: $Ke^{-r\tau} = 99.501$, half the forward intrinsic value 0.249, $\sigma \approx (4.50 - 0.249)\sqrt{2\pi}/50 = 21.31\%$
- Step 4: $BS(20\%) = 4.2322$, $\mathcal{V} = 19.848$, so $\sigma_{\mathbf{1}} = 0.20 - (4.2322 - 4.50)/19.848 = 21.3495\%$; exact 21.3495%
- Units: $\mathcal{V} = 19.848$ per unit of σ, i.e. 0.198 per vol point; the step $0.2678/19.848 = 0.0135$ is 1.35 vol points
- **Interpretation**: near the money the price is almost linear in σ, so one step gives four decimals; deep out of the money vega ≈ 0 and Newton can overshoot: use a bracketing solver (Brent)

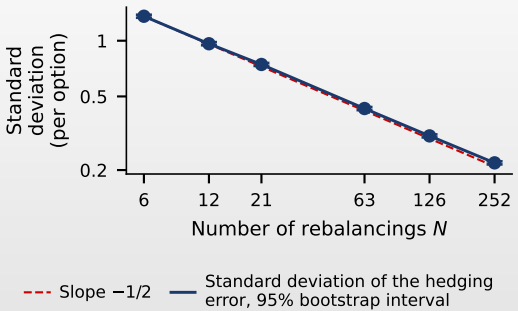
MFM_ch12_seminar

B1 [Solved]: Hedging Error and Rebalancing Frequency — Task

- **Question**: how fast does the hedging error of a delta-hedged call shrink as the dealer rebalances more often?
- **Inputs**
  - Simulated data: 10 000 paths of geometric Brownian motion, $S_0 = 100$, $\mu = 8\%$, $\sigma = 20\%$, 252 steps over three months (seed 42)
  - A three-month ATM call ($K = 100$, $r = 4\%$) sold at its Black–Scholes premium 4.49 and delta-hedged at $\sigma = 20\%$
- **Tasks**
  1. For each $N \in \{6, 12, 21, 63, 126, 252\}$, rebalance the hedge N times and compute on every path the hedging error at expiry: shares plus cash minus the pay-off.
  2. Compute the standard deviation of the error for each N and a 95% confidence interval (CI) by percentile bootstrap from 1 000 resamples of the paths Efron (1979).
  3. Regress $\ln(\text{s.d.})$ on $\ln N$ and give the slope with a bootstrap interval from the same resamples.
  4. Find the smallest N in the list for which the standard deviation is below 10% of the premium.
- **Report**: a table of standard deviations with intervals, the log-log chart and the slope; one sentence on whether the slope is consistent with $-1/2$ and what transaction costs would change
- Notebook: section B1

Back Exercise list

B1 [Solved]: Solution (1/2) — Chart and Table



- Step 1: on each path, sell the call at 4.49, buy Δ_0 shares, finance with cash at r; at each rebalancing date reset Δ; at expiry, error = shares + cash − pay-off
- Step 2: the s.d. falls from 1.356 ($N = 6$) to 0.219 ($N = 252$); the intervals are narrow because 10 000 paths give a precise s.d.

N	6	12	21	63	126	252
s.d.	1.356	0.963	0.744	0.430	0.306	0.219
95% CI	[1.331, 1.381]	[0.945, 0.983]	[0.731, 0.758]	[0.422, 0.438]	[0.300, 0.312]	[0.215, 0.223]

B1 [Solved]: Solution (2/2) — Slope and Interpretation

- Step 3: slope -0.489, 95% CI $[-0.496, -0.482]$: close to $-1/2$ and slightly above it at these finite N; the $1/\sqrt{N}$ law of the primer is asymptotic
- Step 4: the smallest N in the list with an s.d. below 10% of the premium is $N = 63$ (daily): $0.430/4.49 = 9.6\%$ of the premium
  - The lecture reports 9.5% with 20 000 paths; here 10 000 paths, so the gap is Monte Carlo error
  - Four times a day ($N = 252$) halves it again (ratio 1.97)
- **Interpretation**: with a fixed fee per trade, costs grow like N while the error falls like $1/\sqrt{N}$; the best frequency depends on the cost model and on the trade-off chosen between risk and cost

MFM_ch12_seminar

B2 [Solved]: The Variance Risk Premium of the S&P 500 — Task

- **Questions**
  - is the variance risk premium of the S&P 500 significantly positive?
  - was it the same before and after 2008?
- **Inputs**
  - Daily S&P 500 and VIX closes joined on common days, January 1990 – August 2026 (the last 21-day window ends in September 2026)
  - $RV_{t,t+\mathbf{21}}$ (Notation (2/2), $A = 252$, $h = 21$) and VRP_t in $\%^2$ and in vol points, as in "The Variance Risk Premium"
- **Tasks**
  1. Build $RV_{t,t+\mathbf{21}}$ from the next 21 log returns and VRP_t. Retain the common sample with complete future and past 21-return windows, including $RV_{t-\mathbf{21},t}$ used by the forecast comparison.
  2. Compute the mean of VRP_t, its Newey–West standard error with 21 lags Newey & West (1987), a 95% interval, and the naive standard error $s/\sqrt{n}$.
  3. Repeat step 2 in vol points and compute the share of days with a positive premium.
  4. Repeat for 1990–2007 and 2008–2026, then test equality by regressing VRP_t on a constant and a 2008–2026 dummy with Newey–West errors (21 lags).
- **Report**: a table by period and the chart; one sentence each
  - why is the naive standard error wrong here?
  - by how much is it wrong?
  - is the premium stable?
  - why is the variance premium less precise than the vol-point premium?
- Notebook: section B2

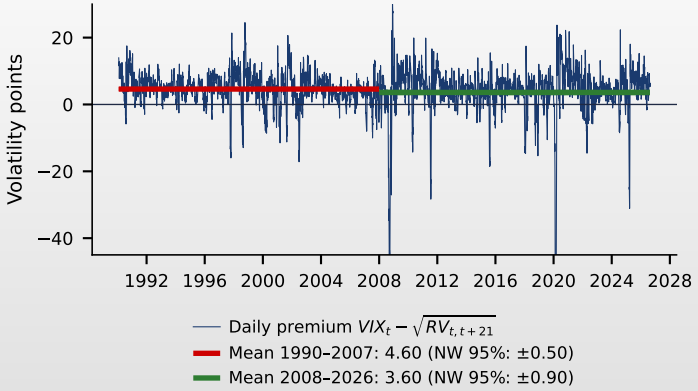
Back Exercise list

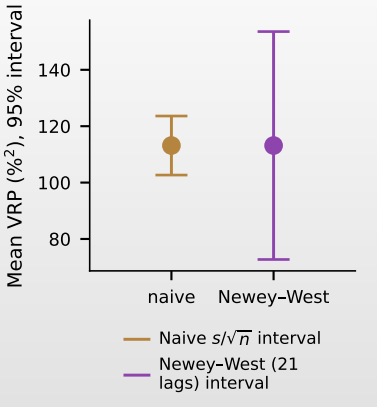
B2 [Solved]: Solution (1/2) — Table, Steps and Interpretation

	n	Mean VRP ($\%^2$)	95% CI	Vol points (s.e.)	Positive
1990–2026	9 203	113.1	[72.7, 153.6]	4.09 (0.27)	86%
1990–2007	4 516	149.3	[125.2, 173.5]	4.60 (0.25)	88%
2008–2026	4 687	78.3	[2.7, 153.9]	3.60 (0.46)	84%

- Steps 1–2: regress VRP_t on a constant with HAC standard errors (Bartlett, 21 lags: the overlap); interval = mean ± 1.96 s.e.; full sample $t = 113.1/20.6 = 5.49$: reject H_0 at 5%
- Naive s.e. 5.3 vs Newey–West 20.6: 3.9 times too small, because consecutive windows share 20 of 21 days
- Step 3: vol points 4.09 (s.e. 0.27), positive on 86% of days
- Step 4: 1990–2007 $t = 12.12$, 2008–2026 $t = 2.03$; difference -71.0 $\%^2$ ($p = 0.08$), in vol points -1.01 ($p = 0.06$)
- **Interpretation**: squaring magnifies the crisis windows (2008, 2020), so the variance premium has heavier tails and a wider interval than the vol-point premium; the premium is positive in both subperiods, but the difference in size is not significant at 5%: the data cannot tell whether it changed

B2 [Solved]: Solution (2/2) — Variance Premium over Time





MFM_ch12_seminar

C1 [Proposed]: Critique an AI Answer — the Answer

- A student asks an AI assistant: "Price a European call with 90 calendar days to expiry on a non-dividend-paying stock with $S = 100$, $K = 105$, volatility 20%, continuous rate $r = 4\%$; then give the put from put–call parity"
- The assistant answers with this code:

```
import numpy as np
from scipy.stats import norm
S, K, r, sigma = 100.0, 105.0, 0.04, 20         # sigma = 20%
T = 90 / 252                                    # 90 days
d1 = (np.log(S / K) + (r + 0.5 * sigma**2) * T) / (sigma * np.sqrt(T))
d2 = d1 - sigma * np.sqrt(T)
call = S * norm.cdf(d1) - K * np.exp(-r * T) * norm.cdf(d2)
put = call + S - K * np.exp(-r * T)             # put-call parity
print(f"call = {call:.2f}, put = {put:.2f}")
```

- "The code follows Black–Scholes exactly; the put comes from parity, so no second formula is needed"

Back Exercise list

C1 [Proposed]: Critique an AI Answer — Task

Solution discussed in class

- **Questions**
  - which three errors does the AI answer contain?
  - which check catches each of them?
- **Inputs**: the code on the previous slide; everything except three lines is correct
- **Tasks**
  1. Before running the code, estimate what a 90-day call 5% out of the money should cost: 0.2, 2 or 20.
  2. For each error, say what is wrong and why, and how you would detect it: a no-arbitrage bound, a unit check, today's primer, or a second method.
  3. Correct the three lines, run the code, report C and P, and check put–call parity numerically.
  4. Log the three errors in `AI_ERRORS.md` with the fields: what was asked, what the AI answered, what was wrong, how you found out, the correct version.
- **Report**: the corrected code and output; one sentence on which error still gives a plausible-looking price, and one on why "the put comes from parity" is not a check of the call
- **Model**: A1, A2 (Black–Scholes step by step)

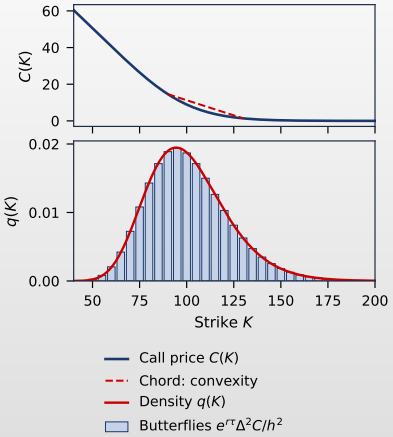
B3 [Proposed]: Is the VIX an Efficient Forecast? — Task

Solution discussed in class

- **Question**: is VIX_t^2 an unbiased and efficient forecast of the variance realised over the next 21 trading days?
- **Inputs**
  - The daily series of B2: $RV_{t,t+21}$ and VIX_t^2, 1990–2026; past variance $RV_{t-21,t}$
  - For step 4: the log HAR-RV, its level forecast and the QLIKE loss in "Comparing Variance Forecasts: HAR-RV, QLIKE, Diebold–Mariano" (each average $RV_t^{(j)}$ floored at 0.01 $\%^2$)
- **Tasks**
  1. Estimate $RV_{t,t+21} = \alpha + \beta\, VIX_t^2 + \varepsilon_t$ Mincer & Zarnowitz (1969) with Newey–West errors (21 lags); test $\alpha = 0$ and $\beta = 1$ separately and jointly (Wald, 5%), and report $\mathrm{corr}(\hat{\alpha}, \hat{\beta})$.
  2. Add $RV_{t-21,t}$ as a second regressor; then re-estimate in logs, $\ln RV$ on $\ln VIX^2$, and test $\beta = 1$.
  3. Instrument $\ln VIX_t^2$ with $\ln VIX_{t-21}^2$ Christensen & Prabhala (1998) (instrumental variable; just-identified GMM, Generalized Method of Moments; HAC covariance); report the first-stage t and test $\beta = 1$.
  4. From 2000, re-estimate every 21 days on an expanding window of already observed targets; compare VIX^2 with HAR-RV by QLIKE loss and a Diebold–Mariano test with HAC (21 lags) Diebold & Mariano (1995), Patton (2011).
- **Report**: all estimates and tests with decisions at 5%; one sentence each
  - why can both t-statistics be small while the joint test rejects?
  - what must the instrument satisfy?
  - does in-sample significance survive out of sample?
- **Model**: B2 (Newey–West inference)
- Notebook: section B3

Back Exercise list

Convexity in the Strike and the Risk-Neutral Density



- **Risk-neutral density** q of S_T: prices are discounted expectations under q

$$C(K) = e^{-r\tau} \int_K^\infty (x - K)\, q(x)\, dx$$

- Differentiating twice in K Breeden & Litzenberger (1978)

$$q(K) = e^{r\tau} \frac{\partial^2 C}{\partial K^2}$$

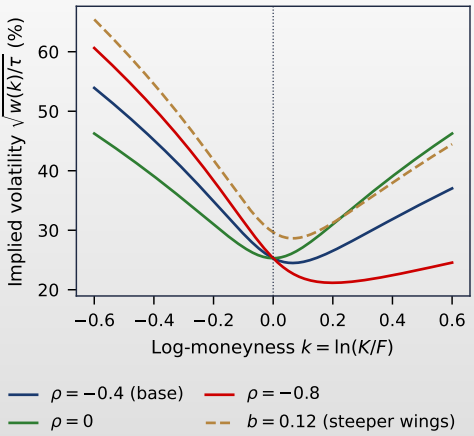
- **Butterfly** of step h: buy the calls $K - h$ and $K + h$, sell two calls K

$$C(K - h) - 2C(K) + C(K + h) \geq 0$$

  - divided by h^2, it approximates $e^{-r\tau} q(K)$; the plotted bars also multiply by $e^{r\tau}$
  - a negative butterfly is a negative density: an arbitrage or a quote error (A3, B8)

- Chart: Black–Scholes calls, $S = 100$, $\tau = 0.5$, $\sigma = 30\%$, $r = 2\%$, $h = 5$

The Smile and the SVI Curve



- If Black–Scholes were exact, σ_{imp} would not depend on K; in the data it does: the **smile** (or skew)
- **SVI** (Stochastic Volatility Inspired, parameterisation) Gatheral & Jacquier (2014): total implied variance $w = \sigma_{imp}^2 \tau$ as a function of $k = \ln(K/F)$

$$w(k) = a + b\Big(\rho(k-m) + \sqrt{(k-m)^2 + s^2}\Big)$$

  - a: level; $b \geq 0$: slope of the wings
  - $\rho \in (-1, 1)$: tilt; $\rho < 0$ makes the put wing dearer (skew)
  - m: horizontal shift; $s > 0$: curvature at the minimum
  - in the wings w is linear in k, with slopes $b(1 \pm \rho)$
- Chart: $a = 0.004$, $b = 0.08$, $m = 0$, $s = 0.15$, $\tau = 0.25$; one parameter changed at a time

Fitting a Smile and Reading a Density from It

- **Least squares** on n quotes (k_i, w_i), $\theta = (a, b, \rho, m, s)$

$$\hat{\theta} = \arg\min_{\theta} \sum_{i=1}^{n} \big(w(k_i;\theta) - w_i\big)^2$$

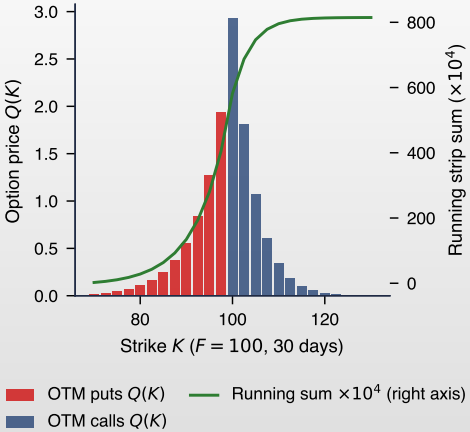
  - fit quality: root mean squared error of $\sqrt{w/\tau}$, in vol points
  - **pairs bootstrap**: resample the quotes with replacement, refit, read 2.5% and 97.5% percentiles Efron (1979)
- **Butterfly condition** Gatheral & Jacquier (2014): the density is nonnegative iff $g(k) \geq 0$ for all k

$$g(k) = \Big(1 - \frac{k\,w'}{2w}\Big)^2 - \frac{w'^2}{4}\Big(\frac{1}{w} + \frac{1}{4}\Big) + \frac{w''}{2}$$

  - w', w'': first and second derivatives of $w(k)$
- **Black-76** prices from the fitted w (forward F, discount $e^{-r\tau}$), then $q(K) = e^{r\tau}\partial^2 C/\partial K^2$ (B8)

$$C = e^{-r\tau}\big[F\Phi(d_1) - K\Phi(d_2)\big], \qquad d_{1,2} = \frac{-k \pm w/2}{\sqrt{w}}$$

Pricing Variance with Options: the VIX



- **Variance swap**: pays $RV - K_{var}$ at T; its fair strike is a strip of OTM options Bakshi, Kapadia & Madan (2003)

$$K_{var} = \frac{2e^{rT}}{T}\Big[\int_0^F \frac{P(K)}{K^2}dK + \int_F^\infty \frac{C(K)}{K^2}dK\Big]$$

- **Cboe formula** Cboe (2025) for a finite strip

$$\sigma^2 = \frac{2}{T}\sum_i \frac{\Delta K_i}{K_i^2} e^{rT} Q(K_i) - \frac{1}{T}\Big(\frac{F}{K_0} - 1\Big)^2$$

  - $Q(K_i)$: OTM mid price (puts below K_0, calls above); K_0: first strike at or below F; ΔK_i: half the gap between neighbouring strikes internally, the full adjacent gap at each endpoint; at K_0 use the mean put/call mid price
  - VIX $= 100\sqrt{\sigma^2}$ interpolated to 30 days; VIX^2: the risk-neutral expected variance, in $\%^2$ a year

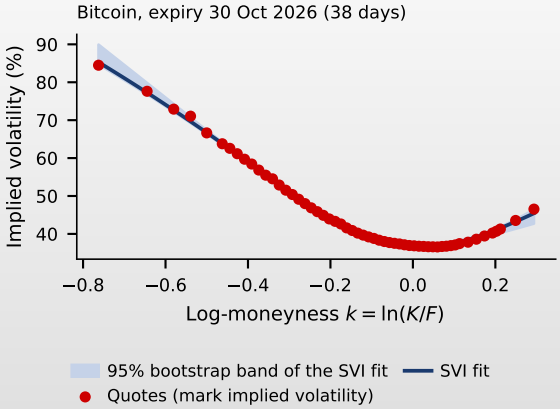
- VIX $= 20$: about 20% a year, $20/\sqrt{12} \approx 5.8\%$ a month

B4 [Solved]: An SVI Smile for Bitcoin — Task

- **Questions**
  - does a five-parameter SVI curve describe the Bitcoin smile?
  - how precisely are its at-the-money level and its skew estimated?
- **Inputs**
  - Option snapshot of the lecture; expiry closest to 30 days: 2026-10-30, 38 days, 56 quoted OTM options (puts below the forward, calls above); for each quote $k = \ln(K/F)$ and $w = \sigma_{mark}^2 \tau$, with σ_{mark} the mark IV
  - Retain quotes with positive bid and mark IV, at least two days to expiry, and $|k| < 1.2$; a fitted expiry needs at least seven quotes. Convert mark IV from percent to decimal before computing w.
  - the SVI curve $w(k)$ in "The Smile and the SVI Curve" Gatheral & Jacquier (2014); the butterfly function $g(k)$ in "Fitting a Smile and Reading a Density from It": $g \geq 0$ for all k means a nonnegative density
- **Tasks**
  1. Fit the five SVI parameters by least squares on w, with $b \geq 0$, $|\rho| < 1$, $s > 0$, and report the root mean squared error in vol points.
  2. Resample the quotes with replacement 300 times, refit SVI each time, and report 95% intervals for the ATM volatility and the skew $IV(0.15) - IV(-0.15)$.
  3. Evaluate $g(k)$ on $k \in [-1.5, 1.5]$ and report its minimum.
- **Report**: parameters, intervals, $\min g$ and the chart of quotes, fit and band; one sentence each: which parameter controls the skew and its sign; is the skew significantly different from 0; why fit w and not prices
- Notebook: section B4

Back Exercise list

B4 [Solved]: Solution (1/2) — Fit and Chart



- Step 1: $\hat{a} = -0.0204$, $\hat{b} = 0.1016$, $\hat{\rho} = -0.2615$, $\hat{m} = -0.0787$, $\hat{s} = 0.3482$; root mean squared error 0.42 vol points
- Step 2: ATM volatility 36.37%, 95% CI [36.10, 36.87]; skew -2.11, 95% CI $[-2.73, -1.52]$
- Step 3: $\min_k g(k) = 0.1359 > 0$ on $k \in [-1.5, 1.5]$: no butterfly arbitrage on this grid

B4 [Solved]: Solution (2/2) — Interpretation

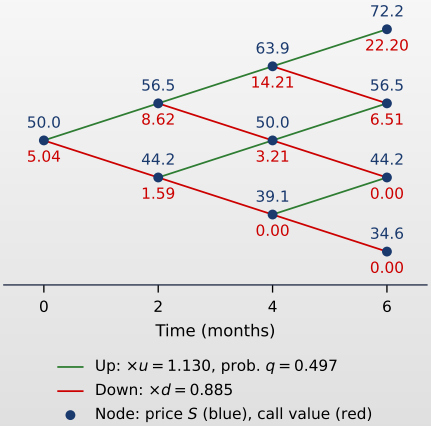
- ρ controls the tilt: $\hat{\rho} < 0$ makes the put wing dearer; under the pairs-bootstrap scheme, the skew interval excludes 0
- Why w: it is comparable across strikes and nearly linear in the wings, while option prices span four orders of magnitude and would weight the fit towards the dearest options
- A positive density everywhere also needs $g \geq 0$ outside the grid and the wing conditions of Gatheral & Jacquier (2014)
- **Interpretation**: the bootstrap resamples smoothed mark volatilities that are dependent across strikes: the intervals measure the sensitivity of the fit to the points selected, not market uncertainty
- For economic uncertainty: sensitivity to bid–ask, to the estimator and to the extrapolated wings

MFM_ch12_seminar

Checks and Project Seed

- Check your results
  - Dividends and the interest rate are in parity; vega and theta carry their units (per vol point, per day)
  - No naive standard errors on overlapping windows; no risk-neutral probability read as a forecast
  - A Sharpe ratio of a short-volatility strategy is reported with its skewness and worst month
- Team-project seed: repeat B4, B6 and B8 on a new Deribit snapshot of Ether options and compare with Bitcoin Alexander & Imeraj (2020), Pele et al. (2023); delta-hedge one-week Bitcoin options on historical paths Matic, Packham & Härdle (2023)

The Binomial Tree



- **Cox–Ross–Rubinstein** Cox, Ross & Rubinstein (1979): in each step of length $\Delta t = T/N$ the price moves to Su or Sd

$$u = e^{\sigma\sqrt{\Delta t}}, \quad d = \frac{1}{u}, \quad q = \frac{e^{r\Delta t} - d}{u - d}$$

  - q: risk-neutral probability of an up move; the expected price grows at the rate r
- **Backward induction**: start from the pay-offs at T, then at each node

$$V = e^{-r\Delta t}\left[qV_u + (1-q)V_d\right]$$

  - V_u, V_d: option values at the two next nodes
- Chart: $S = K = 50$, six months, $r = 4\%$, $\sigma = 30\%$, $N = 3$: $C_0 = 5.04$
- As $N \to \infty$, C_N tends to the Black–Scholes price (A6)

Tools for the Extension Packs

- **Delta method** (A4): standard error of a smooth function $g(\hat{\theta})$ of estimates with covariance $\hat{V}$

$$\operatorname{se}(g(\hat{\theta})) \approx \sqrt{\nabla g(\hat{\theta})^\top \, \hat{V} \, \nabla g(\hat{\theta})}$$

- **Instrumental variable** z_t (IV, not implied volatility) (B3) Christensen & Prabhala (1998): correlated with the regressor, not with the error; $\hat{\beta}_{IV} = \widehat{\operatorname{cov}}(z, y) / \widehat{\operatorname{cov}}(z, x)$
- **Delta-hedged option gain** Bakshi & Kapadia (2003) (B5): option value change minus the hedge gains over $[t, t+\tau]$

$$\Pi_{t,t+\tau} = C_{t+\tau} - C_t - \sum_{n=0}^{N-1} \Delta_{t_n} \left(S_{t_{n+1}} - S_{t_n} \right)$$

  - the formula assumes zero interest and no dividends; otherwise the financing of the cash position is added
  - mean zero needs continuous hedging in the correct model and realised volatility equal to implied along the path, not only in aggregate: the exposure is gamma-weighted
  - negative for the buyer if a variance premium exists
- **Martin bound** Martin (2017) (C3): a lower bound on the expected excess return of the market

$$E_t\left[R_{t,T}\right] - R_{f,t} \geq \frac{1}{R_{f,t}} \operatorname{var}_t^*\left(R_{t,T}\right) = \frac{2}{S_t^2} \int_0^\infty \operatorname{OTM}_t(K) \, dK$$

  - $R_{t,T}$: gross market return; $R_{f,t}$: gross risk-free return; var^*: risk-neutral variance; $\operatorname{OTM}_t(K)$: OTM option price at strike K
  - the inequality needs $\operatorname{cov}_t(M_{t,T} R_{t,T}, R_{t,T}) \leq 0$ ($M_{t,T}$: stochastic discount factor); it does not hold for every distribution

Extension Packs: What Each One Needs

Exercise	Question	Type	Model	Data
A3	are Bitcoin call prices convex in the strike?	Solved	–	option snapshot
A4	which forward and rate does parity imply?	Proposed	A3	option snapshot
A5	what does a dealer earn when $\sigma_{real} \neq \sigma_{imp}$?	Solved	–	–
A6	how fast does the binomial tree converge?	Proposed	A5	–
A7	can one option remove gamma and vega together?	Proposed	A1	–
A8	how large are the errors of a variance strip?	Proposed	A3	–
B5	did delta-hedged option sellers earn a premium?	Proposed	B1, B2	S&P 500, VIX
B6	what is a VIX for Bitcoin?	Proposed	A8, B4	option snapshot
B7	is there a variance risk premium in Bitcoin?	Proposed	B2	DVOL, Bitcoin
B8	how fat are the tails the market prices?	Proposed	B4	option snapshot
C2	is the variance premium a tradable signal?	Proposed	B2, B7	S&P 500, VIX, DVOL
C3	is the Martin bound valid and tight?	Proposed	B2	case-study files

A3 [Solved]: Bounds and Convexity in the Strike — Task

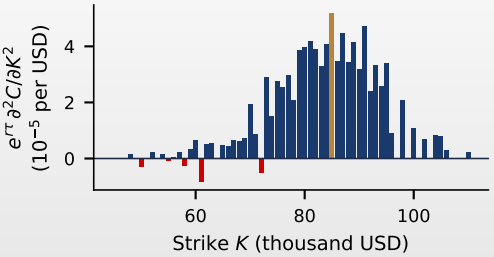
- **Questions**
  - which model-free restrictions must call prices satisfy across strikes?
  - do Bitcoin quotes satisfy them?
- **Inputs**
  - European calls on a non-dividend-paying asset, price $C(K)$, maturity τ, rate r; no arbitrage, no frictions
  - Bitcoin calls expiring 30 October 2026 in the option snapshot; USD price = BTC premium $\times$ index $S = 85\,265$; forward $F = 85\,742$, so $e^{r\tau} = F/S$; **mark** = the exchange's reference price, **bid/ask** = the best buy and sell quotes
- **Tasks**
  1. Prove $\max(S - Ke^{-r\tau}, 0) \le C \le S$ by comparing the pay-offs of two portfolios in every state at T.
  2. For $K_1 < K_2 < K_3$ with $K_2 = \lambda K_1 + (1-\lambda)K_3$, prove $\lambda C(K_1) + (1-\lambda)C(K_3) \ge C(K_2)$.
  3. Show that the butterfly $[C(K-h) - 2C(K) + C(K+h)]/h^2$ tends to $e^{-r\tau}q(K)$, with q the risk-neutral density Breeden & Litzenberger (1978).
  4. Compute the 84/85/86 thousand butterfly from mark prices, $q(85\,000)$ and $P^Q(84\,000 < S_T < 86\,000)$; then count the negative butterflies among all 732 equally spaced triples, at mark prices and at executable prices (wings at the ask, body at the bid).
- **Report**: the proofs, the butterfly, $q(85\,000)$ and the violation counts; is any violation an arbitrage you could trade?
- Notebook: section A3

Back Exercise list

A3 [Solved]: Solution (1/2) — Proofs

- Step 1: a call plus $Ke^{-r\tau}$ in cash is worth $\max(S_T, K) \geq S_T$ at T, so $C + Ke^{-r\tau} \geq S$; the call pays $(S_T - K)^+ \geq 0$, so $C \geq 0$
  - Upper bound: the share pays $S_T \geq (S_T - K)^+$ in every state, so $C \leq S$
- Step 2: for fixed S_T the map $K \mapsto (S_T - K)^+$ is convex, so $\lambda(S_T - K_1)^+ + (1-\lambda)(S_T - K_3)^+ \geq (S_T - K_2)^+$ in every state
  - A portfolio that pays at least as much in every state costs at least as much today
- Step 3: $C(K) = e^{-r\tau} \int_K^\infty (x - K) q(x)\, dx$, so $\partial_K C = -e^{-r\tau} \int_K^\infty q(x)\, dx$ and $\partial_K^2 C = e^{-r\tau} q(K)$
  - The butterfly quotient is the central second difference, so it converges to $e^{-r\tau} q(K)$: convexity is exactly $q \geq 0$

A3 [Solved]: Solution (2/2) — Bitcoin Numbers and Chart



- Step 4: $C(84\,000) = 4\,957.26$, $C(85\,000) = 4\,414.38$, $C(86\,000) = 3\,922.98$ USD; butterfly $= 51.47$ USD > 0
- $q(85\,000) \approx 1.0056 \cdot 51.47/1\,000^{\mathbf{2}} = 5.18 \cdot 10^{-\mathbf{5}}$ per USD; $P^{Q}(84\,000 < S_T < 86\,000) \approx 2h\,q(K) = 10.4\%$
- 7 of 732 mark butterflies are negative (worst -11.80 USD, deep in the money, median spread 255.80 USD); at executable prices: 0
- **Interpretation**: no tradable arbitrage, but a density read from raw marks is negative in places (red bars, 5 of 54 adjacent triples): smooth the smile before differentiating

MFM_ch12_seminar

A4 [Proposed]: Parity as a Regression: the Implied Forward — Task

Solution discussed in class

- **Question**: which forward price and which USD interest rate do Bitcoin option prices imply through put–call parity?
- **Inputs**: the expiry of A3; USD mid prices $(bid + ask)/2 \times S$; only strikes where both the call and the put have a bid and an ask; $D = e^{-r\tau}$ and F unknown
- **Tasks**
  1. Show from put–call parity that $C(K) - P(K) = DF - DK$ for every strike.
  2. Write the Ordinary Least Squares (OLS) regression $C_i - P_i = a + bK_i + e_i$ and derive $\hat{D}$ and $\hat{F}$ from $(\hat{a}, \hat{b})$.
  3. Derive the delta-method standard error of $\hat{F}$ with a heteroskedasticity-robust covariance.
  4. Estimate on the data and compare $\hat{F}$ with the exchange forward, in standard errors.
- **Report**: $\hat{D}$, $\hat{F}$ with standard errors and the number of strikes
  - which USD rate does $\hat{D}$ imply?
  - why might $\hat{F}$ differ from the exchange forward?
- **Model**: A3 (portfolios with identical pay-offs)
- Notebook: section A4

Back Exercise list

A5 [Solved]: The Black–Scholes Equation and the Hedging P&L — Task

- **Question**: what does a dealer who sells an option at σ_{imp} and delta-hedges it earn when the price moves with σ_{real}?
- **Inputs**
  - $dS_t = \mu S_t\,dt + \sigma S_t\,dW_t$, constant r, continuous trading without costs; option value $V(S,t)$ twice differentiable Black & Scholes (1973)
  - Numbers: the call of A1, sold at $\sigma_{imp} = 25\%$, gamma $\Gamma = 0.0222$ (from A1); 252 trading days a year
- **Tasks**
  1. Replicate V with a self-financing portfolio of Δ shares and cash B at r; apply Itô's lemma, choose Δ to remove the dW_t term, and derive the Black–Scholes equation from no arbitrage.
  2. For a dealer who sells at σ_{imp} and holds $\Delta = \partial_S V^{imp}$ while S moves with σ_{real}, derive the instantaneous P&L.
  3. Compute today's P&L rate per trading day if $\sigma_{real} = 30\%$, then annualise it with S and Γ frozen.
- **Report**: the equation, the P&L formula and the two numbers
  - who gains?
  - why does the P&L depend on the path of S?
- Notebook: section A5

Back Exercise list

A5 [Solved]: Solution (1/2) — The Equation

- Step 1: Itô: $dV = \left(V_t + \mu S V_S + \frac{1}{2}\sigma^2 S^2 V_{SS}\right) dt + \sigma S V_S \, dW_t$; the self-financing $\Pi = \Delta S + B$, started at $\Pi = V$, has $d\Pi = \Delta \, dS + rB \, dt$
- $d(V - \Pi) = \left(V_t + \frac{1}{2}\sigma^2 S^2 V_{SS} - rB\right) dt + (V_S - \Delta)(\mu S \, dt + \sigma S \, dW_t)$; with $\Delta = V_S$ both the risk and μ disappear
- A riskless position that costs 0 earns 0: $V_t + \frac{1}{2}\sigma^2 S^2 V_{SS} + rSV_S - rV = 0$
- Step 2: the book $-V + \Delta S + (V - \Delta S)$ in cash changes by $-dV + \Delta \, dS + r(V - \Delta S) \, dt$ plus interest on the P&L; Itô with σ_{real} and the equation at σ_{imp} give
- $d\,\text{P\&L} = r\,\text{P\&L}\, dt + \frac{1}{2}\Gamma S^2 \left(\sigma_{imp}^2 - \sigma_{real}^2\right) dt$

A5 [Solved]: Solution (2/2) — Numbers

- Step 3: $\frac{1}{2} \cdot 0.02223 \cdot 100^2 \cdot (0.25^2 - 0.30^2)/252 = \frac{1}{2} \cdot 0.02223 \cdot 10^4 \cdot (-0.0275)/252 = -0.0121$ per option and trading day
- Annualised with S and Γ frozen: -3.06, about 39% of the call price 7.76: a rate, not the realised P&L of this six-month option, since ΓS^2 changes along the path
- **Interpretation**: the seller loses when realised exceeds implied volatility; the loss accrues through Γ, largest near the strike and near expiry, so two paths with the same realised variance give different P&L

A6 [Proposed]: The Binomial Tree Converges to Black–Scholes — Task

Solution discussed in class

- **Questions**
  - how fast does the Cox–Ross–Rubinstein price approach the Black–Scholes price?
  - why does the error oscillate?
- **Inputs**: Cox–Ross–Rubinstein (CRR) tree with N steps of length $\Delta t = T/N$ Cox, Ross & Rubinstein (1979); European call of the lecture: $S = K = 100$, $T = 1$, $r = 5\%$, $\sigma = 20\%$
- **Tasks**
  1. Show that $u = e^{\sigma\sqrt{\Delta t}}$, $d = 1/u$, $q = (e^{r\Delta t} - d)/(u - d)$ match the risk-neutral mean exactly and the variance of the log step up to $O(\Delta t^2)$.
  2. Price the call with $N = 50, 51, 100, 101, 200, 201$ and compute $N(C_N - C_{BS})$ for each.
  3. Average the prices for $N = 200$ and $N = 201$ and compare the error with that of $N = 200$.
- **Report**: a table of C_N and $N(C_N - C_{BS})$; the order of the error, the reason for the even/odd oscillation, and why averaging helps
- **Model**: A5 (from the tree to the continuous-time equation)
- Notebook: section A6

Back Exercise list

A7 [Proposed]: A Delta-Gamma Hedge — Task

Solution discussed in class

- **Questions**
  - how many options and shares make a short option book delta- and gamma-neutral?
  - is it then vega-neutral too?
- **Inputs**
  - A dealer has **sold** 1 000 calls of A1 ($S = K = 100$, six months, $r = 3\%$, $\sigma = 25\%$)
  - Hedging instruments: a call with $K = 110$, same expiry and flat volatility $\sigma = 25\%$, Black–Scholes price 3.90, and the stock; Greeks from the Black–Scholes formulas, gamma per unit of S
- **Tasks**
  1. Compute the delta and the gamma of the dealer's position.
  2. Find the number n_2 of $K = 110$ calls that makes the gamma zero, then the number of shares that makes the delta zero.
  3. Compute the vega of the hedged book and explain the result; retain fractional quantities until the final displayed table.
- **Report**: a positions table (instrument, quantity, delta, gamma, vega); is the book protected against a change in volatility?
- **Model**: A1
- Notebook: section A7

Back Exercise list

A8 [Proposed]: The Variance-Swap Strike and Its Estimation Error — Task

Solution discussed in class

- **Question**: how large are the errors of a finite strip of options when it is used to price variance, in a case where the answer is known?
- **Inputs**
  - Spanning formula of the lecture: $f(S_T) = f(F) + f'(F)(S_T - F) + \int_0^F f''(K)(K - S_T)^+ dK + \int_F^\infty f''(K)(S_T - K)^+ dK$ Bakshi, Kapadia & Madan (2003)
  - Log-normal benchmark: $F = 100$, $r = 0$, 90 days ($T = 90/365$), flat volatility 20%, so the exact variance is 0.04; OTM prices from Black–Scholes; **discretisation** error = Cboe sum minus the dense integral on $[80, 120]$; **truncation** error = dense integral on $[80, 120]$ minus 0.04
- **Tasks**
  1. Apply the formula with $f(S) = -2\ln(S/F)$ and show $K_{var} = \frac{2e^{rT}}{T}\left[\int_0^F \frac{P(K)}{K^2} dK + \int_F^\infty \frac{C(K)}{K^2} dK\right]$ for continuous paths.
  2. Compute the Cboe sum with $\Delta K = 5$ on $[80, 120]$ Cboe (2025), the dense integral on $[80, 120]$ and the dense integral on $(0, \infty)$.
  3. Split the total error into discretisation and truncation.
- **Report**: the three variances and volatilities and the two errors
  - which error dominates here?
  - what changes for a longer maturity?
- **Model**: A3 (portfolios of options as prices of pay-off shapes)
- Notebook: section A8

Back Exercise list

B5: Cash Flows of the Monthly Hedged Call

- Trade calendar
  - first trade: the first common S&P 500/VIX date
  - then sell a call every 21 observed trading days
  - strike: the index on the sale date; maturity: 21/252 years
- Pricing and hedging
  - fixed volatility for pricing and for daily delta hedging: the VIX at the opening of the trade
  - interest and dividends are zero in this simplified experiment
  - rebalance after each observed daily price, up to the day before expiry
- At expiry
  - liquidate the stock hedge and pay the call payoff
  - an expired option is not rebalanced
- ATM
  - in this exercise: strike equal to spot
  - in the option-strip exercises: strike near the forward; the benchmark and the discounting convention are stated together

B5 [Proposed]: Delta-Hedged S&P 500 Options, 1990–2026 — Task

Solution discussed in class

- **Questions**
  - did sellers of delta-hedged one-month S&P 500 calls earn a significant premium?
  - at what tail risk?
- **Inputs**
  - Daily S&P 500 and VIX closes, January 1990 – September 2026; interest and dividends ignored
  - Cash-flow convention: B5 cash flows. Strategy of the lecture: every 21 trading days sell a synthetic one-month ATM call, priced at the VIX, and hedge it daily at the same volatility; P&L in % of the index Bakshi & Kapadia (2003)
- **Tasks**
  1. Compute the mean monthly P&L and its Newey–West standard error with 3 lags, and test H_0: mean $= 0$ at 5%.
  2. Build a 95% moving-block bootstrap interval for the mean (blocks of 6 months, 2 000 resamples) Künsch (1989).
  3. Compute the share of positive months, the skewness, and the share of the total gain lost in the five worst months.
- **Report**: the mean with both intervals and the tail statistics; why can a strategy with a high t-statistic still be dangerous?
- **Model**: B1 (bootstrap), B2 (Newey–West)
- Notebook: section B5

Back Exercise list

B6 [Proposed]: A VIX for Bitcoin — Task

Solution discussed in class

- **Questions**
  - what does the Cboe VIX formula give for Bitcoin?
  - why does it differ from the at-the-money volatilities and from DVOL?
- **Inputs**: the option snapshot; the two expiries around 30 days (2026-10-09, 2026-10-30); USD mid Q = BTC mid $\times$ index S; F and $e^{RT} = F/S$ from the Deribit forward
- **Tasks**
  1. Select the strip of each expiry with the Cboe rules Cboe (2025): K_0 the first strike at or below F; puts below K_0, calls above, the average of both at K_0; drop zero bids and stop each wing after two consecutive zero bids.
  2. Compute each implied variance σ^2 with the Cboe formula in "Pricing Variance with Options: the VIX" (R, the rate implied by $e^{RT} = F/S$, in place of r).
  3. Interpolate the total variances linearly in time to 30 days and report $100\sigma_{30}$.
  4. Compare with the ATM volatilities of the two expiries and with DVOL at the same moment Alexander & Imeraj (2020).
- **Report**: the two variances, the weight of the near expiry and the index
  - why is the index above both ATM volatilities?
  - which choices could explain the gap to DVOL?
- **Model**: A8 (strip formula and its errors), B4 (the same snapshot)
- Notebook: section B6

Back Exercise list

B7: Matching the Bitcoin Dates

- Bitcoin prices
  - the daily close and the date labels of the supplied BTC-USD.CC.csv
  - do not substitute an intraday snapshot
- DVOL
  - the close field of the Deribit daily candle (resolution 1D)
  - convert its millisecond timestamp to a UTC calendar date
  - join on that date without shifting either series
- Realised variance at date t uses the next 30 daily log returns, excluding the return ending at t
- Limit of the alignment
  - the BTC file records dates, not intraday timestamps
  - the comparison is aligned by date; it does not claim perfectly synchronous prices

B7 [Proposed]: The Variance Risk Premium of Bitcoin — Task

Solution discussed in class

- **Questions**
  - is the Bitcoin variance risk premium significantly positive?
  - is it larger than the equity premium?
- **Inputs**: DVOL (Deribit), March 2021 – August 2026, and daily Bitcoin closes (`BTC-USD.CC`); realised volatility over the next 30 calendar days, $A = 365$
- **Tasks**
  1. Compute the mean premium $\mathrm{DVOL}_t - \sqrt{RV_{t,t+\mathbf{30}}}$ in vol points and its Newey–West standard error with 30 lags, and test it at 5%.
  2. Compute the share of positive days and the ratio of mean implied to mean realised volatility, and compare both with the S&P 500 of B2.
  3. Count the non-overlapping 30-day periods in the sample.
- **Report**: premium, test, ratios and effective sample size
  - in which unit is the Bitcoin premium larger?
  - how much can five years of data tell us?
- **Model**: B2
- Notebook: section B7

Back Exercise list

B8 [Proposed]: A Risk-Neutral Density with Confidence Bands — Task

Solution discussed in class

- **Questions**
  - how much more probability than the log-normal density does the Bitcoin market put on large moves?
  - how precise is that estimate?
- **Inputs**: the option snapshot; expiry closest to 90 days: 2026-12-25, 94 days, 58 OTM quotes; SVI fitted as in B4; forward (undiscounted) prices
- **Tasks**
  1. Compute the density $q(K) = \partial^2 C/\partial K^2$ from the SVI call prices (Black-76 with F) on a fine strike grid Breeden & Litzenberger (1978).
  2. Compute $P^Q(S_T < 0.8F)$ and $P^Q(S_T > 1.2F)$ from the distribution function $F_Q(K) = 1 + \partial C/\partial K$, with no truncation or renormalisation. Give 95% pairs-bootstrap intervals (300 refits; reject refits with $g(k) < 0$) and the log-normal values at the ATM volatility (mean F, log variance $\sigma_{ATM}^2 \tau$).
  3. In each refit, recompute the log-normal value at that refit's ATM volatility and give a 95% interval for the difference SVI minus log-normal.
  4. Report where the quotes end in K/F and which part of each probability relies on extrapolated wings.
- **Report**: the two probabilities with intervals vs log-normal
  - is the difference significant?
  - what does the bootstrap band not measure?
- **Model**: B4 (SVI fit and pairs bootstrap)
- Notebook: section B8

Back Exercise list

C2: A Monthly Variance Strategy with a Fixed Calendar

- Monthly grid
  - align IV and daily log returns on common observed dates; the grid starts at the first aligned return
  - S&P 500: every 21 trading days; the next $H = 21$ returns give the realised variance
  - Bitcoin: every 30 calendar days; the next $H = 30$ returns give the realised variance
- Units
  - annualise the mean squared daily log return by 252 (S&P 500) or 365 (Bitcoin), then multiply by 10^4 for percent-squared units
  - IV is in annualised volatility percent, so the P&L of the formula is in volatility points
- Sharpe interval
  - 2 000 moving-block resamples of three monthly observations, seed 42, annualisation $\sqrt{12}$
  - report the sample endpoints and the count of complete periods
- VIX/VIX3M rule
  - S&P 500 only, after VIX3M (Cboe 3-Month Volatility Index) begins in 2007
  - skipped periods stay on the same calendar with zero P&L
  - compare with the unfiltered strategy over exactly those dates

C2 [Proposed]: Is the Variance Risk Premium a Tradable Signal? — Task

Solution discussed in class

- **Questions**
  - would you sell variance on the S&P 500 or on Bitcoin every month (Carr & Wu (2009), Bollerslev, Tauchen & Zhou (2009))?
  - when would you stop?
- **Inputs**
  - A synthetic variance swap with vega notional 1: monthly P&L $= (IV^2 - RV)/(2\,IV)$ vol points
  - S&P 500 with the VIX (1990–2026, VIX3M from 2007); Bitcoin with DVOL (2021–2026); C2 calendar convention
- **Tasks**
  1. Compute the monthly P&L series and the Sharpe ratio with a block-bootstrap interval, the skewness, the worst month and the maximum drawdown.
  2. Test a timing rule: skip months that start in backwardation (VIX above VIX3M); is it better?
  3. Check how fragile the result is to the day on which the monthly grid starts.
  4. Compare with the delta-hedged option of B5 and propose a position size.
- **Report**: a table and a chart of cumulative P&L; a verdict with its limitations
- **Model**: B2, B7; the predictive regressions of the lecture appendix

Back Exercise list

C3: The Nine Conditioning Variables

- Source: the monthly Welch–Goyal fields and the transformations in the data construction script
- Three transformed variables: $5 + \ln(D12/Index)$, $5 + \ln(E12/Index)$, b/m
- The remaining six: tbl, $1 + lty - tbl$, $BAA - AAA$, $svar$, $1 + ntis$, $1 + infl$; keep the units of the supplied fields
- Definitions
  - $D12$, $E12$: trailing annual dividends and earnings
  - tbl: Treasury-bill rate; lty: long-term yield; BAA, AAA: corporate bond yields
  - $svar$: stock variance; $ntis$: net equity issuance; $infl$: inflation
  - b/m: book equity over market equity; $Index$: the stock-index level
- The added constants keep the variables positive without changing their monthly timing

C3: Validity and Tightness Statistics

- Σ estimates the covariance of the *mean* vector $\bar{\lambda}$
  - not the unscaled long-run covariance
  - do not multiply the statistics by the sample size again
- Projection onto nonnegative moments: $\hat{\lambda} = \arg\min_{\lambda \geq 0} (\bar{\lambda} - \lambda)^\top \Sigma^{-1} (\bar{\lambda} - \lambda)$
- Statistics
  - $D_1 = (\bar{\lambda} - \hat{\lambda})^\top \Sigma^{-1} (\bar{\lambda} - \hat{\lambda})$: violation of the inequalities
  - $D_2 = \hat{\lambda}^\top \Sigma^{-1} \hat{\lambda}$: distance from zero within the permitted region
- P-values: `KPClass.py`, `p_values`, for the chi-square mixture
  - weights are simulated with 100 000 draws; record the seed
  - the returned p-values are percentages
- Two-moment illustration: $\Sigma = I_2$ and $\bar{\lambda} = (-1, 2)$
  - the projection is $(0, 2)$, hence $D_1 = 1$ and $D_2 = 4$
  - the first moment violates its bound; the second has positive slack

C3 [Proposed]: Testing the Martin Bound — Task

Solution discussed in class

- **Question**: is the lower bound on the equity premium of Martin (2017) valid and tight on the course data, as in the case study Back, Crotty & Kazempour (2022)?
- **Inputs**
  - The authors' daily Martin bound, annualised %, $h = 1, 3, 6, 12$ months, 1990–2020 (columns `lb_m_h` of `market_bounds.csv`, replication package)
  - S&P 500 (`GSPC.INDX`); daily risk-free rate (Kenneth French Data Library); monthly Welch–Goyal predictors Welch & Goyal (2008) (Amit Goyal's data page)
- **Tasks**
  1. Compute the daily excess return (price return minus the risk-free rate), compound it over 21, 63, 126 and 252 trading days and annualise with $12/h$, for formation dates 31 January 1990 – December 2020 ($N = 7789$ at 1 month, 7598 at 12 months).
  2. Build the nine positive conditioning variables of Table 3 of the case study (carried forward from month-ends) and form the 10 moments $s_{t,T}(1, z_t)$ with $s_{t,T} = R^e_{t,T} - b_{t,T}$ and their mean $\bar{\lambda}$.
  3. Estimate Σ by Hansen–Hodrick Hansen & Hodrick (1980) with $21h$ lags (Newey–West with $1.5 \times 21h$ lags if not positive semi-definite) and compute $D_{\mathbf{1}}$, $D_{\mathbf{2}}$ and their p-values Kodde & Palm (1986).
  4. Regress $R^e_{t,T}$ on a constant and $b_{t,T}$ (Hansen–Hodrick errors) and test $\alpha = 0$, $\beta = 1$ jointly; repeat step 3 with the bound replaced by its time-series mean (Section 4.3 of the study).
- **Report**: a table of means, $D_{\mathbf{1}}$, $D_{\mathbf{2}}$ and p-values next to Tables 1 and 3 of the study, and one sentence on why the regression test and the tightness test can disagree
- **Model**: B2 (overlapping windows and HAC errors)

Back Exercise list

Acronyms used in this material

- **ACF**: AutoCorrelation Function
- **AI**: Artificial Intelligence
- **ATM**: At-The-Money (strike near the stated spot or forward reference)
- **CI**: Confidence Interval
- **CRR**: Cox–Ross–Rubinstein (binomial tree)
- **DM**: Diebold–Mariano (test)
- **DVOL**: Deribit Volatility Index
- **GBM**: Geometric Brownian Motion
- **GMM**: Generalized Method of Moments
- **HAC**: Heteroskedasticity and Autocorrelation Consistent (standard errors)
- **HAR-RV**: Heterogeneous AutoRegressive model of Realised Volatility
- **IV**: Implied Volatility
- **MSE**: Mean Squared Error
- **MZ**: Mincer–Zarnowitz (regression)
- **OLS**: Ordinary Least Squares
- **OTM**: Out-of-The-Money (call strike above, put strike below the stated spot or forward reference)
- **QLIKE**: Quasi-LIKElihood (loss function)
- **RV**: Realised Variance
- **SVI**: Stochastic Volatility Inspired (parameterisation)
- **US**: United States
- **UTC**: Coordinated Universal Time (Temps Universel Coordonné)
- **VIX**: Cboe Volatility Index
- **VIX3M**: Cboe 3-Month Volatility Index
- **VRP**: Variance Risk Premium

References (1/3)

- Alexander, C., Imeraj, A. (2020). The Bitcoin VIX and its variance risk premium. *Journal of Alternative Investments*, 23(4), 84–109.
- Back, K., Crotty, K., Kazempour, S.M. (2022). Validity, tightness, and forecasting power of risk premium bounds. *Journal of Financial Economics*, 144(3), 732–760.
- Bakshi, G., Kapadia, N. (2003). Delta-hedged gains and the negative market volatility risk premium. *Review of Financial Studies*, 16(2), 527–566.
- Bakshi, G., Kapadia, N., Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies*, 16(1), 101–143.
- Black, F., Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637–654.
- Bollerslev, T., Tauchen, G., Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies*, 22(11), 4463–4492.
- Breeden, D.T., Litzenberger, R.H. (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4), 621–651.
- Brenner, M., Subrahmanyam, M.G. (1988). A simple formula to compute the implied standard deviation. *Financial Analysts Journal*, 44(5), 80–83.
- Carr, P., Wu, L. (2009). Variance risk premiums. *Review of Financial Studies*, 22(3), 1311–1341.
- Cboe Global Indices (2025). Cboe Volatility Index: Methodology. Cboe Global Markets.
- Christensen, B.J., Prabhala, N.R. (1998). The relation between implied and realized volatility. *Journal of Financial Economics*, 50(2), 125–150.

References (2/3)

- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2), 174–196.
- Cox, J.C., Ross, S.A., Rubinstein, M. (1979). Option pricing: a simplified approach. *Journal of Financial Economics*, 7(3), 229–263.
- Diebold, F.X., Mariano, R.S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263.
- Efron, B. (1979). Bootstrap methods: another look at the jackknife. *Annals of Statistics*, 7(1), 1–26.
- Gatheral, J., Jacquier, A. (2014). Arbitrage-free SVI volatility surfaces. *Quantitative Finance*, 14(1), 59–71.
- Hansen, L.P., Hodrick, R.J. (1980). Forward exchange rates as optimal predictors of future spot rates: an econometric analysis. *Journal of Political Economy*, 88(5), 829–853.
- Kodde, D.A., Palm, F.C. (1986). Wald criteria for jointly testing equality and inequality restrictions. *Econometrica*, 54(5), 1243–1248.
- Künsch, H.R. (1989). The jackknife and the bootstrap for general stationary observations. *Annals of Statistics*, 17(3), 1217–1241.

References (3/3)

- Martin, I. (2017). What is the expected return on the market?. *Quarterly Journal of Economics*, 132(1), 367–433.
- Matic, J.L., Packham, N., Härdle, W.K. (2023). Hedging cryptocurrency options. *Review of Derivatives Research*, 26(1), 91–133.
- Mincer, J., Zarnowitz, V. (1969). The evaluation of economic forecasts. In *Economic Forecasts and Expectations*, NBER, 3–46.
- Newey, W.K., West, K.D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.
- Patton, A.J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1), 246–256.
- Pele, D.T., Wesselhöfft, N., Härdle, W.K., Kolossiatis, M., Yatracos, Y.G. (2023). Are cryptos becoming alternative assets?. *European Journal of Finance*, 29(10), 1064–1105.
- Welch, I., Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4), 1455–1508.