








Modelling Financial Markets

Chapter 12: Options and the Volatility Surface



Daniel Traian PELE

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

What do option prices reveal about the future distribution?



Fischer Black. Photo: Dalmatine at English Wikipedia (2009 upload); public domain; Wikimedia Commons

"If options are correctly priced in the market, it should not be possible to make sure profits by creating portfolios of long and short positions in options and their underlying stocks."

— *Black & Scholes (1973)*

- Black–Scholes: replication gives a single implied volatility per asset
  - daily delta hedging, three-month call: error standard deviation 9.5% of the premium
- The smile: one-month implied volatility from the historical distribution of the S&P 500 (Standard & Poor's)
  - 30.6% at $K/S = 0.8$, against 13.9% at the money
- Bitcoin, 38-day expiry: risk-neutral probability of a fall of more than 30%
  - 1.24%, against 0.14% under the log-normal benchmark
- S&P 500, 1990–2026: mean VIX (CBOE Volatility Index) 19.4, mean realised volatility 15.3
  - variance risk premium: 4.09 vol points, positive on 86% of days

▸ Central chart: implied vs realised volatility

The Idea of the Chapter

- An option pays only in some states: its price is a bet on the **distribution**
  - not only on the direction
  - so it reveals what the market expects about volatility and crashes
- Part I, pricing by replication
  - parity, static spanning, the binomial tree
  - the Black–Scholes equation, the Greeks and the discrete-hedging error
- Then implied volatility, the smile and a full Bitcoin surface with its risk-neutral density
- Part II, implied volatility as information
  - variance swaps and the VIX (the Cboe index of 30-day S&P 500 implied volatility)
  - the variance risk premium and 0DTE options (zero days to expiry)
- Case studies: option-implied bounds on the equity premium and how the VIX reacts to volatility shocks
- AI (artificial intelligence) for scientific discovery
- **What you should be able to do**
  - derive put–call parity, the binomial price and the Black–Scholes equation
  - compute the Greeks, the discrete-hedging error and an arbitrage-free SVI (Stochastic Volatility Inspired, parameterisation) surface with its density and bands
  - interpret option-implied forecasts (VIX, variance risk premium) with the right standard errors
  - Further Topics and Appendix: reading for familiarity

Chicago, 1973: Listed Options Begin



Trading floor of the Chicago Board of Trade in 2006; its members founded the Cboe in 1973

Photo: Richie Diesterheft (2006); CC BY 2.0; Wikimedia Commons

- The Chicago Board Options Exchange (Cboe) opened in 1973 MacKenzie & Millo (2003)
  - the first exchange for standardised, listed options
  - Before: options were negotiated one by one, over the counter
- The same year: the Black–Scholes and Merton pricing papers Black & Scholes (1973), Merton (1973)
- Traders soon used the formula on the floor: theory changed the market it described MacKenzie & Millo (2003)

The Formula and Its Authors



Fischer Black (1938–1995)

Photo: Dalmatine at English Wikipedia; public domain; Wikimedia Commons



Myron Scholes (2008)

Photo: Nobellaureatesphotographer (2008); CC BY-SA 3.0; Wikimedia Commons



Robert C. Merton (2006)

Photo: Digarnick (2006); public domain; Wikimedia Commons

- Key idea: a portfolio of the stock and cash can **replicate** the option Black & Scholes (1973)
  - so its price follows from no-arbitrage
- Scholes and Merton received the 1997 Nobel Memorial Prize in Economic Sciences
  - Black had died in 1995

Put–Call Parity: What Does No Arbitrage Alone Imply?

- Two portfolios with the same pay-off at T (European options, same K and T, no dividends):
  - A: one call + cash $Ke^{-r\tau}$ (worth K at T) $\Rightarrow \max(S_T, K)$
  - B: one put + one share $\Rightarrow \max(K, S_T)$
- Same pay-off in every state, so the same price today, otherwise there is an arbitrage:
  - $C_t - P_t = S_t - Ke^{-r\tau}$ (with dividend yield δ: $S_t e^{-\delta\tau}$ instead of S_t)
- No model for S is needed: parity uses only no-arbitrage Merton (1973)
- Consequence: a call and a put with the same K and T have the **same** implied volatility

Static Replication: Any Pay-Off from a Bond, a Forward and Options

- Carr–Madan spanning: any twice-differentiable pay-off $f(S_T)$ and any point F:
  - Taylor with integral remainder; proof in the Appendix
  - $f(S_T) = f(F) + f'(F)(S_T - F) + \int_0^F f''(K)\,(K - S_T)^+ dK + \int_F^\infty f''(K)\,(S_T - K)^+ dK$
  - Reading: a bond, a forward and a **strip** of options: $f''(K)\,dK$ puts below F and calls above F
  - Priced today (F the forward): $V_0 = e^{-r\tau} f(F) + \int_0^F f''(K)\,P(K)\,dK + \int_F^\infty f''(K)\,C(K)\,dK$
  - $P(K)$, $C(K)$: today's put and call prices at strike K, whatever the model for S Bakshi, Kapadia & Madan (2003)
- Why this price: each piece has a market price today
  - the bond $e^{-r\tau}$ per unit, the forward 0, the options $P(K)$ and $C(K)$
  - No model for S and no pricing measure are needed for V_0; the measure Q of Section 3 is needed only to read option prices as expectations, as in the density bullet of the next slide

Static Replication: Three Results Used Later

- "Risk-neutral": a price is the discounted mean pay-off under the pricing measure Q
  - not under real-world probabilities
  - Q is constructed in Section 3
- **Variance**: $f(S) = -2\ln(S/F)$ has $f''(K) = 2/K^2$
  - the strip with weights $1/K^2$ is the variance swap and the VIX (Section 8)
  - Low strikes get the largest weights, so an expensive put wing makes variance expensive
- **Density**: a butterfly $[C(K-h) - 2C(K) + C(K+h)]/h^2 \to \partial^2 C/\partial K^2 = e^{-r\tau} q(K)$ Breeden & Litzenberger (1978)
  - h: the strike step; the risk-neutral density is a price (Section 7)
- Straddles, spreads and protective puts are finite strips
  - f'' concentrated at one or two strikes
  - moments and SVIX (simple VIX, the variance index of Martin (2017)): ▸ Appendix

The Binomial Tree: How Do We Price an Option by Replication?

- One period of length Δt: $S \to uS$ or dS, with $d < R = e^{r\Delta t} < u$
  - u, d: up and down factors
  - R: gross riskless return
- Replicate the call with Δ shares and cash B: $\Delta uS + RB = C_u$ and $\Delta dS + RB = C_d$
  - C_u, C_d: the call values after an up and a down move
  - Solution: $\Delta = \dfrac{C_u - C_d}{(u-d)S}$, and the price is $C = \Delta S + B = \dfrac{q\,C_u + (1-q)\,C_d}{R}$
  - $q = \dfrac{R-d}{u-d} \in (0,1)$ (the risk-neutral probability, the discrete analogue of the density $q(K)$) is the unique probability under which S_t/R^t is a martingale
- The real probability of an up move never enters
  - risk preferences are already in the price S
- **Fundamental theorem**: no arbitrage $\Leftrightarrow$ an equivalent martingale measure Q exists Harrison & Pliska (1981), Delbaen & Schachermayer (1994)
  - Q is unique exactly when the market is complete; with jumps or stochastic volatility there are many Q, and option prices carry risk premia (Section 8)

The CRR Tree: N Steps and Convergence to Black–Scholes

- Why N steps: one period has only two final states
  - with N steps the $N+1$ final states approximate the distribution of S_T
- The CRR (Cox–Ross–Rubinstein) tree Cox, Ross & Rubinstein (1979): $\Delta t = T/N$, $u = e^{\sigma\sqrt{\Delta t}}$, $d = 1/u$
  - This choice gives the one-step log return the variance $\sigma^2\Delta t$ (up to terms of order Δt^2), so the tree converges to the geometric Brownian motion of the Black–Scholes model
  - $q = (e^{r\Delta t} - d)/(u - d)$ is the same at every node
- European call: $C_0 = e^{-rT} \sum_{j=0}^{N} \binom{N}{j} q^j (1-q)^{N-j} \left(S_0 u^j d^{N-j} - K\right)^+$
  - j: the number of up moves; the sum is the discounted risk-neutral mean of the pay-off
- American options: value backwards from T
  - taking max(continuation value, exercise value) at every node

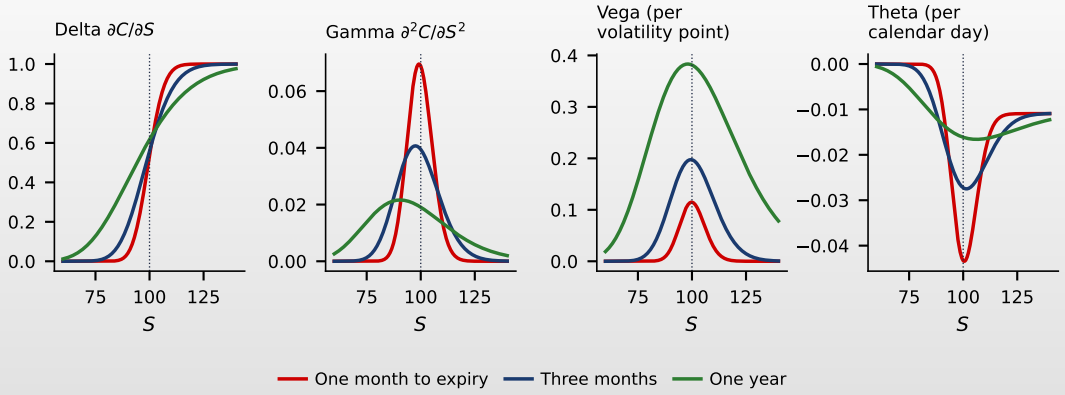
The Black–Scholes Equation: What Price Does Continuous Replication Give?

- Geometric Brownian motion $dS_t = \mu S_t\,dt + \sigma S_t\,dW_t$
  - constant σ and r
  - continuous trading, no costs, short selling allowed
- Replicate the option value $V(S,t)$ with a **self-financing** portfolio Π
  - Δ_t shares and cash B_t at r: $d\Pi = \Delta\,dS + rB\,dt$
- Itô's lemma: $dV = \left(V_t + \mu S V_S + \tfrac{1}{2}\sigma^2 S^2 V_{SS}\right)dt + \sigma S V_S\,dW_t$
  - V_t, V_S, V_{SS}: the partial derivatives
- Choose $\Delta = V_S$, start from $\Pi = V$ ($B = V - SV_S$)
  - $d(V-\Pi) = \left(V_t + \tfrac{1}{2}\sigma^2 S^2 V_{SS} - rB\right)dt$
  - both dW_t and μ drop out
- A riskless position that costs nothing must earn nothing
  - hence the **Black–Scholes equation** Black & Scholes (1973), Merton (1973): $V_t + \tfrac{1}{2}\sigma^2 S^2 V_{SS} + rSV_S - rV = 0$
  - with $V(S,T) =$ pay-off
- That μ is irrelevant is a result, not an assumption
  - Feynman–Kac: $V = e^{-r\tau}E^Q[\text{pay-off}]$ with S drifting at r under Q

The Black–Scholes Formula

- European call and put, no dividends (with yield δ: replace S by $Se^{-\delta\tau}$):
  - $C = S\,\Phi(d_1) - Ke^{-r\tau}\,\Phi(d_2), \qquad P = Ke^{-r\tau}\,\Phi(-d_2) - S\,\Phi(-d_1)$
  - $d_1 = \dfrac{\ln(S/K) + (r + \sigma^2/2)\tau}{\sigma\sqrt{\tau}}, \qquad d_2 = d_1 - \sigma\sqrt{\tau}$, with Φ the standard Normal distribution function
- Reading the two terms:
  - $\Phi(d_2)$ = risk-neutral probability that the call ends in the money
  - $\Phi(d_1)$ = number of shares in the replicating portfolio (the delta)
- Five inputs: S, K, τ, r observed
  - σ is the only one not observed: it must be estimated or implied

The Greeks: One Identity, All Sensitivities (1/2)



MFM_ch12_greeks ▶ Interactive

The Greeks: One Identity, All Sensitivities (2/2)

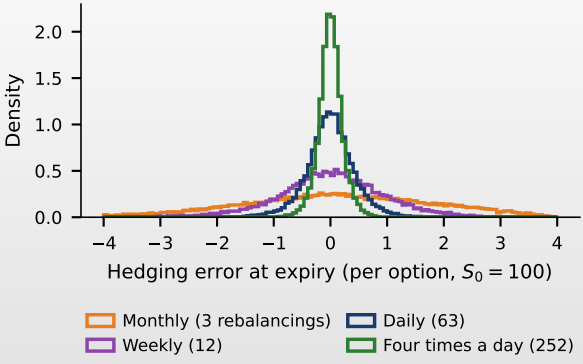
- Greeks: partial derivatives of the price
  - $\Delta = \partial C/\partial S$, $\Gamma = \partial^2 C/\partial S^2$
  - vega $\mathcal{V} = \partial C/\partial \sigma$, $\Theta = \partial C/\partial t$, rho $= \partial C/\partial r$
- Delta $\Delta = \Phi(d_1)$ (put: $\Phi(d_1) - 1$)
  - gamma $\Gamma = \varphi(d_1)/(S\sigma\sqrt{\tau})$, φ the standard Normal density
  - key identity $S\varphi(d_1) = Ke^{-r\tau}\varphi(d_2)$
- Vega $= S\varphi(d_1)\sqrt{\tau} = \Gamma S^2\sigma\tau$: for one expiry gamma and vega are proportional
  - Seminar 12, A1, A7
  - theta $\Theta = rC - rS\Delta - \frac{1}{2}\sigma^2 S^2\Gamma$
  - rho $= K\tau e^{-r\tau}\Phi(d_2)$
- Call with $K = 100$
  - short maturities concentrate gamma and theta at the strike
  - long maturities carry the most vega
  - vega per vol point, theta per calendar day

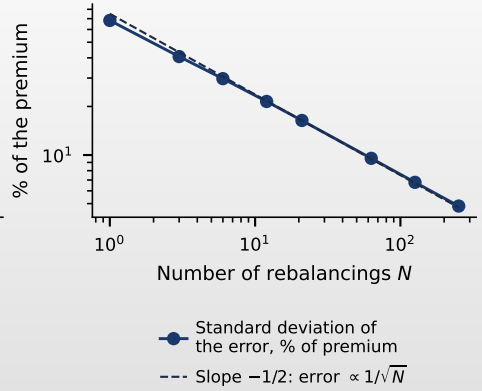
MFM_ch12_greeks

Discrete Delta Hedging: How Much Risk Is Left with N Rebalancings?

- Black–Scholes replication needs continuous rebalancing
  - in practice the hedge is adjusted N times, at $t_i = iT/N$
- At $t_0 = 0$: sell the call at its Black–Scholes price C_0
  - buy $\Delta_0 = \Phi(d_1)$ shares
  - keep the cash $B_0 = C_0 - \Delta_0 S_0$
- At t_i, $i = 1, \dots, N-1$: the portfolio is self-financing
  - $\Delta_i = \Phi\big(d_1(S_{t_i}, T - t_i)\big)$
  - $B_i = B_{i-1} e^{rT/N} - (\Delta_i - \Delta_{i-1}) S_{t_i}$
- **Hedging error** at expiry: $\varepsilon_N = \Delta_{N-1} S_T + B_{N-1} e^{rT/N} - (S_T - K)^+$
- Simulation: geometric Brownian motion with $\mu = 8\%$ and $\sigma = 20\%$
  - $S_0 = K = 100$, $r = 4\%$, $T = 0.25$
  - 20 000 paths; $N = 3, 12, 63, 252$
- Why $\mu \neq r$: replication never uses μ
  - a real drift different from r shows that the error stays centred at 0

Discrete Delta Hedging: Algorithm and Error (1/2)





MFM_ch12_hedge_error ▶ Interactive

Discrete Delta Hedging: Algorithm and Error (2/2)

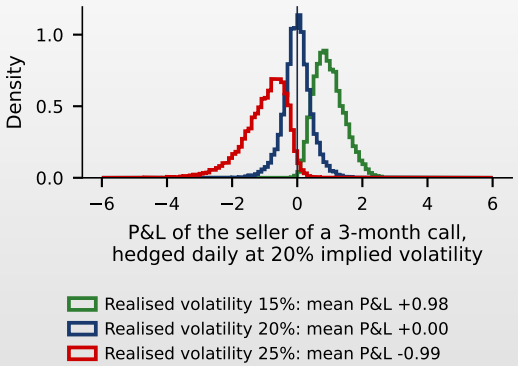
- Left: distribution of the error ε_N for four rebalancing frequencies
  - right: standard deviation of the error, in % of the premium, against N on log axes
- Black–Scholes premium 4.49 and $d_1 = 0.15$
- Mean error near 0
  - standard deviation (s.d.) in % of the premium: monthly 40.8%, weekly 21.5%, daily 9.5%, four times a day 4.8%
  - worst 1% of paths, in currency units per option: -2.68 (weekly), -1.17 (daily)
- Log-log slope -0.48
  - theory (Appendix): s.d. $\approx \sqrt{\pi/4}\,\mathcal{V}\sigma/\sqrt{N}$ ($\mathcal{V}$: vega)
  - theory 1.01, 0.44, 0.22 vs simulated 0.96, 0.43, 0.22 ($N = 12, 63, 252$)
- Transaction costs depend on the cost model
  - with a fixed cost per rebalancing, the total cost grows like N
  - with costs proportional to the traded value, in the standard diffusion setting, the expected cost grows roughly like $\sqrt{N}$

MFM_ch12_hedge_error

Derivation: the Gamma-Weighted Volatility P&L

- Sell the option at σ_{imp} and hold $\Delta = \partial_S V^{imp}$
  - V^{imp}: its Black–Scholes value at σ_{imp}
  - the price moves with σ_{real}: $dS = \mu S\,dt + \sigma_{real} S\,dW_t$
- Itô on V^{imp}: $dV = V_t\,dt + \Delta\,dS + \frac{1}{2}\Gamma\sigma_{real}^2 S^2 dt$
  - the equation at σ_{imp} gives $V_t = rV - rS\Delta - \frac{1}{2}\sigma_{imp}^2 S^2\Gamma$
- Short option + Δ shares + cash at r, self-financing; its profit and loss (P&L):
  - $d\,\text{P\&L} = r\,\text{P\&L}\,dt + \frac{1}{2}\,\Gamma_t S_t^2\left(\sigma_{imp}^2 - \sigma_{real}^2\right)dt$, so over the life
  - $\text{P\&L} = \displaystyle\int_0^T e^{r(T-t)}\,\frac{1}{2}\,\Gamma_t S_t^2\left(\sigma_{imp}^2 - \sigma_{real,t}^2\right)dt$: path-dependent, weighted by Γ, large near K and near T
- A short variance swap pays $\sigma_{imp}^2 - \sigma_{real}^2$ with constant weights
  - a delta-hedged option is a variance bet with random weights
- First-order size: vega $\times$ (implied $-$ realised), checked on the next slide
  - Why: vega $= \Gamma S^2\sigma\tau$ (Section 4), so $\frac{1}{2}\Gamma S^2(\sigma_{imp}^2 - \sigma_{real}^2)\tau = \frac{1}{2}\Gamma S^2(\sigma_{imp} + \sigma_{real})(\sigma_{imp} - \sigma_{real})\tau \approx \Gamma S^2\sigma(\sigma_{imp} - \sigma_{real})\tau = \text{vega} \times (\sigma_{imp} - \sigma_{real})$

Selling Volatility at 20%: Three Outcomes



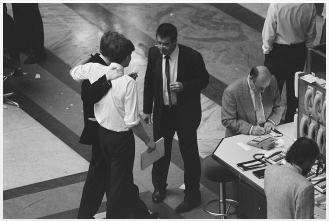
- Three-month call sold at 20% implied volatility, hedged daily
  - realised volatility 15%, 20% or 25%
- Realised 15%: mean P&L $+0.98$, positive on 99.8% of paths
- Realised 20%: mean $+0.00$; only discrete-hedging noise remains (s.d. 0.43)
- Realised 25%: mean -0.99, positive on only 2.2% of paths
- Each mean is close to vega $\times$ 5 vol points $= 0.197 \times 5 \approx \pm 0.99$
  - a clean bet on volatility

MFM_ch12_vol_mismatch ▶ Interactive

Implied Volatility: Which σ Explains the Market Price?

- **Implied volatility** σ_{imp}: the σ that makes the Black–Scholes price equal the market price
  - $BS(S, K, \tau, r, \sigma_{imp}) = C^{market}$; $BS(\cdot)$: the Black–Scholes formula, C^{market}: the observed price
  - The solution is unique because vega > 0: the price increases in σ
  - There is no closed form: σ enters both $\Phi(d_1)$ and $\Phi(d_2)$, so the equation is solved numerically
- Traders quote options in implied volatility, not in money
  - it is comparable across strikes, maturities and assets
- If Black–Scholes were right, all options on one asset would have the **same** implied volatility

After October 1987: the Smile Appears



Amsterdam stock exchange after the crash, 21 October 1987

Photo: Bart Molendijk / Anefo (1987); CC0; Wikimedia Commons

- Before 1987 index options had a nearly flat implied volatility across strikes Rubinstein (1994)
  - Since the crash, low-strike puts are much more expensive: the **skew**
- October 1986 – August 1987: option prices showed crash fears Bates (1991)
  - Not strongly in the two months immediately before the crash
- **Smile**: implied volatility as a function of K
  - **term structure**: as a function of T
  - **surface**: both
- Crash risk on the Bucharest Stock Exchange: Pele & Mazurencu-Marinescu (2012)

Why Smiles Exist and How Models Produce Them (1/2)

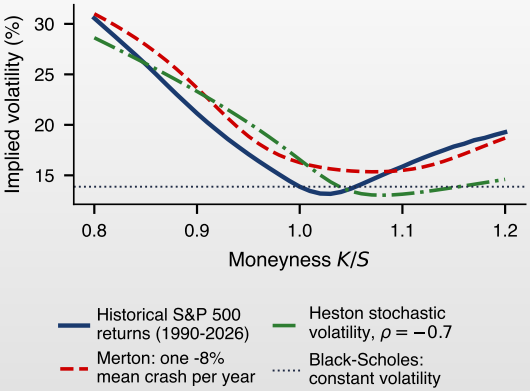
- **Fat tails**: large moves are more frequent than the Normal distribution allows (Chapter 1)
  - so deep OTM (Out-of-The-Money) options are worth more than Black–Scholes says
- **Leverage and volatility feedback**: when prices fall, volatility rises Christie (1982)
  - negative skewness makes low-strike puts dearer
- **Crash fear**: the market prices a negative jump risk, visible in put prices Bates (1991)
- **Demand for insurance**: investors buy puts Gârleanu, Pedersen & Poteshman (2009)
  - intermediaries, who cannot fully hedge the risk, charge a premium for them

Why Smiles Exist and How Models Produce Them (2/2)

Model	Idea	Fits best
Jumps Merton (1976)	Poisson jumps added to the Brownian motion	steep short-dated smiles
Stochastic volatility Heston (1993), Hull & White (1987)	variance follows its own mean-reverting process, correlated with S	skew and term structure of longer maturities
Local volatility Gatheral (2006)	$\sigma(S,t)$ read off the whole surface (the formula of Dupire (1994))	matches every quoted price today, exotic options
Parametric smile, SVI (stochastic volatility inspired) Gatheral & Jacquier (2014)	five parameters per expiry, no dynamics	interpolation and quoting (Section 7)

- Processes: Chapter 11
  - the implied-volatility surface in detail: Franke, Härdle & Hafner (2019), Čížek, Härdle & Weron (2011)

Three Sources of a Smile (1/2)



MFM_ch12_smile_models ▶ Interactive

Three Sources of a Smile (2/2)

- One-month implied volatility from three sources
  - the historical S&P 500 distribution, a jump model and a stochastic-volatility model
- History: 9 225 overlapping monthly returns
  - skewness -1.27, excess kurtosis 6.8
  - the excess kurtosis requires a finite fourth moment, which Chapter 1 questions
  - rescaled to a martingale ($r = 0$, no risk adjustment)
  - 30.6% at $K/S = 0.8$, 13.9% ATM (At-The-Money), 19.3% at $K/S = 1.2$
- Merton jumps (one crash of about -8% a year): 31.0% at $K/S = 0.8$ vs 16.3% ATM
- Heston (price–variance correlation $\rho = -0.7$)
  - 28.6% at $K/S = 0.8$, 16.6% ATM, 14.6% at $K/S = 1.2$
- Risk-neutral laws are more left-skewed than realised ones Jackwerth (2000)
  - priced crash risk steepens market smiles further

MFM_ch12_smile_models

The Bitcoin Surface: How to Describe the Smile Without Arbitrage?

- Deribit lists European options on Bitcoin, traded 24/7 Deribit DVOL
  - The options are **inverse**: premium and pay-off in bitcoin (BTC), not in dollars (USD)
  - Each expiry has its own forward F
  - DVOL: Deribit's 30-day implied-volatility index
- Our snapshot: 22 September 2026, 07:51 UTC (Coordinated Universal Time)
  - 9 expiries that have at least two days left
  - 375 OTM options with a positive bid
- Why Bitcoin: the only asset in the course data with a full option chain
  - across strikes and expiries
- Crypto risk: Pele et al. (2023)
- Hedging Bitcoin options: Matic, Packham & Härdle (2023)

The SVI Smile: Total Variance and Parameters

- Work with the **total implied variance** $w(k) = \sigma_{imp}^2(k)\,\tau$ at log-moneyness $k = \ln(K/F)$
- SVI parameterisation Gatheral & Jacquier (2014): $w(k) = a + b\left(\rho\,(k-m) + \sqrt{(k-m)^2 + s^2}\right)$
  - SVI: stochastic volatility inspired
  - a: level
  - b: slope of the wings
  - ρ: tilt (negative = put skew)
  - m: shift
  - s: curvature at the bottom

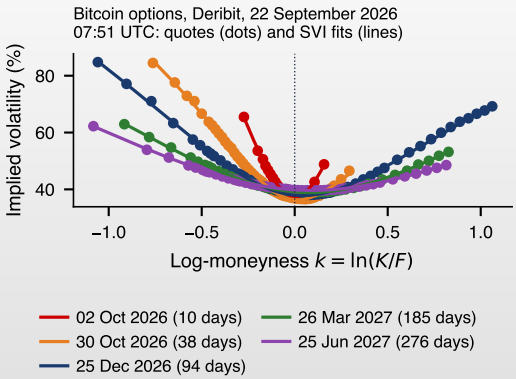
The SVI Smile: Why This Shape and How We Fit It

- Why w and k: the no-arbitrage conditions (next slide) are written directly in $w(k)$
  - different expiries become comparable
- Why SVI: as $|k| \to \infty$, w becomes linear in k
  - the shape that Lee (2004) allows for the wings
  - five parameters are estimated from 24–58 quotes per expiry
- Wing slopes: $b(1-\rho)$ on the left and $b(1+\rho)$ on the right
- Fit, for each expiry: $\min \sum_j \left(w(k_j) - \sigma^2_{imp,j}\tau\right)^2$
  - least squares on w
  - with parameter bounds and six starting points
- Each fit is then checked against the no-arbitrage constraints

No-Arbitrage Constraints on the Smile

- **Wings**: Lee (2004) bounds the asymptotic slope, $\limsup_{|k|\to\infty} w(k)/|k| \le 2$
  - for SVI $b(1+|\rho|) \le 2$ is necessary
  - $w \ge 0$ needs $a + bs\sqrt{1-\rho^2} \ge 0$
  - A steeper wing means fatter tails of S_T under Q (moment formula: Appendix)
- **Butterfly**: call prices convex in K, density $q \ge 0$ Gatheral & Jacquier (2014)
  - with $w' = \partial_k w$, $w'' = \partial_k^2 w$
  - $g(k) \ge 0$ is exactly a nonnegative density
  - $g(k) = \left(1 - \frac{kw'}{2w}\right)^2 - \frac{w'^2}{4}\left(\frac{1}{w} + \frac{1}{4}\right) + \frac{w''}{2} \ge 0$ for all k, plus $d_+(k) = -k/\sqrt{w} + \sqrt{w}/2 \to -\infty$ as $k \to \infty$ (Gatheral & Jacquier (2014), Lemma 2.2)
- **Calendar**: $\partial_\tau w(k,\tau) \ge 0$ for every k: total variance cannot fall with maturity
- Slice-by-slice least squares imposes none of them
  - SSVI (Surface SVI), the version of SVI for the whole surface, has simple sufficient conditions (Gatheral & Jacquier (2014), Theorems 4.1–4.2)

The Bitcoin Smile



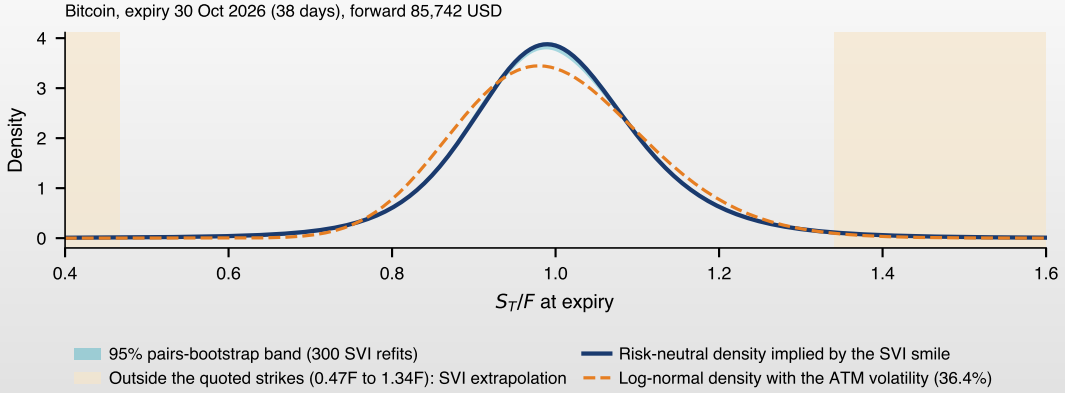
- Dots: OTM options at the exchange's reference implied volatility (mark IV); lines: SVI fits, five expiries
- A true **smile**: both wings rise, unlike the one-sided equity skew
  - Bitcoin can jump up as well as down Zulfiqar & Gulzar (2021)
- Short expiries have the steepest wings
  - a jump over a few days weighs more relative to $\sigma\sqrt{\tau}$
- SVI root mean squared error: at most 0.70 vol points
- Puts dearer than calls at equal $|k|$: skew is negative at all 9 expiries
- Type of evidence: description (it shows what happened in the data, without testing a cause)

MFM_ch12_btc_smile ▶ Interactive

The Risk-Neutral Density: From Prices to a Density

- From Breeden & Litzenberger (1978) (Section 2): $q(K) = e^{r\tau}\,\partial^2 C(K)/\partial K^2$
  - $q \geq 0$ is the butterfly condition above
- From noisy, discrete quotes this is an ill-posed inverse problem
  - differentiating twice amplifies the noise
  - Fit a smooth smile (SVI), then differentiate (other estimators: Appendix)
- Uncertainty: pairs bootstrap of the quotes, refit, differentiate
  - pointwise bands for q and intervals for tail probabilities
- Beyond the last quoted strike the tail is an extrapolation
  - report where the quotes end
- q mixes probabilities and risk aversion Aït-Sahalia & Lo (1998), Grith, Härdle & Park (2013)
  - pricing kernel $\propto q/p$, with p the physical density
  - Its non-monotone shape is the pricing-kernel puzzle Jackwerth (2000)

What the Market Prices for Bitcoin (1/2)



MFM_ch12_btc_rnd ▶ Interactive

What the Market Prices for Bitcoin (2/2)

- Density from the SVI smile of the expiry closest to 30 days
  - 38 days, forward 85 742 USD, ATM 36.4%
- Band: 95% pairs bootstrap; benchmark: a log-normal density
  - shaded: outside the 56 quotes ($0.47F$ to $1.34F$)
- Probabilities: $F_Q(K) = 1 + e^{r\tau}\,\partial C/\partial K$, analytic in SVI (mass 1, mean F)
  - bootstrap refits with butterfly arbitrage ($g(k) < 0$) are rejected: 5 of 305
- Fall of more than 30%: 1.24% [1.16, 1.27] vs 0.14% log-normal
  - about 9 times more
  - 0.16 of the 1.24 points lie below the lowest quoted strike
- Fall of more than 20%: 4.35% vs 3.27%
  - rise above 20%: 5.3% in both
- Narrow band: it shows fitting noise, not bid–ask uncertainty
  - $\pm 2.0\%$ at the centre, $\pm 7\%$ at $0.7F$
  - mark volatilities are already smoothed
- Prices of insurance, not forecasts: they include the premium for crash risk
- Type of evidence: description

MFM_ch12_btc_rnd

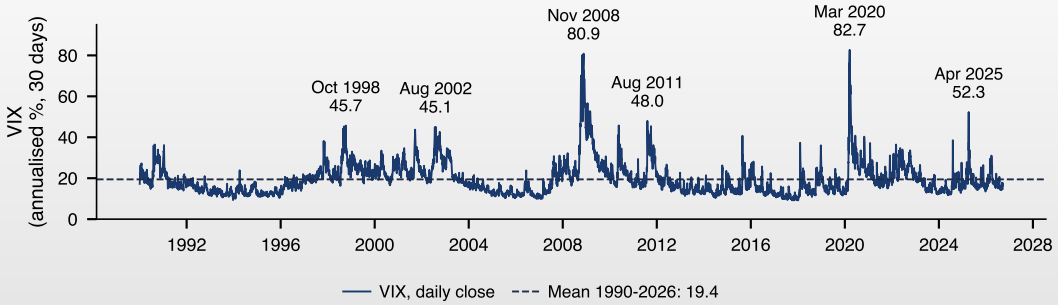
Variance Swaps: What Does Future Variance Cost?

- A **variance swap** pays at T: notional $\times$ (realised variance $-$ strike K_{var})
  - Realised variance $\frac{252}{n}\sum r_t^2$, r_t: the n daily log returns to T
  - Fair strike: the risk-neutral expected variance Demeterfi, Derman, Kamal & Zou (1999)
  - $K_{var} = \frac{2e^{r\tau}}{\tau}\left(\int_0^F \frac{P(K)}{K^2}\,dK + \int_F^\infty \frac{C(K)}{K^2}\,dK\right)$ (spanning with $f = -2\ln(S/F)$, Section 2)
- Why it works: Itô on $\ln S$ (continuous paths, $r = \delta = 0$) Britten-Jones & Neuberger (2000)
  - $-2\ln(S_T/F) = \int_0^T \sigma_t^2\,dt - 2\int_0^T dS_t/S_t$
  - the last term is a Q-martingale
  - "Model-free" means free of a model for σ_t, not of continuity: with crash jumps the strip understates expected realised variance Carr & Wu (2009) (▸ Appendix)
- Buyer of variance = insurance against turbulence Carr & Wu (2009)
  - seller = earns a premium, pays in crashes

The VIX Index: the Cboe Formula

- The discrete version of the strip K_{var} from the previous slide Cboe (2025):
  - $\sigma^2 = \frac{2}{T}\sum_i \frac{\Delta K_i}{K_i^2} e^{RT} Q(K_i) - \frac{1}{T}\left(\frac{F}{K_0} - 1\right)^2$
  - T: time to expiry in years
  - R: risk-free rate
  - K_i: the strikes
  - K_0: the first strike at or below the forward index level F
  - $Q(K_i)$ (Cboe notation, not the measure Q): bid–ask mid of the put for $K_i < K_0$, of the call for $K_i > K_0$, the average of both at K_0
  - The second term corrects for the strip being split at K_0 rather than exactly at F

The VIX, 1990–2026



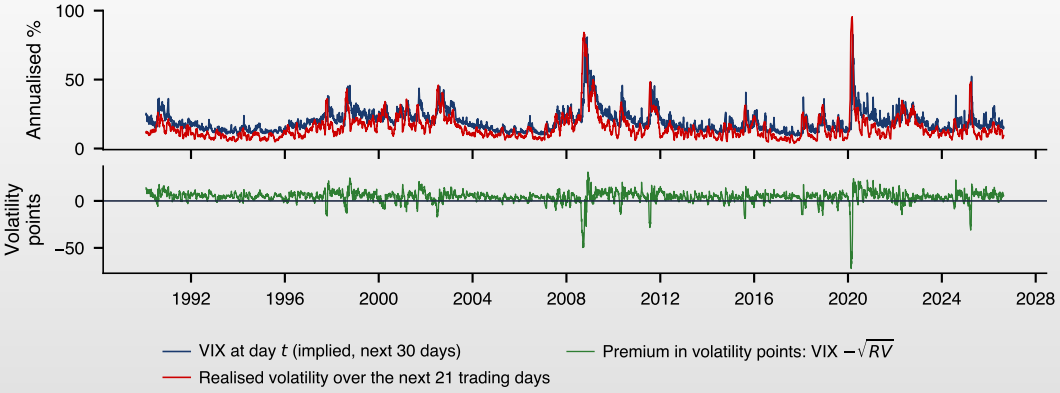
- Daily close; the six highest peaks at least 500 days apart are labelled
- Mean 19.4 (5.6% a month), range 9.1–82.7 (2020-03-16, COVID-19)
  - above 30 on 7.9% of days
  - sudden spikes, slow decay, as in GARCH (Chapter 5)

MFM_ch12_vix_history ▶ Interactive

The Variance Risk Premium

- Variance risk premium (VRP) Carr & Wu (2009), Bollerslev, Tauchen & Zhou (2009)
  - three versions that must not be confused:
  - **Economic (ex ante)**: $E_t^Q[RV_{t,t+\mathbf{21}}] - E_t^P[RV_{t,t+\mathbf{21}}]$; $VIX_t^{\mathbf{2}}$ estimates the first term, the second needs a forecast
  - **Ex post** (used here): $VIX_t^{\mathbf{2}} - RV_{t,t+\mathbf{21}}$, the realised P&L of the seller of a variance swap, known only at $t+21$
  - **Ex-ante proxy**: $VIX_t^{\mathbf{2}} - RV_{t-\mathbf{21},t}$, past variance as the forecast of E_t^P; known at t, so it can serve as a predictor (Section 11)
- Realised variance (RV): $RV_{t,t+\mathbf{21}} = \frac{\mathbf{252}}{\mathbf{21}} \sum_{i=\mathbf{1}}^{\mathbf{21}} r_{t+i}^{\mathbf{2}}$ (annualised, in $\%^{\mathbf{2}}$)
  - r_{t+i}: the daily log returns in %
  - roughly the same month as the 30 calendar days of the VIX
- Positive on average if investors dislike states with high volatility
  - these are also bad states for wealth
- Overlapping 21-day windows are autocorrelated: Newey–West with at least 21 lags Newey & West (1987)
  - 21 lags cover the overlap, not the persistence of volatility
  - standard error of the premium, in vol points: 0.27 with 21 lags, 0.31 with 42, 0.32 with 63, 0.33 with 126

Implied vs Realised Volatility, S&P 500 (1/2)



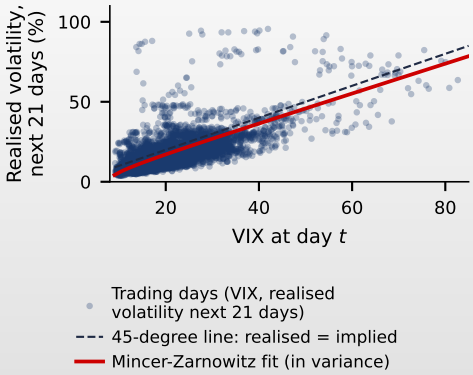
MFM_ch12_vrp ▶ Interactive

Implied vs Realised Volatility, S&P 500 (2/2)

- Top: VIX and the volatility realised over the next 21 trading days
  - bottom: their difference VIX − RV, 1990–2026
  - here RV $= \sqrt{RV_{t,t+\mathbf{21}}}$, the realised volatility, in vol points
- 9 203 days: mean VIX 19.4, mean realised 15.3
  - Premium 4.09 vol points (Newey–West standard error, SE, 0.27); in variance, 113 $\%^{\mathbf{2}}$, $t = 5.5$
  - the t in variance requires a finite fourth moment, which Chapter 1 questions
- Positive on 86% of days
  - the exceptions are crashes
  - Worst −72.0 vol points, window from 2020-02-20 (COVID-19)
- Same message as the delta-hedged option
  - sellers of volatility are paid on average and lose in crises
- Type of evidence: description
  - an ex-post premium; the economic premium needs a forecast of E_t^P (slide *The Variance Risk Premium*)

MFM_ch12_vrp

Is the VIX a Good Forecast?



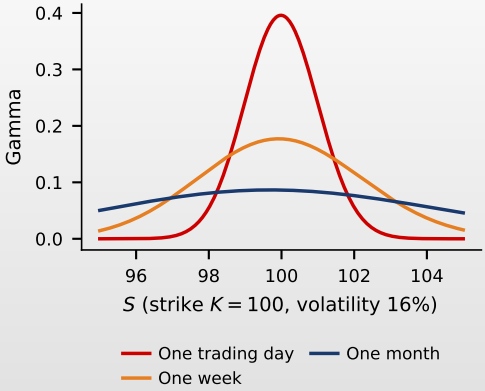
- Mincer–Zarnowitz Mincer & Zarnowitz (1969): $RV_{t,t+21} = \alpha + \beta\, VIX_t^2 + \varepsilon_t$
  - Unbiased: $\alpha = 0$, $\beta = 1$; (s.e.): Newey–West, 21 lags
  - inference on coefficients needs $E\|x_t\varepsilon_t\|^2 < \infty$ (x_t: the regressors) and weak dependence; in variance units, eighth moments of returns, which Chapter 1 questions
- $\hat{\alpha} = -51.9$ (34.4), $\hat{\beta} = 0.860$ (0.105), $R^2 = 0.37$
  - Separately neither differs from $(0, 1)$: $t = -1.51$ and -1.34
- But the estimates correlate -0.90: unbiasedness rejected
  - joint Wald $\chi^2_2 = 41.7$, $p = 8.9 \cdot 10^{-10}$
- Chart axes in volatility (%)
  - the fit estimated in variance drawn back in volatility
  - most points lie below the 45-degree line
- Past realised variance alone: $R^2 = 0.29$
  - With both: VIX keeps its weight (0.77), past variance adds little (0.08, SE 0.13)

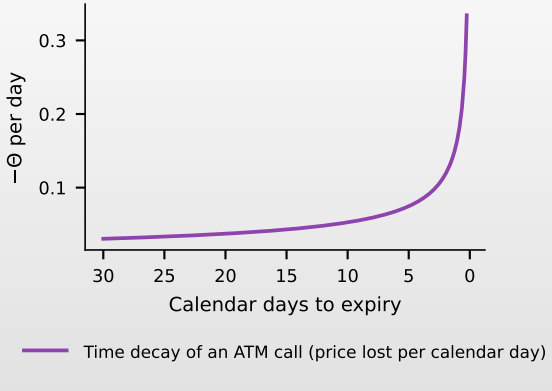
MFM_ch12_vix_rv_mz ▶ Interactive

0DTE Options: What Changes When Expiry Is Today?

- **0DTE** options are traded on the day they expire Cboe (2023)
  - S&P 500 options now expire every trading day
- More than 75% of retail trades in S&P 500 options involve 0DTE contracts Beckmeyer, Branger & Gayda (2023)
  - retail investors lose money on them
- 0DTE options were 48% of index-option trading in August 2023 Brogaard, Han & Won (2023)
- More 0DTE trading is followed by higher index volatility Brogaard, Han & Won (2023)
  - mostly through speculative retail trades
- Ultra-short pricing needs skewness and kurtosis corrections Bandi, Fusari & Renò (2023)
- From Section 4: as $\tau \to 0$, gamma and theta explode at the money

Gamma and Time Decay Near Expiry (1/2)





MFM_ch12_0dte_gamma ▶ Interactive

Gamma and Time Decay Near Expiry (2/2)

- Left: gamma of a call with $K = 100$ one day, one week and one month before expiry
  - right: daily time decay of an ATM call
- Inputs: $S = K = 100$, $\sigma = 16\%$, $r = 0$
  - gamma counts trading days ($\tau = 1/252$ and $21/252$)
  - time decay counts calendar days ($\tau = 1/365$ and $30/365$), as on the two axes
- ATM gamma one day before expiry 0.396 vs 0.086 one month before
  - 4.6 times larger ($= \sqrt{21}$)
  - a 1% move in the last day changes delta by about 0.4
- Time decay 0.167 in the last day vs 0.030 with 30 days left
  - about 5.5 times more
  - the seller collects fast decay and carries large gamma risk
  - The ATM value with one day left, about $0.4\, S\sigma\sqrt{\tau} = 0.33$, is lost entirely during that day: twice the decay rate at the start of the day (0.167)

MFM_ch12_0dte_gamma

Case Study: the Research Question (1/2)

- **Questions** Back, Crotty & Kazempour (2022): option prices give lower bounds on the equity premium
  - Are these bounds valid conditionally?
  - Are they tight conditionally?
  - Do they forecast returns out of sample?
- Why: a tight bound is a real-time equity-premium forecast with no estimated parameter
  - regression of realised returns on the bound
  - Martin (2017) could not reject $\alpha = 0$, $\beta = 1$ on 1996–2012
  - a test with low power

Case Study: the Research Question (2/2)

- **Martin bound** (eq. 3):
  - $R_{t,T}$ the gross market return, $R_{f,t,T}$ the gross riskless return
  - var^*_t the risk-neutral variance ($\mathrm{SVIX}^2_{t,T}$ not annualised here)
  - $E_t[R_{t,T}] - R_{f,t,T} \geq R_{f,t,T}\,\mathrm{SVIX}^2_{t,T}, \qquad \mathrm{SVIX}^2_{t,T} = \mathrm{var}^*_t(R_{t,T}/R_{f,t,T})$
  - Valid if $\mathrm{cov}_t(M_{t,T}R_{t,T}, R_{t,T}) \leq 0$ ($M_{t,T}$ the stochastic discount factor); tight if this covariance is zero, e.g. a log-utility investor who holds the market
- **Chabi-Yo–Loudis bound** Chabi-Yo & Loudis (2020): built from the first four risk-neutral moments
  - free of preference parameters in its restricted form
- **Link to this chapter**: SVIX is the static-spanning strip with $f = S^2$
  - Section 2 and ▸ Appendix
  - Realised slackness $s_{t,T} = R^e_{t,T} - b_{t,T}$ (excess return minus bound): valid means $E_t[s_{t,T}] \geq 0$, tight means $E_t[s_{t,T}] = 0$

Case Study: Data, Moments and Forecasting

- **Data** (Section 3 and Appendix C of the paper)
  - filters: ▸ Appendix
  - Daily Martin and Chabi-Yo–Loudis bounds, January 1990 – December 2020, from S&P 500 index options, interpolated to 30, 90, 180 and 360 days
  - S&P 500 excess returns over 21, 63, 126 and 252 trading days, annualised, in %
  - compounded as $\prod(1 + r - r_f) - 1$, the convention of the authors' code; it differs from $\prod(1 + r) - \prod(1 + r_f)$
- **Moments** (eq. 10): $E[(R^e_{t,T} - b_{t,T})\, z_t] \geq 0$
  - for any nonnegative z_t known at t
  - z_t: a constant and nine Welch–Goyal predictors made positive Welch & Goyal (2008) (Table 3)
- **Forecasting**: bias-corrected predictive regressions Amihud & Hurvich (2004) (Tables 5, 6)
  - Amihud–Hurvich: a persistent predictor biases the OLS (Ordinary Least Squares) slope (Stambaugh bias), so the innovation of the predictor, estimated from a bias-corrected AR(1), is added as a regressor
- Monthly out-of-sample forecasts against the historical mean
  - measured by R^2_{OS}, the out-of-sample R^2 against that mean
  - Table 8; protocol: ▸ Appendix

Case Study: the Kodde–Palm Tests

- Kodde & Palm (1986) (eqs. 11–14): $\bar{\lambda}$ is the vector of the 10 sample means
  - Σ: their Hansen–Hodrick covariance Hansen & Hodrick (1980)
  - with lags = days in the horizon
  - The Hansen–Hodrick covariance keeps, with equal weights, exactly the autocovariances created by the overlap of the return windows
- $D_1 = \min_{\lambda \geq 0}(\lambda - \bar{\lambda})^\top \Sigma^{-1}(\lambda - \bar{\lambda})$ tests validity
  - H_0: $\lambda_0 \geq 0$
  - the distance to the region where all moments are ≥ 0
  - zero if all are already positive
- $D_2 = \bar{\lambda}^\top \Sigma^{-1} \bar{\lambda} - D_1$ tests tightness (H_0: $\lambda_0 = 0$ vs $\lambda_0 \geq 0$): a large D_2 rejects tightness
  - Example with two moments and $\Sigma = I$: $\bar{\lambda} = (0.5, -0.2)$ gives $D_1 = 0.2^2 = 0.04$ and $D_2 = 0.29 - 0.04 = 0.25$; with $\bar{\lambda} = (0.5, 0.2)$, $D_1 = 0$ and $D_2 = 0.29$
- Under H_0 some inequality constraints bind and others do not
  - both statistics follow weighted mixtures of $\chi^2_0, \ldots, \chi^2_{10}$
  - bounds on the critical values tabulated by Kodde & Palm (1986)

Case Study: Results, Validity and Tightness (Tables 1 and 3)

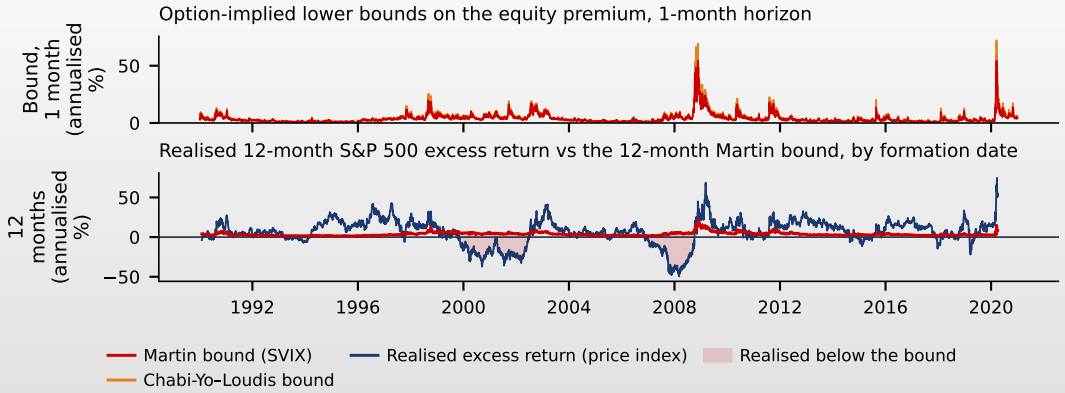
Horizon	Mean R^e (%)	Mean b (%)	D_1	D_2	$\hat{\beta}$ (SE)	R^2_{OS} (%)
1 month	8.76	3.81	0.00	29.93	1.36 (1.23)	−1.2
3 months	8.73	3.86	0.00	39.51	1.39 (1.47)	−1.2
6 months	8.67	3.83	0.00	49.28	1.78 (0.85)	−0.2
12 months	8.86	3.68	0.00	71.21	1.04 (0.88)	−3.0

- Source: Back, Crotty & Kazempour (2022), *Journal of Financial Economics*, Tables 1, 3, 5, 6 and 8
- **Validity not rejected**: $D_1 = 0$ and $p = 100\%$ at every horizon
  - all 10 sample moments are positive
  - this does not prove validity in every state
- **Not tight**: D_2 exceeds 28.86 at every horizon
  - the upper 0.1% Kodde–Palm critical value for 10 restrictions
  - finite-sample p: 3.6, 0.4, 2.6, 8.3%
  - mean slackness 4.8 to 5.2 points a year: the bound is less than half of the realised premium
- Type of evidence: description
  - mean slackness over 1990–2020, not forecasting power: out of sample the bound loses to the historical mean (Further Topics)

Case Study: Do the Two Tightness Tests Agree?

- Same D_2, two yardsticks
  - the asymptotic critical value rejects tightness at every horizon
  - the finite-sample p-values reject at 5% at 1, 3 and 6 months, but not at 12 months
- With a Holm correction for four horizons (Section 11), the finite-sample test rejects only at 3 months
  - Holm compares the smallest p with $5\%/4$, the next with $5\%/3$: $0.4\% < 1.25\%$, then $2.6\% > 1.67\%$, so it stops
- The asymptotic test rejects at every horizon even after the correction
- Why they can disagree: Σ is imprecise
  - 1990–2020 holds only about 31 non-overlapping 12-month windows
  - the moments use persistent predictors
  - so asymptotic critical values can be too small
- The finite-sample p is the more cautious reading
  - report both, never only the one that rejects (Seminar 12, C3)
- The mean slackness of 4.8–5.2 points a year does not depend on the test
  - economically, the bound covers less than half of the realised premium at every horizon

Case Study: the Bound and the Realised Premium (1/2)



MFM_ch12_bck_bounds ▶ Interactive

Case Study: the Bound and the Realised Premium (2/2)

- Top: the authors' daily 1-month bounds, 1990–2020 (replication package)
  - bottom, course data: 12-month excess return of the S&P 500 price index
  - compared with the 12-month Martin bound
- 1-month bound: median 2.7%, peak 54.4% (16 March 2020)
  - 12-month bound: peak 21.1% (20 November 2008)
  - A risk-neutral variance, the bound jumps with the VIX in crises
- Chabi-Yo–Loudis $\geq$ Martin on 99.99% of days
  - negative skewness and fat tails raise it
- Realised 12-month excess return below the bound on 33% of dates (2000–2002, 2007–2008)
  - No contradiction: the bound restricts $E_t[R^e_{t,T}]$, not each realisation
- Correlation bound–realised 0.15 (overlapping, descriptive only)
  - return noise dwarfs the bound
  - hence imprecise slopes and negative R^2_{OS}

MFM_ch12_bck_bounds

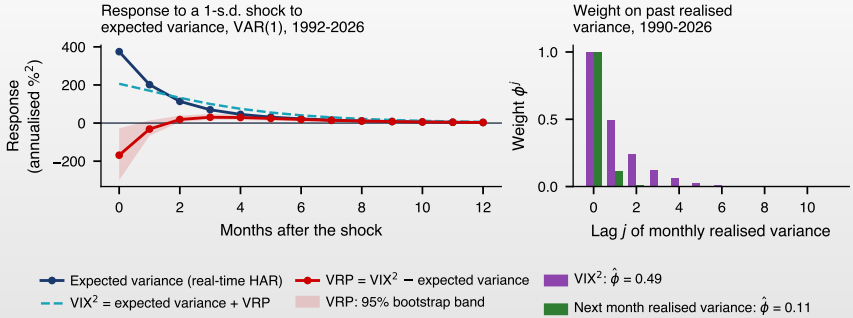
Case Study: Replication and Limitations

- **Replicate** (Seminar 12, C3)
  - the authors' daily bound series (replication package)
  - the S&P 500 daily closes of the course data
  - The daily risk-free rate (Kenneth French Data Library) and the Welch–Goyal predictors (Amit Goyal's data page)
  - Rebuild Table 1 and D_1, D_2 of Table 3
- **Limitations stated by the authors**
  - The series starts only in 1990 because the bounds need option prices: the short sample leaves the out-of-sample tests little power (paper, Section 6)
  - Validity holds even when the bound is replaced by its time-series mean, so it says nothing about predictive power (paper, Section 4.3)
  - More recent bounds and formulas for risk premia are not tested (paper, Introduction)

Does the VIX Underreact to Volatility?

- Lochstoer & Muir (2022): investors update their expectations of market volatility too slowly
  - underreaction first, after a volatility shock, then delayed overreaction
- Why it matters: VIX^2 is the risk-neutral price of 30-day variance
  - if the VIX rises less than expected variance, the VRP (variance risk premium) falls exactly when risk rises
  - VRP = VIX^2 minus the objective expected variance $\hat{\sigma}_t^2$
- The market's expectation as a weighted average of the past (equation 4 of the paper): $\text{VIX}_t^2 = a + b \sum_{j=0}^{11} \phi^j \, rv_{t-j} + \varepsilon_t$
  - rv_t: realised variance of month t, the sum of squared daily S&P 500 log returns, annualised
  - $\phi = 0$: only the current month counts; a larger ϕ means expectations that move more slowly
- The paper's numbers (Journal of Finance, Section I): $\phi = 0.42$ (standard error 0.12) for VIX^2, January 1990 to April 2020
  - after a shock to expected variance, the VRP response is negative at first and changes sign after the third month (Figure 2)
- The consequence the authors propose: a weak risk–return tradeoff, because the premium rises only once volatility has started to fall

How the VIX Responds to a Shock to Expected Variance



- **Left: the monthly response, in annualised %², to a one-standard-deviation shock to $\hat{\sigma}^2$ in a vector autoregression (VAR), 1992–2026**
- **Right: the weights $\hat{\phi}^j$ on the variance of month $t-j$, for VIX² and for next month's realised variance**

MFM_ch12_vol_expectations ▶ Interactive

The Paper and the Course Replication

- The paper: ϕ estimated from equation 4 for VIX^2 and, as the rational benchmark, for subsequent realised variance
  - $\hat{\sigma}_t^2$: the real-time forecast of next month's variance, a HAR (heterogeneous autoregressive) model on 5-minute returns
  - Table III: regressions on $\hat{\sigma}_t^2$ and on the weighted average of the last 6 values of $\hat{\sigma}^2$, with weights proportional to 0.5^j
  - a monthly VAR(1) with $\hat{\sigma}^2$ ordered first (Cholesky) gives the responses to the shock
- The course replication: daily S&P 500 and VIX, January 1990 to Aug 2026 (440 months)
  - $\hat{\sigma}_t^2$: HAR on daily returns (last day, last 5 and last 21 days), re-estimated each month on the data available then
  - the first 24 months are the burn-in, so $\hat{\sigma}^2$ starts in January 1992
  - the VAR has three variables: $\hat{\sigma}^2$, the VRP and the excess return of the price index (risk-free rate: Kenneth French)
  - the paper adds the log price–dividend ratio, not in the course data
- Inference
  - standard errors of $\hat{\phi}$: moving block bootstrap with 18-month blocks, like the paper's 6 quarters
  - Table III: Newey–West standard errors with 12 lags; VAR bands: residual bootstrap
  - HAC (Heteroskedasticity and Autocorrelation Consistent) standard errors such as Newey–West need $E\|x_t u_t\|^2 < \infty$ (x_t: regressors, u_t: errors) and weak dependence; in variance units, eighth moments of returns, which Chapter 1 questions

The VIX Depends More Persistently on Past Variance

- Replication: $\hat{\phi}$ for VIX^2 is 0.54 (standard error 0.08) on the paper's sample (paper: 0.42), 0.49 (0.08) up to 2026
  - for next month's realised variance: 0.26 (0.30) and 0.11 (0.20), not significantly different from zero
- Table III, 334 months, 1992–2020: next month's variance depends on $\hat{\sigma}_t^2$ (0.82, s.e. 0.12); the past-average coefficient is not significant (0.04, s.e. 0.10)
  - in the paper (Panel A, column 6): 1.53 (0.54) and -0.58 (0.46)
- VIX^2 depends on both: 0.56 (0.09) and 0.41 (0.04), $R^2 = 0.79$
  - in the paper (Panel A, column 7): 0.30 (0.12) and 0.64 (0.07), $R^2 = 0.79$
- VAR: a shock of 375 $\%^2$ to $\hat{\sigma}^2$ raises VIX^2 by only 206 $\%^2$ (55% of the shock)
  - VRP response: -169 $\%^2$ on impact (pointwise 95% band: $[-300, -27]$), positive from month 2, $+30$ in month 3
- Type of evidence: a description consistent with slow expectations, on daily instead of 5-minute data
  - it does not identify the mechanism: a rationally time-varying price of variance risk could give the same signs
- The limits
  - the daily HAR forecasts less well: $R^2 = 0.32$ for next month's variance, against 0.47 in the paper
  - part of the VIX weight on the past may be information our HAR lacks
  - months of very high variance (2008, 2020) dominate the regressions in levels; the Cholesky ordering changes the initial sign (next slide)

Inference in the Case Study: What the Bands Cover

- The difference $\hat{\phi}_{\mathrm{VIX}} - \hat{\phi}_{RV}$: bootstrap on the same 18-month blocks for both regressions
  - paper's sample: 0.29, 95% interval $[-0.05, 0.91]$, so the difference is not significant at 5%
  - up to 2026: 0.38, 95% interval $[0.08, 0.80]$, significant at 5%
- The VAR bands are conditional on the estimated $\hat{\sigma}^2$ series
  - the bootstrap resamples the VAR residuals without re-estimating the first-stage HAR
  - the bands are pointwise, horizon by horizon, not simultaneous over all horizons
  - the Newey–West errors of Table III also leave out the uncertainty of $\hat{\sigma}^2$
- VAR(1) diagnostics
  - stable: the largest eigenvalue has modulus 0.73
  - the return residuals are conditionally heteroskedastic (Engle test, 12 lags, $p < 0.001$)
  - resampling the residuals independently is therefore only an approximation
- The Cholesky ordering matters
  - with $\hat{\sigma}^2$ ordered last, the VRP response is 0 on impact, by construction, and $+33$ %^{2} in month 1
  - the negative initial sign rests on $\hat{\sigma}^2$ not responding within the month to VRP shocks

The Section's Question and the Tools to Test It

- The section's question: does the VRP (Section 8) still predict equity returns after 2007?
- **Hodrick 1B** Hodrick (1992): keep the regression of overlapping h-month returns on VRP_t
  - only the covariance estimator of the coefficient changes
  - its score: the one-month residual under H_0 times the sum of the last h values of the regressor
  - so the covariance needs no autocovariances of overlapping errors
- $R^2_{OS} > 0$: the predictive model beats the historical mean
  - lower mean squared forecast error than the historical mean
  - **Clark–West** Clark & West (2007) tests this for nested models
  - it adjusts for the noise that estimating a truly zero slope adds to the larger model
- **Holm correction** for m tests (e.g. four horizons)
  - sort the p-values
  - compare the smallest with α/m, the next with $\alpha/(m-1)$, and so on
  - stop at the first that fails
  - Example: $p = 0.01, 0.03, 0.04$, $\alpha = 0.05$, $m = 3$: $0.01 < 0.0167$ rejects, $0.03 > 0.025$ stops, so only one hypothesis is rejected, although all three $p < 0.05$

An Open Question: Does the VRP Still Predict Returns?

- Bollerslev, Tauchen & Zhou (2009), Table 2: monthly data 1990–2007
  - $y_t^{(h)} = \frac{12}{h} \sum_{j=1}^{h} (r_{m,t+j} - r_{f,t+j})$ on $\mathrm{VRP}_t = \mathrm{VIX}_t^2/12 - \mathrm{RV}_t$
  - RV_t: realised variance of month t; both terms in monthly, not annualised, $\%^2$ (a factor 12 below Section 8)
  - At $h = 3$: slope 0.47, Hodrick $t = 2.86$, adjusted $R^2 = 6.82\%$, the peak of a hump-shaped horizon profile
  - the Hodrick t on a variance regressor requires a finite fourth moment, which Chapter 1 questions
- Goyal, Welch & Zafirov (2024): 29 equity-premium predictors proposed after 2008, re-examined with data to 2021
  - over one third lose in-sample significance
  - half of the rest fail out of sample
- Open: is VRP predictability a stable price of variance and jump risk, or a property of 1990–2007?
  - the 2008–2026 sample adds two crashes and the 0DTE era
- Formally: $H_0 : \beta_3 = 0$ vs $H_1 : \beta_3 > 0$ in $y_t^{(3)} = \alpha + \beta_3 \mathrm{VRP}_t + u_t$
  - Out of sample: H_0: the VRP forecast does not lower the mean squared prediction error of the historical mean; Clark–West statistic Clark & West (2007)

The Discovery Loop and What the Researcher Must Check (1/2)

- **Literature**: ask an assistant for peer-reviewed tests of VRP predictability since 2009
  - DOI (digital object identifier), sample, table
  - Check: resolve every DOI on Crossref and compare the cited table with the paper, digit by digit
- **Hypotheses**: three candidate explanations
  - time-varying risk aversion
  - jump-risk compensation
  - or a small-sample artefact of persistent regressors and overlapping returns
  - Check: each must imply something testable for the horizon profile, for 2008–2026 and for non-US markets

The Discovery Loop and What the Researcher Must Check (2/2)

- **Data and code**: build $\text{VIX}^2/12$ and monthly RV from the course data
  - from the daily VIX and S&P 500 closes
  - reproduce Table 2 ($h = 3$) before extending anything
  - Check timing (VIX at the end of month t, returns from $t + 1$) and match the published means (implied 33.23, realised 14.93)
- **Robustness and critique**: a dated pre-analysis plan
  - horizons, start date, benchmark, Holm correction
  - a hostile-referee prompt
  - Check overlap (Hodrick errors Hodrick (1992)), horizon and start-date mining (R^2_{OS} against the start date), the weight of 2008 and 2020; full prompts: Appendix

Mini-Case: Replication and Extension (1/2)

- A typical assistant proposal, tested on the course data
  - "the VRP predicts the quarterly equity premium; extend the 2009 result to today"
- Replication, February 1990 to September 2007, $h = 3$, daily data
  - $\hat{\beta}_3 = 0.42$, $t_{\text{Hodrick}} = 2.25$, adjusted $R^2 = 4.6\%$
  - published: 0.47, 2.86, 6.82%
  - Mean implied variance 33.1 is close to 33.23; mean realised 20.6 vs 14.93: daily close-to-close data
- Extension to May 2026 (436 months)
  - $\hat{\beta}_3 = 0.09$, $t_{\text{Hodrick}} = 1.33$, adjusted $R^2 = 1.0\%$
  - all horizons with four kinds of standard errors: Appendix

Mini-Case: Replication and Extension (2/2)

- Out of sample from January 2008: 221 expanding-window forecasts
  - $R^2_{OS} = 1 - \sum_t (y_t - \hat{y}_t)^2 / \sum_t (y_t - \bar{y}_t)^2 = -22.2\%$ Campbell & Thompson (2008); $\bar{y}_t$: historical mean up to t
  - Clark–West $t = 0.64$ ($p = 0.26$): H_0 not rejected
  - the Clark–West test on squared errors requires a finite fourth moment, which Chapter 1 questions
- Reading: sign and size replicate in-sample
  - the longer sample does not support stable predictability
  - outliers, measurement or a vanished premium: that is the open part
- Type of evidence: comparison of procedures (it shows which method works better on the same data, not why)
  - the VRP forecast against the historical mean

MFM_ch12_ai_discovery

Your Turn: a Project on the VRP After 2008 (1/2)

- Questions
  - Does the VRP predict the equity premium after 2008?
  - Does the Bitcoin variance risk premium (DVOL, Further Topics) predict Bitcoin returns?
- Replicate first: Bollerslev, Tauchen & Zhou (2009), Table 2, $h = 3$
  - the published numbers of the open question
  - explain every gap

Your Turn: a Project on the VRP After 2008 (2/2)

- Deliverables
  - A dated pre-analysis plan; the replication table; out-of-sample results for all horizons with a Holm correction
  - Robustness (subsamples, without 2008 and 2020; optionally 5-minute realised variance, if you have intraday data); a referee-style critique and your reply
- Log each prompt that shaped a decision in `AI_USE.md`
  - log each invented reference, wrong number or look-ahead caught in `AI_ERRORS.md`
  - with the test that caught it (templates)

What Option Prices Measure, and Their Limits

- Parity, bounds and static spanning need no model
  - the tree and the Black–Scholes equation price by dynamic replication
- Discrete hedging leaves an error of order $1/\sqrt{N}$
  - a delta-hedged option is a bet on realised vs implied variance
- Implied variance exceeded realised variance on average, on the S&P 500 and, in Further Topics, on Bitcoin
- The surface and the risk-neutral density are estimates
  - check the no-arbitrage constraints, reject estimates that violate them and report bands
- The VIX is a strip estimator with jump and truncation biases
  - as a forecast it needs joint tests; logs and HAC loss comparisons: in Further Topics
- Option-implied equity-premium bounds: validity not rejected for the moments and sample tested; not tight in 1990–2020
  - evidence rests on average slackness, not on forecasting power
- "Not tight" is robust on the asymptotic test
  - less certain on the finite-sample tests after a multiple-testing correction

Check your understanding (1/2)

- Which pricing restrictions follow from no arbitrage alone?
- Which of the restrictions used in the chapter depend on the Black–Scholes assumptions?
- Why can a delta-hedged option lose money even when the model is right?
- What role does the gap between realised and implied variance play in the hedging result?
- Does a fitted smile automatically exclude arbitrage?

Check your understanding (2/2)

- How do you check the butterfly constraint for an SVI fit?
- Does a risk-neutral density forecast the real-world frequencies of returns?
- Why is the standard error $s/\sqrt{n}$ wrong for a mean over overlapping 21-day windows?
- Which standard error replaces it?
- How can two small t-statistics coexist with a large joint Wald statistic?

Open Questions in Option Research (1/2)

- **Frontier**
  - Andersen, Fusari & Todorov (2015): a left-tail factor in index-option surfaces, not spanned by volatility, predicts future equity and variance premia
  - Bali et al. (2023): nonlinear machine learning on option and stock characteristics predicts equity-option returns out of sample
  - Open: jointly estimated, arbitrage-free surfaces with valid confidence bands; the pricing of 0DTE options
- **Next**: Chapter 13, machine learning for financial prediction
  - volatility and option-implied features among the inputs

Open Questions in Option Research (2/2)

- **Further reading**
  - Textbooks: Franke, Härdle & Hafner (2019), Härdle, Chen & Overbeck (2017), Čížek, Härdle & Weron (2011), Gatheral (2006)
  - Pricing kernels: Aït-Sahalia & Lo (1998), Grith, Härdle & Park (2013); crypto options: Matic, Packham & Härdle (2023), Alexander & Imeraj (2020), Pele et al. (2023)
  - Quantinar: XFG (Quantlet prefix of Applied Quantitative Finance, Härdle, Chen & Overbeck) Advanced Methods in Quantitative Finance, Pricing Kernels and Risk Premia, Hedging Cryptocurrency Options
  - Quantlets: github.com/danpele/MFM/tree/main/Quantlets/Ch_12
- Detailed examples and advanced topics are in Further Topics, after the Conclusions
  - the footer button leads there; the Back button returns to the starting slide

The Practical Decision

- **What I choose**
  - for hedging: delta hedging at a frequency set by the error and the cost model; the s.d. of the error falls like $1/\sqrt{N}$ (weekly 21.5%, daily 9.5% of the premium)
  - for the smile: an SVI fit per expiry, accepted only after the butterfly and calendar checks; the density with bootstrap bands and the quoted range marked
  - for next month's variance: VIX^2, read with the joint Mincer–Zarnowitz test ($\hat{\beta} = 0.86$, unbiasedness rejected: $\chi^2_2 = 41.7$)
  - not the VRP as a return predictor after 2008: $R^2_{OS} = -22.2\%$, Clark–West $p = 0.26$
- **Under which conditions**
  - the density and the VRP are risk-neutral: prices of insurance, not forecasts of frequencies
  - overlapping 21-day windows: Newey–West with 21 lags and sensitivity to 42, 63 and 126 (standard error of the premium: 0.27 to 0.33 points)
  - mean premium: finite fourth moment of returns; coefficients and loss differences: finite second moment of the score or of the difference, and weak dependence (Chapter 1)
- **What result would make me drop it**
  - an SVI fit: $g(k) < 0$ or total variance falling with maturity; I then move to SSVI
  - the VIX as a forecast: past variance taking over its weight in the joint regression (here 0.08, SE 0.13, against 0.77)

The Data: Volatility Indices and a Bitcoin Option Snapshot

Series	What it measures	Period	Owner
S&P 500	US equity index, daily close	1990–2026	S&P Dow Jones Indices
SPY	exchange-traded fund (ETF) on the S&P 500: open and close	1993–2026	State Street
VIX	30-day implied volatility of the S&P 500	1990–2026	Cboe
VIX9D, VIX3M	the same over 9 days and 3 months	VIX9D: 2011–2026; VIX3M: 2007–2026	Cboe
VVIX	implied volatility of the VIX	2006–2026	Cboe
Bitcoin options	all listed contracts, one dated snapshot	22 Sep 2026	Deribit
DVOL	30-day implied volatility of Bitcoin	2021–2026	Deribit

- Daily data up to 18 September 2026
  - the option snapshot was taken on 22 September 2026, 07:51 UTC (UTC: Coordinated Universal Time)
- Implied volatility is quoted in annualised %; "vol points" = percentage points of volatility

Notation: Prices, Expiry and Moneyness (1/2)

- S_t: price of the underlying at t; K: strike; T: expiry
  - $\tau = T - t$: time left, in years
  - r: continuously compounded risk-free rate
  - δ: dividend yield
- C, P: prices of a call and a put
  - $(x)^+ = \max(x, 0)$
  - σ_{imp}: implied volatility (Section 6)
- European call and put pay off $(S_T - K)^+$ and $(K - S_T)^+$ at T
  - American options can be exercised at any time up to T
- Both pay-offs are convex in S_T, so an option's value increases with volatility (vega > 0, Section 4)

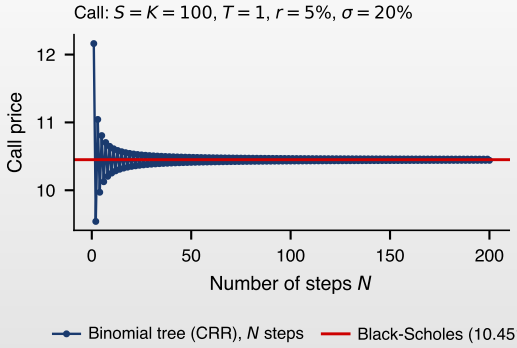
Notation: Prices, Expiry and Moneyness (2/2)

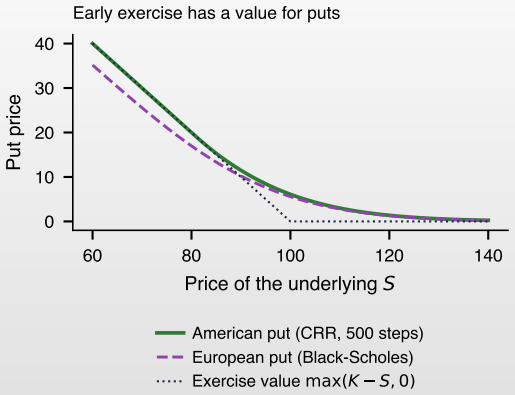
- Volatility σ: the annualised standard deviation of log returns $\ln(S_t/S_{t-1})$
  - Normal under Black–Scholes
- Used in most of the chapter: forward $F = Se^{(r-\delta)\tau}$
  - log-moneyness $k = \ln(K/F)$
  - total implied variance $w = \sigma_{imp}^2 \tau$
- ATM (at the money): $k \approx 0$
  - OTM (out of the money): puts with $K < F$ and calls with $K > F$
  - ITM (in the money): the opposite

Context-dependent Symbols, Model-Free Bounds and Prerequisites

- $q(K)$: the risk-neutral density of S_T (Sections 2 and 7)
  - in the binomial tree, q is the risk-neutral up-probability (Section 3); the dividend yield is δ
- Q: the pricing measure; $Q(K_i)$: the Cboe mid quote at strike K_i (Section 8)
- Model-free bounds (no dividends): $\max(S - Ke^{-r\tau}, 0) \leq C \leq S$ and $\max(Ke^{-r\tau} - S, 0) \leq P \leq Ke^{-r\tau}$
  - With $r > 0$ an American call on a stock without dividends is never exercised early: it is worth at least $S - Ke^{-r\tau} > S - K$
- Moneyness is measured here against the forward F
  - Seminar 12 and the smile chart of Section 6 use the spot S (OTM call: $K > S$)
  - The two differ by the factor $F/S = e^{(r-\delta)\tau}$, close to 1 for short expiries: only strikes between S and F change class
- Assumed knowledge
  - Brownian motion, Itô's lemma and martingales (Chapter 11)
  - HAC (heteroskedasticity and autocorrelation consistent) standard errors (Chapter 8)

From the Tree to Black–Scholes





MFM_ch12_binomial

Convergence and the Early-Exercise Premium

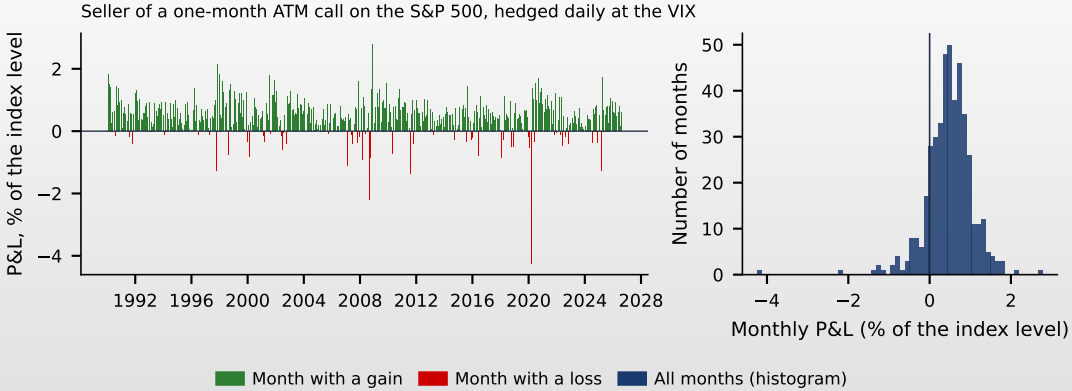
- Left: CRR price of a European call for $N = 1, \ldots, 200$ steps
  - $S_0 = K = 100$, $T = 1$, $r = 5\%$, $\sigma = 20\%$
  - right: American vs European put
- Converges to Black–Scholes (10.451)
  - $N = 50$: 10.411; $N = 200$: 10.441; $N = 1000$: 10.449
  - The error is of order $1/N$ and oscillates: $N = 10$ gives 10.253, $N = 11$ gives 10.611 (strike on or between nodes; Seminar 12, A6)
- American put at $S = K = 100$: 6.090 vs European 5.574
  - early-exercise premium 0.516
  - deep in the money it pays to receive K now

MFM_ch12_binomial

Selling S&P 500 Volatility: the Strategy Tested

- The question: is implied volatility priced, on average, above realised volatility?
  - a sold, delta-hedged option answers it directly
- Every 21 trading days t_0: sell a synthetic 21-day ATM call ($K = S_{t_0}$)
  - priced by Black–Scholes with $\sigma = VIX_{t_0}/100$
  - the VIX comes from a strip of options, not the quoted implied volatility of the ATM call: the premium collected is set by construction
- Hedge it daily with $\Delta_i = \Phi(d_1)$ at the same volatility
  - interest and dividends are ignored ($r = \delta = 0$)
- Monthly P&L, in % of the index:
  - $\text{P\&L} = 100\left[C_0 + \sum_{i=0}^{20} \Delta_i \left(S_{t_{i+1}} - S_{t_i}\right) - \left(S_{t_{21}} - K\right)^+\right] / S_{t_0}$
- Why the VIX: it is the 30-day implied volatility available daily since 1990 in the course data
- Why ATM: vega is largest near the money, so the P&L mainly measures volatility

Thirty-Six Years of Selling S&P 500 Volatility (1/2)



MFM_ch12_delta_hedged_sp500

Thirty-Six Years of Selling S&P 500 Volatility (2/2)

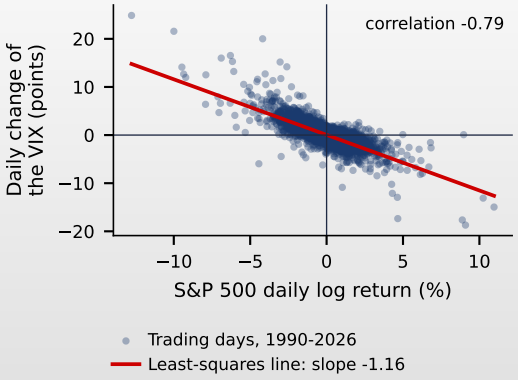
- Left: P&L of every month, in % of the index (green: gain, red: loss)
  - right: histogram of the monthly P&L
- 440 months from 1990-01: mean 0.48% a month, Newey–West $t = 16.0$ Newey & West (1987)
- Positive in 86% of months
- Worst month -4.25% (2020-03), skewness -1.46
  - many small gains, a few large losses
- Consistent with a **volatility risk premium**
  - On real option quotes, delta-hedged option gains are negative for the buyer Bakshi & Kapadia (2003)
  - here: a synthetic call priced at the VIX, without historical ATM quotes, bid–ask, dividends and financing; not the result of an executable strategy

MFM_ch12_delta_hedged_sp500

Implied Volatility: Newton Inversion

- Newton's method: $\sigma_{k+\mathbf{1}} = \sigma_k - \dfrac{BS(\sigma_k) - C^{market}}{\mathcal{V}(\sigma_k)}$
  - σ_k: the k-th iterate; $\mathcal{V}(\sigma_k)$: vega at σ_k, the slope of the price in σ
  - Why Newton: vega is the exact derivative of the price in σ, so convergence is quadratic near the solution
- Example: $S = K = 100$, three months, $r = 2\%$, market price 4.50
  - start at $\sigma_{\mathbf{0}} = 20.00\%$
  - First step: $BS(\sigma_{\mathbf{0}}) = 4.2322$, $\mathcal{V} = 19.848 \Rightarrow \sigma_{\mathbf{1}} = 21.3495\%$
  - The second step changes σ by less than $10^{-\mathbf{6}}$: $\sigma_{imp} = 21.3495\%$
- When vega is close to zero, a Newton step can leave the no-arbitrage range
  - vega is close to zero deep in or out of the money and at very short expiries
  - then use a bracketing solver

The Leverage Effect in the Data



- Daily S&P 500 log return (%) vs the change of the VIX (index points) on the same day, 1990–2026
- Correlation -0.79 over 9 244 days
  - implied volatility rises when the market falls
- Slope -1.16 VIX points per 1% return
  - Steeper on down days (-1.39) than on up days (-1.01)
  - Mean VIX change: $+0.88$ on down days, -0.76 on up days
- The association is consistent with the leverage effect and motivates asymmetric models (Heston with $\rho < 0$, GJR-GARCH)
  - it does not estimate Heston's ρ or the GJR coefficient directly: ΔVIX also contains changes in expectations and risk premia
  - GJR: Glosten–Jagannathan–Runkle, Chapter 5
  - Low-strike puts gain both from the fall and from the rise in volatility, hence they are expensive

MFM_ch12_leverage

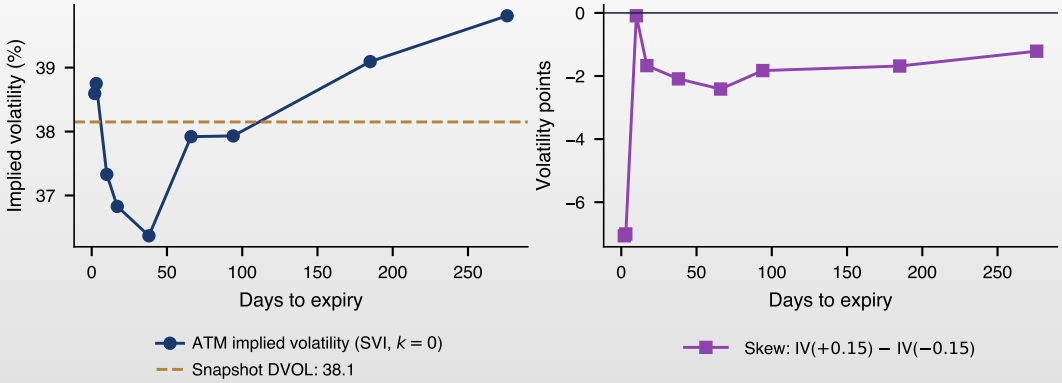
Inverse Options: Why the USD Formulas Still Apply

- An inverse call pays $(S_T - K)^+ / S_T$ BTC at T
  - converted at S_T USD per BTC this is exactly $(S_T - K)^+$ USD
  - the pay-off of a standard USD call
  - Example: $K = 90\,000$, $S_T = 100\,000$: the holder receives $10\,000/100\,000 = 0.1$ BTC, worth $10\,000$ USD
- Same pay-off, same value: $p \times S$ is the USD price of the standard call
  - a premium of p BTC paid today costs $p \times S$ USD at today's index S
  - Black–Scholes/Black-76 and the exchange's mark IV (its reference implied volatility) therefore apply to $p \times S$ unchanged; this is the conversion used in the smile, density and VIX slides
- The rate: Deribit quotes a forward F for each expiry
  - this chapter uses $e^{r\tau} = F/S$, the carry implied by the futures basis
  - in place of a Treasury rate
  - Check: an inverse put with $K = 90\,000$ ends at $S_T = 75\,000$. What does it pay in BTC and in USD? (Answer: 0.2 BTC, i.e. 15 000 USD)

Are the Nine Bitcoin Fits Arbitrage-Free?

- For 2 October, 9 October and 27 November, a hits its lower bound (-1)
  - the bound imposed in the optimisation
  - the individual parameters are not precisely identified
- Wings: largest slopes 0.51 (right) and 2.30 (left)
  - the left wing exceeds Lee's upper bound of 2
- Butterfly: $g(k) < 0$ somewhere on the grid $k \in [-1.5, 1.5]$ for 2 of 9 expiries
  - lowest -0.017, 2 October 2026
- Calendar: 1 of 8 adjacent pairs crosses
  - w(2 October 2026) $>$ w(9 October 2026) on 9% of $k \in [-0.5, 0.5]$
  - by at most 0.0018
  - Separate unconstrained fits can violate cross-expiry restrictions; these fitted curves do not establish arbitrage in executable market quotes
- Remedy: estimate the surface jointly under the constraints
  - SSVI, or least squares with $g \geq 0$ and $\partial_\tau w \geq 0$
  - also imposing the wing conditions

Term Structure and Skew



MFM_ch12_btc_term

Expiry, Skew and the DVOL Index

- Left: ATM implied volatility by expiry and the DVOL index
  - right: skew = IV($k = +0.15$)− IV($k = -0.15$)
  - IV: the implied volatility at that log-moneyness
- ATM between 36.4% and 39.8%; longest expiry (276 days) 39.8%
  - the curve falls at short maturities, then rises beyond about one month
- DVOL 38.1 vs ATM 36.8% and 36.4% around 30 days
  - a variance-strip index need not lie between ATM volatilities
- Skew closest to 0 at 10 days (-0.1), most negative at 66 days (-2.4), then -1.2 at 276 days
  - At 2 and 3 days, $k = \pm 0.15$ lies beyond the quotes (at most $k = 0.11$ and 0.14): the skew of about -7 points is an SVI extrapolation there
- Limitation: one snapshot is one day

MFM_ch12_btc_term

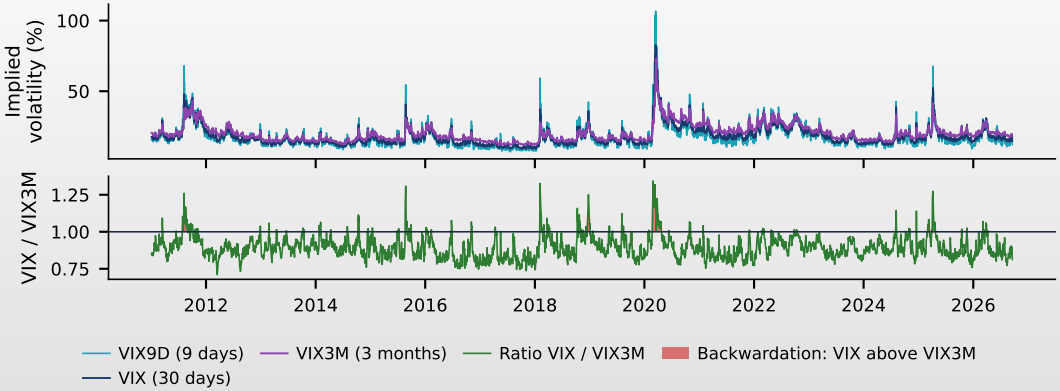
The VIX Index: Computation Rules

- Cboe Volatility Index (VIX), published since 1993 Cboe (2025)
  - since 2003 computed from S&P 500 options with the Cboe formula
  - the series under this rule is computed back to January 1990, the start of the charts
- $\Delta K_i = (K_{i+1} - K_{i-1})/2$, one-sided at the two ends
  - at the lowest (highest) strike, ΔK is the full gap to the next higher (lower) strike
- Options with a zero bid are dropped, and each wing stops after two consecutive zero bids
- Two expiries around 30 days, interpolated to exactly 30 days; VIX $= 100\,\sigma$
- VIX futures since 2004, VIX options since 2006 Cboe (2025)
- The VIX is called the "fear gauge" Whaley (2000)

Worked Example: a VIX for Bitcoin

- Cboe rules on our Deribit snapshot: USD mid Q = BTC mid $\times$ index S
  - F and $e^{RT} = F/S$ from Deribit's forward
  - Cboe: F from parity, R from Treasury yields
  - Near expiry: 17 days, 26 OTM options, $\sigma = 39.5\%$; next expiry: 38 days, 56 OTM options, $\sigma = 40.2\%$
- Interpolate the total variances $\sigma_i^2 T_i$ linearly in time to 30 days:
  - $\sigma_{30}^2 = \frac{1}{30}\left[T_1\sigma_1^2 \frac{T_2 - 30}{T_2 - T_1} + T_2\sigma_2^2 \frac{30 - T_1}{T_2 - T_1}\right]$ (days)
- Our Bitcoin VIX: 40.0; Deribit's DVOL at the same moment: 38.1 Alexander & Imeraj (2020)
  - both exceed the ATM volatilities
  - the variance swap also pays for the wings
- The gap to DVOL, 1.9 vol points, cannot be attributed to truncation
  - the SVI extrapolation it would require violates Lee's wing condition
  - Appendix; Seminar 12, A8

The VIX Term Structure (1/2)



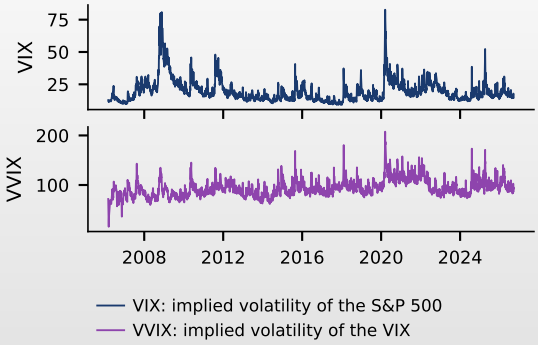
MFM_ch12_vix_term

The VIX Term Structure (2/2)

- Top: implied volatility over 9 days, 30 days and 3 months
  - bottom: the ratio VIX / VIX3M (Cboe 3-Month Volatility Index), 2011–2026
- Normally upward sloping (**contango**): mean VIX / VIX3M = 0.89
  - VIX9D (Cboe 9-Day Volatility Index) below the VIX on 73% of days
- Inverted (**backwardation**) on 7.6% of days, always in crises
  - fear is concentrated in the near term
- Over the next 21 days, after backwardation vs in contango:
  - S&P 500 log return 2.88% vs 0.80%; realised volatility 26.7% vs 13.6%
  - High risk and high expected return after panics Bollerslev, Tauchen & Zhou (2009)
- The comparison is descriptive
  - the 21-day windows overlap
  - backwardation days are few and clustered in a handful of crises

MFM_ch12_vix_term

The Volatility of Volatility



- VIX (top) and VVIX (bottom), daily closes, 2006–2026
- Mean VVIX 93.5; mean daily ratio VVIX/VIX 5.3
  - Both are annualised implied volatilities, of percentage changes of the VIX and of the S&P 500: in relative terms the VIX is expected to move about 5.3 times more than the index
- VVIX peak 207.6 on 2020-03-16, the day of the VIX peak
- Daily changes of VIX and VVIX: correlation 0.71
- Volatility is itself random, the idea behind stochastic-volatility models such as Heston Heston (1993)

MFM_ch12_vvix

Forecast-Evaluation Tools

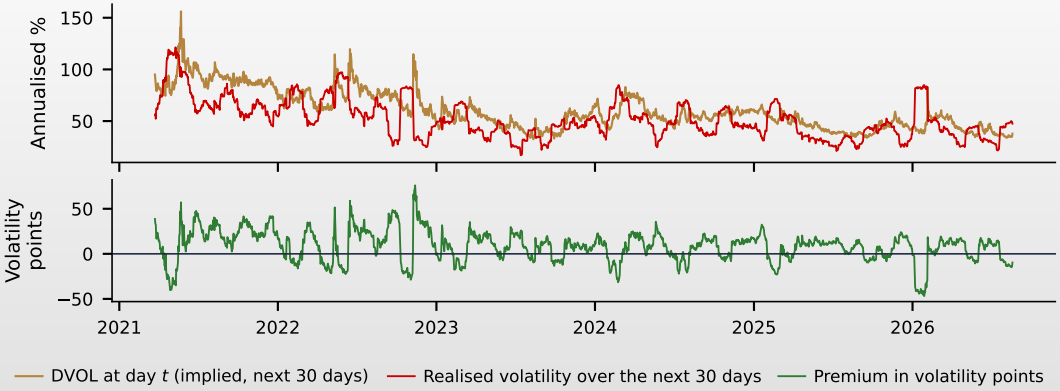
- **Instrumental variables**: the lagged VIX^2 instruments the VIX^2 measured with error
  - a large first-stage t (regression on the instrument) means a strong instrument
- **HAR-RV**: forecast the realised variance of the next month
  - by a regression on its averages over the past day, week and month
- **QLIKE** (Quasi-LIKElihood, loss function) loss $L = x - \ln x - 1$ with $x = RV/\hat{\sigma}^2$
  - zero for a perfect forecast
  - it punishes forecasts that are too low more than forecasts that are too high
  - Example: $RV = 100$; forecast 50 gives $x = 2$, $L = 0.31$; forecast 200 gives $x = 0.5$, $L = 0.19$
- **Diebold–Mariano**: a t-test that the mean loss difference between two forecasts is zero
  - with a HAC standard error
  - overlapping targets make the differences autocorrelated

Logs, Instruments and HAR-RV: Inference Done Right

Specification	$\hat{\beta}$	SE	$t(\beta = 1)$	Note
Levels, OLS	0.860	0.105	-1.34	joint Wald $p = 8.9 \cdot 10^{-10}$
Logs, OLS	1.06	0.035	1.65	intercept -0.94, $R^2 = 0.57$
Logs, instrumental variables	1.02	0.042	0.60	first-stage $t = 37.7$

- Levels are dominated by 2008 and 2020
  - in logs the slope is close to 1
  - the premium is a proportional shift (intercept < 0)
- VIX^2 measures expected variance with error, which attenuates the ordinary least squares (OLS) slope
  - The instrument, VIX^2 21 days earlier Christensen & Prabhala (1998), barely moves the slope
  - Valid only if the lag is uncorrelated with the forecast error (untestable with one instrument)
- Against HAR-RV Corsi (2009), in sample the HAR terms add information: $\chi^2_3 = 14.2$, $p = 0.003$
  - Out of sample from 2000: VIX about 6% better in QLIKE loss, Diebold–Mariano $p = 0.35$ Diebold & Mariano (1995): not significant
  - Test details: Appendix

The Variance Risk Premium of Bitcoin (1/2)



MFM_ch12_vrp_btc

The Variance Risk Premium of Bitcoin (2/2)

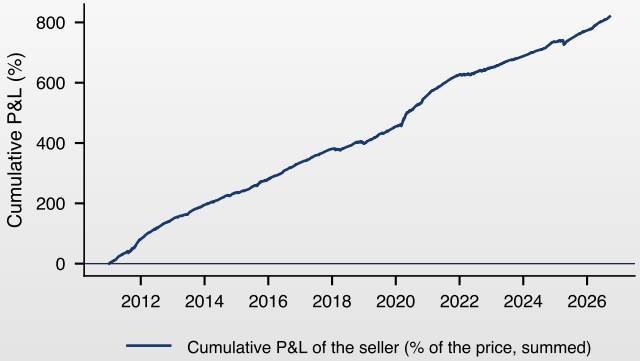
- Top: DVOL and Bitcoin volatility realised over the next 30 calendar days
  - Bitcoin trades every day
  - bottom: their difference, 2021–2026
- 1 975 days: mean DVOL 60.3 vs realised 51.7
  - Premium 8.5 vol points (SE 1.6, $t = 5.3$)
  - About 2.1 times the S&P 500 premium in vol points, but smaller relative to the level: 14% of mean DVOL vs 21% of mean VIX
- Positive on 71% of days, less often than the S&P 500
  - worst -46.9 vol points (window from 2026-01-28)
- Short sample: 66 non-overlapping 30-day periods Alexander & Imeraj (2020)

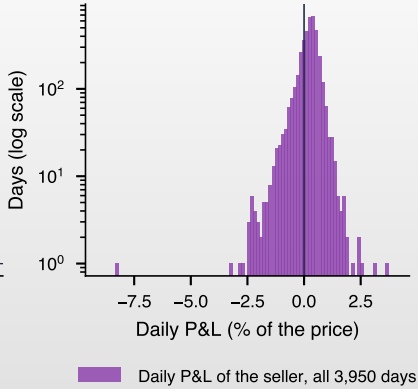
MFM_ch12_vrp_btc

Selling a Same-Day Straddle: the Method

- The question: does the seller of an option expiring today also earn the volatility premium?
  - the premium of one-month options
- The course data contain no 0DTE prices: we build a stylised test
  - on SPY (SPDR S&P 500 ETF Trust), 2011–2026 (VIX9D starts in 2011)
  - sell an ATM straddle (a call and a put with $K = S_t^{open}$) at the open
  - settle it at the close
- Premium: $\Pi_t = \sqrt{2/\pi}\,\sigma_{day}\,S_t^{open}$
  - the approximate Black–Scholes value of an ATM straddle
  - $E|Z| = \sqrt{2/\pi}$ for a standard normal Z
- $\sigma_{day} = \dfrac{VIX9D_{t-\mathbf{1}}}{100}\sqrt{\omega/252}$
  - the previous day's 9-day implied volatility, scaled to one session
- $\omega = 0.74$: the share of the session (open to close) in daily variance
  - estimated on 1993–2010, before the test period
- The seller's daily P&L, in % of the opening price:
  - $100\left(\Pi_t/S_t^{open} - \left|S_t^{close}/S_t^{open} - 1\right|\right)$

Selling a Same-Day Straddle: a Stylised Test (1/2)





MFM_ch12_0dte_straddle

Selling a Same-Day Straddle: a Stylised Test (2/2)

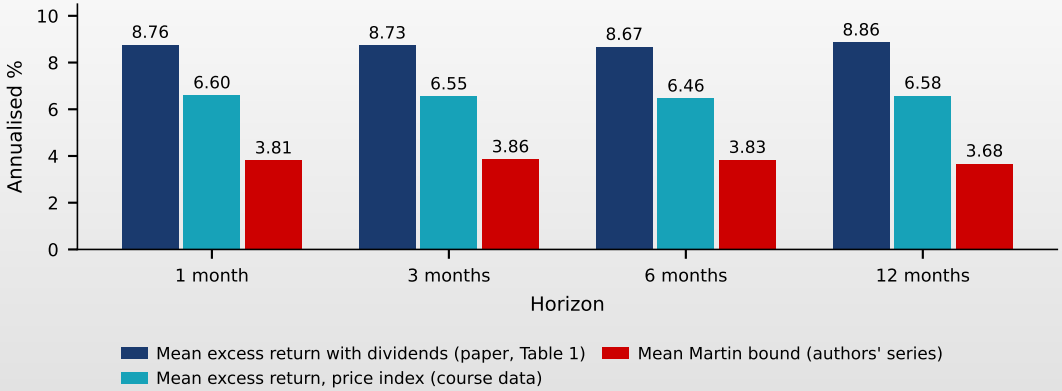
- Left: the seller's cumulative P&L (sum of daily P&L, in % of the price)
  - right: histogram of all daily P&L, log vertical axis
- 3 950 days: mean premium 0.76% vs mean absolute move 0.55%
  - the seller wins on 75% of days
  - 20.7 basis points a day
- Worst day −8.25% (2025-04-09), about 11 times the mean premium
  - stylised pricing measures the premium inside VIX9D, not traded 0DTE prices Bandi, Fusari & Renò (2023), Beckmeyer, Branger & Gayda (2023)

MFM_ch12_0dte_straddle

Case Study: Results, Forecasting (Tables 5, 6 and 8)

- **Forecasting**: slopes positive but imprecise (Table 5)
  - With the nine predictors as controls: a one-s.d. rise in the 12-month bound, a 6.6-point higher 12-month excess return (Table 6)
  - Bootstrap p-values 20.7%, below 0.1%, 0.2%, 9.5%; with a Holm correction for the four horizons (Section 11), only 3 and 6 months stay significant at 5%
- **Out of sample**: the Martin bound loses to the historical mean at every horizon
  - the best case, Chabi-Yo–Loudis at 6 months, reaches 4.3%, not significant (Table 8)
  - low power is one possible explanation

Case Study: Mean Premium and Mean Bound (1/2)



MFM_ch12_bck_bounds

Case Study: Mean Premium and Mean Bound (2/2)

- Dark blue: Table 1 of the paper (S&P 500 with dividends)
  - red: mean of the authors' Martin bound
- Teal: course data (price index, the same formation days)
  - 7 789 formation days at 1 month, 7 598 at 12 months
- Mean slackness on the course data: 2.79, 2.69, 2.63, 2.91 points a year
  - at 1, 3, 6 and 12 months
  - vs 4.8–5.2 in the paper
  - The gap of 2.2–2.3 points is mostly the dividend yield missing from the price index
- The mean bound covers 41–44% of the premium with dividends and 56–59% of the price-index premium
- Single realisations fall below the bound on 41% of days at 1 month
- The average realisation never falls below the mean bound (Seminar 12, C3)

MFM_ch12_bck_bounds

Case Study: an Open Extension

- The course data have VIX, not SVIX (Martin (2017), Result 3)
  - VIX weights OTM options by $1/K^2$
  - VIX exceeds SVIX when the left tail is priced
  - Is the Martin bound still valid on 2021–2026 (the 2022 bear market, the 0DTE era of Section 9)?
  - Is a VIX-based bound still valid?
  - If not, can the Kodde–Palm test detect it?

Proof of Spanning, Moments and SVIX

- Proof for $S_T > F$, by parts:
  - $\int_F^{S_T} f''(K)(S_T - K)\,dK = f(S_T) - f(F) - f'(F)(S_T - F)$
  - the put integral vanishes
  - symmetric for $S_T < F$
- Take $e^{-r\tau}E^Q[\cdot]$ with F the forward (a forward contract costs 0)
  - a bond, a forward and a strip of OTM options with weights $f''(K)$ replicate any smooth pay-off
- **Moments**: $f = [\ln(S/F)]^n$, $n = 2, 3, 4$, price the raw risk-neutral moments Bakshi, Kapadia & Madan (2003)
  - raw moments of the log return
  - centring and standardising gives volatility, skewness and kurtosis
- **SVIX** Martin (2017), from $f = S^2$:
  - $\mathrm{SVIX}_t^2 = \dfrac{2}{\tau R_f S_t^2}\left[\int_0^F P\,dK + \int_F^\infty C\,dK\right] = \mathrm{var}_t^Q(R_T/R_f)/\tau$
  - R_T, R_f: gross returns
  - SVIX^2 gives the Martin lower bound on the equity premium (case study, Section 10)

Back

Why $1/\sqrt{N}$? The Discrete Hedging Error

- Over one interval $\Delta t = T/N$: time decay against realised squared return
  - the hedged book earns $\frac{1}{2}\Gamma S^2\left[\sigma^2\Delta t - (\Delta S/S)^2\right]$
- Mean 0 under the model to leading order in Δt
  - $\mathrm{Var}[(\Delta S/S)^2] \approx 2\sigma^4\Delta t^2$
  - the terms are approximately uncorrelated
  - so $\mathrm{Var}(\text{error}) \approx \frac{1}{2}\sigma^4\Delta t\, E\int_0^T \Gamma_t^2 S_t^4\, dt$
- ATM call (derivation): $E[\Gamma_t^2 S_t^4] \approx \dfrac{K^2}{2\pi\sigma^2\sqrt{T^2 - t^2}}$
  - it integrates to $K^2/(4\sigma^2)$
  - so s.d. $\approx \sigma K\sqrt{T/(8N)} = \sqrt{\pi/4}\;\mathcal{V}\sigma/\sqrt{N}$
  - The last equality holds exactly for the vega at $d_1 = 0$, $\mathcal{V}_0 = K\sqrt{T}\,\varphi(0)$; here $d_1 = 0.15$ and $\mathcal{V} = 0.989\,\mathcal{V}_0 = 19.72$, so at $N = 12$ the two forms give 1.02 and 1.01
- Asymptotic theory of the replication error: Bertsimas, Kogan & Lo (2000)
  - it fails with jumps (the error does not vanish as $N \to \infty$)
  - with stochastic volatility (vega risk) and with costs

Back

The Gamma Moment Behind the $1/\sqrt{N}$ Law

- ATM call, $S_0 = K$, $r = 0$, remaining time $\tau = T - t$
  - $\Gamma_t S_t^2 = S_t\,\varphi(d_1)/(\sigma\sqrt{\tau}) \approx K\varphi(d_1)/(\sigma\sqrt{\tau})$
  - with $d_1 \approx X/(\sigma\sqrt{\tau})$ and $X = \ln(S_t/K) \sim N(0, \sigma^2 t)$
- Gaussian integral: $E[\varphi(d_1)^2] = \dfrac{1}{2\pi} E\left[e^{-X^2/(\sigma^2\tau)}\right] = \dfrac{1}{2\pi}\sqrt{\dfrac{\tau}{\tau + 2t}}$, and $\tau + 2t = T + t$
- Hence $E[\Gamma_t^2 S_t^4] \approx \dfrac{K^2}{2\pi\sigma^2}\,\dfrac{1}{\sqrt{(T-t)(T+t)}} = \dfrac{K^2}{2\pi\sigma^2\sqrt{T^2 - t^2}}$, and $\displaystyle\int_0^T \frac{dt}{\sqrt{T^2 - t^2}} = \frac{\pi}{2}$ gives the integral $K^2/(4\sigma^2)$

Back

The Moment Formula for the Wings

- Moment formula of Lee (2004): $\beta_R = \limsup_{k\to\infty} w(k)/k \in [0,2]$ and $\tilde{p} = \frac{1}{2\beta_R} + \frac{\beta_R}{8} - \frac{1}{2}$
  - $\tilde{p} = \sup\{p : E^Q[S_T^{1+p}] < \infty\}$: a steeper wing means fatter tails; the left wing mirrors it with negative moments
- Bitcoin: the largest right slope 0.511 gives $\tilde{p} = 0.54$
  - a formal extrapolation calculation only
  - this fitted slice fails a butterfly check, so it does not identify a valid risk-neutral moment
- For SVI a right-wing slope below 2 also gives the tail condition $d_+(k) \to -\infty$ of Gatheral & Jacquier (2014), Lemma 2.2

Back

Estimators of the Risk-Neutral Density

- Fit a smile (SVI) and differentiate the implied call prices twice, as in the lecture
- Kernel regression of C on (K, τ) Aït-Sahalia & Lo (1998)
  - the same under monotonicity and convexity in K Aït-Sahalia & Duarte (2003)
- Smoothness-penalised implied trees Jackwerth & Rubinstein (1996); shape-invariant pricing kernels Grith, Härdle & Park (2013)
- Tails beyond the quotes: SVI wings or an extreme-value tail
  - either way an assumption, not an estimate

Back

What the Strip Measures When Prices Jump

- With jumps J: the strip prices $\int \sigma_t^2 dt + 2\sum(e^J - 1 - J)$
  - but realised variance is $\int \sigma_t^2 dt + \sum J^2$
  - Gap $2E^Q \sum(e^J - 1 - J - J^2/2) \approx E^Q \sum J^3/3$: negative for crashes, so the strip understates expected realised variance Carr & Wu (2009)
- Merton model of Section 6 (one jump a year, mean -8%)
  - strip 17.82 vs $\sqrt{E^Q[\mathrm{QV}]} = 18.00$ vol points, -2.0% in variance
  - QV: quadratic variation, the realised variance including jumps

Back

Truncation and Discretisation of the Bitcoin Index

- Only a finite set of strikes is quoted Jiang & Tian (2005)
  - truncating the strip biases the index down
  - a coarse grid biases it up
- The Bitcoin 30-day index from the same two expiries, in three valid variants:
  - from quoted prices: 40.03
  - from SVI prices at the same strikes: 39.99
  - from SVI on a dense grid inside the quoted range: 39.77
- The gap of about 0.26 points between the three variants reflects smoothing and discretisation
  - all use the same quoted range: truncation is not quantified here
  - the omitted OTM contributions are positive, so they would not lower the index towards DVOL (1.9 points below)
  - quote selection and methodology remain
- Extending SVI beyond the quotes ($k \in [-4, 4]$) is not a fourth estimate
  - the left wing of the 17-day expiry has slope $2.30 > 2$ (Lee's bound)
  - so the extrapolated prices do not come from a valid density
  - the integral gives 128.47, a value with no economic meaning

Back

Encompassing and Out-of-Sample Tests

- HAR-RV Corsi (2009): realised variance averaged over the last 1, 5 and 22 trading days
  - daily, weekly, monthly regressors
  - encompassing regression in logs Busch, Christensen & Nielsen (2011)
  - The target is $RV_{t,t+21}$ over the next 21 days: 22 is a look-back window, 21 the forecast horizon
  - VIX slope 0.90 (0.056); R^2: VIX 0.570, HAR 0.474, both 0.575
- Out of sample from 2000 (6 697 days)
  - coefficients re-estimated every 21 days on targets already observed
  - QLIKE ranks forecasts correctly with a noisy RV if RV is a conditionally unbiased proxy of the variance Patton (2011)
  - Loss ratio VIX^2 / HAR: 0.94, Diebold–Mariano $t = -0.93$ ($p = 0.35$) Diebold & Mariano (1995); VIX corrected by the Mincer–Zarnowitz regression: 0.95 ($p = 0.35$)
  - level forecasts from logs use $\exp(\hat{m} + \hat{s}^2/2)$, a log-normal correction, not a general identity
  - with smearing (mean of $e^{\hat{u}}$ on the training data): 0.95 ($p = 0.39$) and 0.94 ($p = 0.29$)
- Targets overlap by 20 of 21 days
  - the Diebold–Mariano variance is HAC with 21 lags
  - instrumental variables by just-identified generalized method of moments (GMM) with a HAC covariance

Back

Does the Ex-Ante VRP Predict Returns? Overlapping-Horizon Inference

Horizon	$\hat{b}$	t_{NW}	t_{HH}	$t_{Hodrick}$	p_{boot}	R^2 (%)
$h = 1$	0.07	0.33	0.36	0.41	0.35	0.2
$h = 3$	0.09	0.69	0.67	1.33	0.07	1.2
$h = 6$	0.03	0.49	0.48	0.65	0.26	0.3
$h = 12$	−0.02	−0.64	−0.79	−0.54	0.69	0.2

- Monthly 1990–2026: $\frac{12}{h}\sum_{j=1}^{h} r^e_{t+j} = a + b\,\mathrm{VRP}_t + u_t$ Bollerslev, Tauchen & Zhou (2009)
  - $\mathrm{VRP}_t = VIX_t^2/12 - RV_t$, ex-ante proxy with past RV_t as the forecast
  - r^e_{t+j}: S&P 500 excess return of month $t + j$, h: horizon in months
- $VIX_t^2/12$ uses the month-end VIX
  - RV_t is the sum of squared daily log returns (%) within month t

Back

VRP: Inference with Overlapping Horizons

- If one-month errors are serially uncorrelated, overlap makes u_t an MA($h-1$)
  - a moving average of order $h-1$
  - Newey–West (t_{NW}, h lags), Hansen–Hodrick (t_{HH}) Hansen & Hodrick (1980), Hodrick 1B ($t_{Hodrick}$) Hodrick (1992): under the null its size is closer to the nominal level
  - In a given sample $t_{Hodrick}$ can still exceed t_{NW}: at $h=3$, 1.33 vs 0.69
  - These bandwidths are choices, so check their sensitivity
- p_{boot}: one-sided, under $b=0$ Stambaugh (1999)
  - VRP simulated as AR(1), a first-order autoregression ($\hat{\phi}=0.13$)
  - residuals resampled jointly
- On the full 1990–2026 sample no horizon is significant
  - on 1990–2007 (Section 11) $h=3$ is significant

Back

Case Study, Filters and Forecasting Protocol

- Option data: Cboe for 1990–1995 and 2020–2021, OptionMetrics for 1996–2019
- Filters: standard options only (no weekly or quarterly), bid > 0
  - the OTM option at each strike
  - 7 to 550 days, at least 10 strikes per expiry
- Returns: S&P 500 with dividends minus the Fama–French risk-free rate
  - up to March 2021
- Out of sample (Table 8): monthly forecasts against the expanding post-1926 market mean
  - 60-month initial window
  - $R^2_{OS} = 1 - \sum \varepsilon_t^2 / \sum \nu_t^2$, ε_t and ν_t the errors of the forecast and of the benchmark; Diebold–Mariano tests Diebold & Mariano (1995)

Back

The Discovery Loop, Reusable Prompts (1/2)

- **Literature**:
  - `"List peer-reviewed tests of VRP return predictability since 2009: DOI, sample, horizon, table, headline estimate."`
  - Search with Semantic Scholar, Google Scholar or Elicit
- **Hypotheses**:
  - `"Give three competing explanations of VRP predictability and the testable implication of each for the horizon profile, for 2008-2026 and for non-US markets."`
- **Data and code**:
  - `"Write Python that builds VIX squared/12 and monthly realised variance from the daily VIX and S&P 500 closes of the course data and reproduces Table 2 of Bollerslev, Tauchen and Zhou (2009), h = 3, with Hodrick (1992) standard errors."`
  - Any assistant (Claude, ChatGPT, Gemini, GitHub Copilot)

Back

The Discovery Loop, Reusable Prompts (2/2)

- **Robustness**, fixed before the results:
  - `"Draft a pre-analysis plan: horizons, out-of-sample start date, benchmark, test statistic and the multiple-testing correction across horizons."`
  - Placebo: VRP lagged 12 months or block-shuffled
  - Subsamples
  - Optional: 5-minute realised variance of SPY (from 2020), if you have intraday data, which the course data lack
- **Critique**:
  - `"Act as a hostile referee: why could this out-of-sample R2 be spurious or biased, and which test settles each objection?"`
  - Expect: the checks of the main slide, plus the bias with persistent regressors
- **Report**: every horizon and specification tried, in one table
  - with the plan dated before the results
  - AI use logged

Back

Acronyms used in this material (1/2)

- **0DTE**: Zero Days To Expiry (options)
- **AI**: Artificial Intelligence
- **AR**: AutoRegressive model
- **ATM**: At-The-Money: strike close to the forward, $k = \ln(K/F) \approx 0$
- **BTC**: bitcoin, currency code of the cryptocurrency
- **COVID-19**: COronaVIrus Disease 2019
- **CRR**: Cox–Ross–Rubinstein (binomial tree)
- **DVOL**: Deribit Volatility Index
- **ETF**: Exchange-Traded Fund
- **GARCH**: Generalised AutoRegressive Conditional Heteroskedasticity
- **GBM**: Geometric Brownian Motion
- **GJR**: Glosten–Jagannathan–Runkle (asymmetric GARCH)
- **GJR-GARCH**: Glosten–Jagannathan–Runkle GARCH
- **GMM**: Generalized Method of Moments
- **HAC**: Heteroskedasticity and Autocorrelation Consistent (standard errors)
- **HAR**: Heterogeneous AutoRegressive (model)
- **HAR-RV**: Heterogeneous AutoRegressive model of Realised Volatility
- **ITM**: In-The-Money: a put with $K > F$ or a call with $K < F$
- **IV**: Implied Volatility (mark IV; label of the Bitcoin term-structure chart), not instrumental variables
- **MA**: Moving Average

Acronyms used in this material (2/2)

- **NBER**: National Bureau of Economic Research
- **OLS**: Ordinary Least Squares
- **OTM**: Out-of-The-Money: a put with $K < F$ or a call with $K > F$
- **QLIKE**: Quasi-LIKElihood (loss function)
- **QV**: Quadratic Variation
- **RV**: Realised Variance
- **SE**: Standard Error
- **SSVI**: Surface SVI (Gatheral–Jacquier parameterisation of the whole surface)
- **SVI**: Stochastic Volatility Inspired (parameterisation)
- **SVIX**: Simple VIX: index of simple variance swaps (Martin, 2017)
- **US**: United States
- **USD**: US dollar, international currency code
- **UTC**: Coordinated Universal Time (Temps Universel Coordonné)
- **VAR**: Vector AutoRegression
- **VIX**: Cboe Volatility Index
- **VIX3M**: Cboe 3-Month Volatility Index
- **VIX9D**: Cboe 9-Day Volatility Index
- **VRP**: Variance Risk Premium
- **VVIX**: Cboe VIX of VIX Index
- **XFG**: Quantlet prefix of Applied Quantitative Finance (Härdle, Chen & Overbeck)

References (1/5)

- Aït-Sahalia, Y., Duarte, J. (2003). Nonparametric option pricing under shape restrictions. *Journal of Econometrics*, 116(1–2), 9–47.
- Aït-Sahalia, Y., Lo, A.W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance*, 53(2), 499–547.
- Alexander, C., Imeraj, A. (2020). The Bitcoin VIX and its variance risk premium. *Journal of Alternative Investments*, 23(4), 84–109.
- Amihud, Y., Hurvich, C.M. (2004). Predictive regressions: a reduced-bias estimation method. *Journal of Financial and Quantitative Analysis*, 39(4), 813–841.
- Andersen, T.G., Fusari, N., Todorov, V. (2015). The risk premia embedded in index options. *Journal of Financial Economics*, 117(3), 558–584.
- 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.
- Bali, T.G., Beckmeyer, H., Mörke, M., Weigert, F. (2023). Option return predictability with machine learning and big data. *Review of Financial Studies*, 36(9), 3548–3602.
- Bandi, F.M., Fusari, N., Renò, R. (2023). 0DTE option pricing. SSRN Working Paper 4503344.
- Bates, D.S. (1991). The crash of '87: was it expected? The evidence from options markets. *Journal of Finance*, 46(3), 1009–1044.
- Beckmeyer, H., Branger, N., Gayda, L. (2023). Retail traders love 0DTE options... but should they?. SSRN Working Paper 4404704.
- Bertsimas, D., Kogan, L., Lo, A.W. (2000). When is time continuous?. *Journal of Financial Economics*, 55(2), 173–204.

References (2/5)

- 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.
- Britten-Jones, M., Neuberger, A. (2000). Option prices, implied price processes, and stochastic volatility. *Journal of Finance*, 55(2), 839–866.
- Brogaard, J., Han, J., Won, P.Y. (2023). Does 0DTE options trading increase volatility?. SSRN Working Paper 4426358.
- Busch, T., Christensen, B.J., Nielsen, M.Ø. (2011). The role of implied volatility in forecasting future realized volatility and jumps in foreign exchange, stock, and bond markets. *Journal of Econometrics*, 160(1), 48–57.
- Campbell, J.Y., Thompson, S.B. (2008). Predicting excess stock returns out of sample: can anything beat the historical average?. *Review of Financial Studies*, 21(4), 1509–1531.
- 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.
- Cboe Global Markets (2023). The evolution of same day options trading, by S. Doar. Cboe Insights, 3 August 2023.
- Chabi-Yo, F., Loudis, J. (2020). The conditional expected market return. *Journal of Financial Economics*, 137(3), 752–786.
- Christensen, B.J., Prabhala, N.R. (1998). The relation between implied and realized volatility. *Journal of Financial Economics*, 50(2), 125–150.
- Christie, A.A. (1982). The stochastic behavior of common stock variances: value, leverage and interest rate effects. *Journal of Financial Economics*, 10(4), 407–432.

References (3/5)

- Čížek, P., Härdle, W.K., Weron, R. (Eds.) (2011). Statistical Tools for Finance and Insurance, 2nd ed. Springer.
- Clark, T.E., West, K.D. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, 138(1), 291–311.
- 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.
- Delbaen, F., Schachermayer, W. (1994). A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*, 300(1), 463–520.
- Demeterfi, K., Derman, E., Kamal, M., Zou, J. (1999). A guide to volatility and variance swaps. *Journal of Derivatives*, 6(4), 9–32.
- Deribit. DVOL: Bitcoin volatility index. Statistics page.
- Diebold, F.X., Mariano, R.S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263.
- Dupire, B. (1994). Pricing with a smile. *Risk*, 7(1), 18–20.
- Franke, J., Härdle, W.K., Hafner, C.M. (2019). Statistics of Financial Markets: An Introduction, 5th ed. Springer.
- Gârleanu, N., Pedersen, L.H., Poteshman, A.M. (2009). Demand-based option pricing. *Review of Financial Studies*, 22(10), 4259–4299.
- Gatheral, J. (2006). The Volatility Surface: A Practitioner's Guide. Wiley.
- Gatheral, J., Jacquier, A. (2014). Arbitrage-free SVI volatility surfaces. *Quantitative Finance*, 14(1), 59–71.
- Goyal, A., Welch, I., Zafirov, A. (2024). A comprehensive 2022 look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 37(11), 3490–3557.
- Grith, M., Härdle, W.K., Park, J. (2013). Shape invariant modeling of pricing kernels and risk aversion. *Journal of Financial Econometrics*, 11(2), 370–399.

References (4/5)

- 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.
- Härdle, W.K., Chen, C.Y.-H., Overbeck, L. (Eds.) (2017). Applied Quantitative Finance, 3rd ed. Springer.
- Harrison, J.M., Pliska, S.R. (1981). Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and their Applications*, 11(3), 215–260.
- 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.
- Hodrick, R.J. (1992). Dividend yields and expected stock returns: alternative procedures for inference and measurement. *Review of Financial Studies*, 5(3), 357–386.
- Hull, J., White, A. (1987). The pricing of options on assets with stochastic volatilities. *Journal of Finance*, 42(2), 281–300.
- Jackwerth, J.C. (2000). Recovering risk aversion from option prices and realized returns. *Review of Financial Studies*, 13(2), 433–451.
- Jackwerth, J.C., Rubinstein, M. (1996). Recovering probability distributions from option prices. *Journal of Finance*, 51(5), 1611–1631.
- Jiang, G.J., Tian, Y.S. (2005). The model-free implied volatility and its information content. *Review of Financial Studies*, 18(4), 1305–1342.
- Kodde, D.A., Palm, F.C. (1986). Wald criteria for jointly testing equality and inequality restrictions. *Econometrica*, 54(5), 1243–1248.
- Lee, R.W. (2004). The moment formula for implied volatility at extreme strikes. *Mathematical Finance*, 14(3), 469–480.
- Lochstoer, L.A., Muir, T. (2022). Volatility expectations and returns. *Journal of Finance*, 77(2), 1055–1096.
- MacKenzie, D., Millo, Y. (2003). Constructing a market, performing theory: the historical sociology of a financial derivatives exchange. *American Journal of Sociology*, 109(1), 107–145.

References (5/5)

- 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.
- Merton, R.C. (1973). Theory of rational option pricing. *Bell Journal of Economics and Management Science*, 4(1), 141–183.
- Merton, R.C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1–2), 125–144.
- 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., Mazurencu-Marinescu, M. (2012). Modelling stock market crashes: the case of Bucharest Stock Exchange. *Procedia – Social and Behavioral Sciences*, 58, 533–542.
- 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.
- Rubinstein, M. (1994). Implied binomial trees. *Journal of Finance*, 49(3), 771–818.
- Stambaugh, R.F. (1999). Predictive regressions. *Journal of Financial Economics*, 54(3), 375–421.
- Welch, I., Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4), 1455–1508.
- Whaley, R.E. (2000). The investor fear gauge. *Journal of Portfolio Management*, 26(3), 12–17.
- Zulfiqar, N., Gulzar, S. (2021). Implied volatility estimation of bitcoin options and the stylized facts of option pricing. *Financial Innovation*, 7, 67.