Volatility

BUSS386. Futures and Options

Professor Ji-Woong Chung

Lecture Outline

- Volatility
    - Does the BSM predict market option price?
    - Volatility Smile/Smirk
    - Volatility Surface

Volatility

Volatility

- In the Black–Scholes–Merton (BSM) model, volatility is the only **unobservable** input:
$$V(S, K, T, r, \sigma)$$
- The model assumes that σ is a *known and constant* parameter.
- **In reality, volatility violates both assumptions:**
  - **Not constant (empirical facts):**
    - It changes over time (stochastic or time-varying volatility).
    - It differs across strikes and maturities (volatility smile/surface).
  - **Not directly observable (conceptual fact):** Future volatility cannot be measured; implied volatility must be inferred from option prices.
- Because volatility is unobservable and time-varying, market participants form different expectations about it—resulting in substantial **dispersion** in implied volatilities across the market.

Stylized Facts about Volatility

- **Volatility is not constant.** It fluctuates significantly over time.
- **Volatility clusters.** High-volatility periods tend to be followed by more high volatility, and low-volatility periods tend to persist as well.
- **Volatility is mean reverting.** Extremely high or low volatility eventually moves back toward long-run average levels.
- **Volatility is asymmetric.** In equity markets, volatility typically rises when prices fall and (often) declines when prices rise — a pattern known as the *leverage effect* or *volatility–return asymmetry*.

Stylized Facts about Volatility (cont'd)

(Today we will focus on the following cross-sectional features of volatility:)

- **Volatility smile/skew.** Implied volatility exhibits a systematic pattern across strikes — the "smile" or "skew" — indicating that options with different moneyness levels embed different volatility expectations.

- **Term structure of volatility.** Implied volatility varies systematically with maturity, reflecting how the market's expectation of future uncertainty changes over different horizons.

Estimating/Forecasting Volatility

Historical Volatility

- Historical (or *realized*) volatility uses past price data to estimate how much the asset typically moves.

- First compute log returns from past prices:

$$R_{t-k} = \ln\left(\frac{S_{t-k}}{S_{t-k-1}}\right), \quad k = 1, \dots, K$$

- The volatility estimate is the **standard deviation** of these returns, scaled (annualized) by h:

$$\hat{\sigma}_{\text{hist}} = \sqrt{h \cdot \frac{\sum_{k=1}^{K} R_{t-k}^2}{K}}$$

where:

  - $h = 252$ for daily data,
  - $h = 52$ for weekly data.

Estimating/Forecasting Volatility

Historical Volatility

- This method assumes:
  - recent past behavior is informative about the near future,
  - the volatility during the lookback window is roughly constant.

	A	B	C	D
1	**Date**	**Price**	**log return**	**Squared**
2	t	102	0.985%	9.707E-05
3	t-1	101	4.041%	0.0016329
4	t-2	97	-2.041%	0.0004165
5	t-3	99	-1.005%	0.000101
6	t-4	100		
7		**average**	0.495%	0.0005619
8		**annualize (*255)**	126.242%	0.14328
9				
10			**volatility**	37.852%

Estimating/Forecasting Volatility

Exponentially Weighted Moving Average (EWMA)

- Historical volatility treats all past returns equally. EWMA improves this by giving **more weight to recent data** and gradually downweighting older observations.

- Compute log returns as before.

- Apply exponentially decreasing weights:

$$w^k, \quad \text{where } 0 < w < 1$$

  For example, RiskMetrics (J.P. Morgan) uses:

$$w = 0.94 \text{ (daily)}.$$

- The volatility estimate is:

$$\hat{\sigma}_{\text{EWMA}} = \sqrt{h \cdot \frac{\sum_{k=0}^{K} w^k R_{t-k}^2}{\sum_{k=0}^{K} w^k}}$$

Estimating/Forecasting Volatility

Exponentially Weighted Moving Average (EWMA)

- Intuition:
  - recent shocks matter more for predicting tomorrow's volatility,
  - helps capture volatility clustering.

	A	B	C	D	E	F
1	Date	Price	log return	Squared	weight (0.9)	w*Squared
2	t	102	0.985%	9.707E-05	1	9.70677E-05
3	t-1	101	4.041%	0.0016329	0.9	0.001469638
4	t-2	97	-2.041%	0.0004165	0.81	0.000337383
5	t-3	99	-1.005%	0.000101	0.729	7.36357E-05
6	t-4	100				
7		average	0.495%	sum	3.439	0.001977724
8		annualize (*255)	126.242%		weighted average	0.000575087
9					annualize (*255)	0.146647175
10						
11					volatility	38.295%

Estimating/Forecasting Volatility

GARCH (Generalized Autoregressive Conditional Heteroskedasticity)

- GARCH models capture a key empirical fact: **volatility today depends on volatility yesterday and recent shocks.**

- The simplest model, GARCH(1,1), has two equations:
  - **Return equation:**
  $$r_t = \mu + \epsilon_t, \quad \epsilon_t \sim (0, \sigma_t^2)$$
  - **Variance equation:**
  $$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

  where:
    - $\alpha =$ reaction to new shocks (volatility jumps),
    - $\beta =$ persistence of past volatility,
    - $\omega =$ long-run average level.

- Typical empirical values:

$$\alpha \approx 0.05\text{–}0.10, \qquad \beta \approx 0.90\text{–}0.95$$

$$\alpha + \beta < 1 \quad \text{(ensures mean reversion)}$$

Estimating/Forecasting Volatility

GARCH (Generalized Autoregressive Conditional Heteroskedasticity)

- Key intuition:
  - A big return shock today raises tomorrow's volatility.
  - High volatility tends to persist (volatility clustering).
  - Much more realistic than the "constant volatility" assumption of BSM.

	A	B	C	D	E	F
1	Date	Price	log return	Squared*255	lagged	Conditional variance
2	t	102	0.985%	0.024752275	0.41639735	5.884%
3	t-1	101	4.041%	0.416397351	0.10621312	3.674%
4	t-2	97	-2.041%	0.106213121	0.02575736	2.942%
5	t-3	99	-1.005%	0.025757359		2.576%
6	t-4	100				

$$\omega = 0.495\%, \ \alpha = 0.05, \text{ and } \beta = 0.9.$$

Comparing Volatility Estimation Methods

Method	Inputs	Formula / Model	Intuition	Pros / Cons
Historical	K past log returns R_t	$\hat{\sigma}^2 = \frac{1}{K}\sum_{k=1}^{K} R_k^2$	Volatility is just the sample variance of past returns.	+ Very simple, easy to compute – Treats all days equally; ignores clustering and changing volatility
EWMA	K past log returns R_t and decay factor w	$\hat{\sigma}^2_{\text{EWMA}} = \frac{\sum_{k=0}^{K} w^k R_{t-k}^2}{\sum_{k=0}^{K} w^k}$	Recent returns get more weight; old data gradually "forgotten."	+ Captures volatility clustering + Still easy to implement – Choice of w is somewhat ad hoc; still not a full model
GARCH(1,1)	Past shocks ϵ_{t-1} and past variance σ^2_{t-1}	Return: $r_t = \mu + \epsilon_t$ Variance: $\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2$	Volatility today depends on yesterday's shock and yesterday's volatility.	+ Flexible; matches clustering and mean reversion + Widely used in practice – More complex; parameters must be estimated statistically

Practical Issues in Estimating Volatility from Past Prices

- In the Black–Scholes world, the underlying follows a lognormal diffusion with a **constant** volatility.

  If this were true:

  - more data $\Rightarrow$ better estimates,
  - higher-frequency data $\Rightarrow$ more precision.

- **In reality, markets do not behave this cleanly.** Several practical issues arise:
  - **1. Choice of observation interval (daily? weekly? intraday?).**
    - Higher frequency contains more information, but ultra-high frequency data introduces noise: bid–ask bounce, microstructure effects, stale quotes on some exchanges.
    - In practice, **daily data is a reasonable compromise** for most asset classes.
  - **2. Should we estimate the mean return?**
    - The sample mean is tiny relative to daily volatility. Using it adds noise without improving accuracy.
    - **Standard practice: assume mean = 0.**

Practical Issues in Estimating Volatility from Past Prices

- **3. How much past data should we include?**
  - More data increases statistical precision. But markets change—volatility in 2008 or 2020 may not reflect volatility today.
  - Use long enough windows to get stability, but avoid mixing very different economic regimes.
- **4. What to do with outliers (e.g., 1987 crash, COVID-19 spike)?**
  - Outliers can dominate historical volatility calculations.
  - No universally correct answer: remove them? downweight them? keep them to capture tail risk?
  - Best practice: **examine sensitivity and use judgment.**

Forecasting Volatility: Practical Suggestions

- **Evaluate forecasting methods out-of-sample.** The real test is not how well a model fits history, but how it performs on future data.
- **All volatility forecasting methods are noisy.** Even the best approaches produce large forecast errors, especially at short horizons.
- **Use zero mean for returns.** Subtracting the sample mean often makes volatility estimates worse, as the sample mean is itself very imprecise.
- **Simple models often work surprisingly well.** Historical volatility over a long window performs about as well as many complex alternatives, and is more robust to structural change.

Forecasting Volatility: Practical Suggestions

- **Forecasts are more reliable for longer horizons.** One-day-ahead volatility is very difficult to predict, but 3–6 month volatility is typically easier due to mean reversion.
- **GARCH models are good for short horizons.** They adapt quickly to new information but require:
  - a long data history for accurate parameter estimation,
  - careful checking of model assumptions.

How Practitioners Estimate Volatility

- **Practitioners do not use historical or GARCH volatility to price options.**
  - These methods are used mainly for *risk management* (VaR, ES, stress testing).
- **Option markets use implied volatility (IV).**
- **The implied volatility surface is the key pricing input.**
- **Workflow on an options desk:**
  1. Observe liquid option prices $\rightarrow$ compute IVs.
  2. Fit a smooth implied volatility surface.
  3. Price and hedge exotic/illiquid options using the surface.
  4. Use historical/EWMA/GARCH only for *risk*, not pricing.

Implied Volatility

- **Definition.** Implied volatility (IV) is the value of σ that makes an option pricing model (e.g., Black–Scholes–Merton) match the option price observed in the market:

$$C_{\text{model}}(S_0, K, T, r, \sigma_{\text{impl}}) = C_{\text{mkt}}.$$

  - IV is often considered the **best estimate of volatility available**, because it incorporates all information and beliefs reflected in market prices (including expectations, risk premia, and supply–demand effects).
  - There is a **one-to-one mapping** between price and IV: higher uncertainty $\Rightarrow$ higher option price $\Rightarrow$ higher implied volatility.
  - In many markets (FX, equity index options), traders quote *volatility*, not option prices. The volatility surface is the central object used in pricing and hedging.

Implied Volatility

- **IV is the same for European calls and puts with the same S_0, K, T.**
  - In the market, put–call parity must also hold (ignoring small arbitrage bounds):

$$C_{\text{mkt}} - P_{\text{mkt}} = S_0 - e^{-rT}K.$$

  - The BSM option prices should satisfy the parity.

$$C_{\text{BS}} - P_{\text{BS}} = S_0 - e^{-rT}K.$$

    Subtract:

$$C_{\text{BS}} - C_{\text{mkt}} \;=\; P_{\text{BS}} - P_{\text{mkt}}.$$

  - Suppose the IV for the put is 20%. I.e., $P_{\text{BS}} = P_{\text{mkt}}$ when IV=20%. Hence, $C_{\text{BS}} - C_{\text{mkt}}$ should be zero as well when IV=20% is used.
  - Therefore: One implied volatility per strike–maturity pair, not one per option.

The VIX Index

Trading Volatility

- The CBOE publishes indices of implied volatility. The most important is the **VIX**, which measures the market's expectation of **annualized 30-day S&P 500 volatility**, extracted from a broad set of SPX options.

- **VIX Futures** (introduced in 2004)
  - Contract size: $1{,}000 \times$ VIX index level
  - Minimum tick: 0.01 (1 volatility basis point)
  - Futures allow investors to trade *expected future volatility*, not the VIX today.

- **VIX Options** (introduced in 2006)
  - European-style settlement in cash.
  - Call payoff:
  $$100 \times \max(VIX_T - K, 0)$$
  - Note: VIX options are options on **VIX futures**, not on the spot VIX.

Trading Volatility

- The VIX is like a "temperature gauge": you can observe it at each moment, but you **cannot store** it or carry it forward. Therefore, **no cost-of-carry pricing model** (like for commodities). VIX derivatives are priced based on **expectations** and the implied volatility surface.

- **Example: VIX futures:** Buy April VIX futures at 18.5, sell at 19.3:

$$\text{Profit} = (19.3 - 18.5) \times 1,000 = \$800.$$

Variance Swaps and Volatility Swapsx

- **Variance swaps**: A forward contract on the *realized variance* of returns over a future period.
  - Strike is quoted as a variance (e.g. $0.20^2 = 0.04$).
  - Realized variance =annualized average of R_t^2, assuming mean return $= 0$.
  - Payoff = Notional $\times$ (Realized Var − Strike Var).
  - Traders who expect higher future volatility **buy variance**.

- **Example:** Expect high volatility through Dec. 31, strike = 20% vol = 0.04 variance. If realized variance ends at 0.06, the payoff = notional $\times$ (0.06 – 0.04).

- **Volatility swaps**: Similar idea, but payoff is based on *realized volatility* instead of variance.
  - More intuitive (investors think in vol), but harder to value.
  - Cannot be perfectly hedged using options $\rightarrow$ less commonly traded.
  - Variance swaps are preferred because realized variance can be **replicated exactly** using a strip of OTM options.

BSM option price = Market option price?

- Let's compare BSM options prices to market prices at a point in time
  - The data is from May 3, 2007
  - The S&P 500 index was at $S = 1502.39$
  - The one-month risk-free rate was at $r = 4.713\%(c.c.)$
  - The dividend yield on the S&P 500 was about $q = 1.91\%$
  - Using the previous 3 months of returns:
    $\sigma = \sqrt{\frac{1}{63}\sum_{i=1}^{63}(r_{t-i} - \bar{r})} \times \sqrt{252} = 12.35\%$

Comparing BSM predictions to market prices

SPX (S&P 500 INDEX)	Today	cc rate	div yield	volatility
1502.39	5/3/2007	0.04713	0.0191	0.1236

				CALLS			PUTS		
Maturity	Time to T	Strike	Moneyness K/S	Mkt Price	B/S	BSC/Mkt	Mkt Price	B/S	BSP/Mkt
6/15/2007	0.12	1430	0.952	83.9	80.12	0.955	6.2	3.19	0.514
6/15/2007	0.12	1435	0.955	79.4	75.74	0.954	6.7	3.78	0.564
6/15/2007	0.12	1440	0.958	75	71.44	0.953	7.3	4.46	0.610
6/15/2007	0.12	1445	0.962	70.6	67.24	0.952	7.9	5.23	0.662
6/15/2007	0.12	1450	0.965	66.3	63.14	0.952	8.7	6.10	0.701
6/15/2007	0.12	1455	0.968	62.1	59.15	0.952	9.3	7.08	0.761
6/15/2007	0.12	1460	0.972	57.9	55.27	0.955	10.1	8.17	0.809
6/15/2007	0.12	1465	0.975	53.8	51.52	0.958	10.9	9.39	0.862
6/15/2007	0.12	1470	0.978	49.8	47.89	0.962	11.9	10.74	0.902
6/15/2007	0.12	1475	0.982	45.9	44.40	0.967	12.6	12.22	0.970
6/15/2007	0.12	1480	0.985	42.1	41.05	0.975	14.1	13.84	0.982
6/15/2007	0.12	1485	0.988	38.4	37.84	0.986	15.4	15.61	1.014
6/15/2007	0.12	1490	0.992	34.8	34.79	1.000	17.05	17.52	1.028
6/15/2007	0.12	1495	0.995	31.4	31.88	1.015	18.55	19.59	1.056
6/15/2007	0.12	1500	0.998	28.05	29.13	1.039	20.35	21.82	1.072
6/15/2007	0.12	1505	1.002	24.55	26.54	1.081	21.95	24.19	1.102
6/15/2007	0.12	1510	1.005	22	24.10	1.095	24	26.73	1.114
6/15/2007	0.12	1515	1.008	19.3	21.81	1.130	26.2	29.41	1.123
6/15/2007	0.12	1520	1.012	16.6	19.68	1.186	28.6	32.25	1.128
6/15/2007	0.12	1525	1.015	14.8	17.70	1.196	31.2	35.24	1.130
6/15/2007	0.12	1530	1.018	12.3	15.86	1.290	34	38.38	1.129
6/15/2007	0.12	1535	1.022	10.3	14.17	1.376	37	41.66	1.126
6/15/2007	0.12	1540	1.025	8.6	12.61	1.467	40.3	45.07	1.118
6/15/2007	0.12	1545	1.028	7.05	11.19	1.587	43.7	48.62	1.113
6/15/2007	0.12	1550	1.032	5.95	9.89	1.663	47.4	52.30	1.103
6/15/2007	0.12	1555	1.035	4.5	8.72	1.937	51.2	56.09	1.096
6/15/2007	0.12	1560	1.038	3.7	7.65	2.068	55.2	60.00	1.087
6/15/2007	0.12	1565	1.042	2.9	6.69	2.308	59.4	64.01	1.078
6/15/2007	0.12	1570	1.045	2.325	5.83	2.509	63.7	68.13	1.070
6/15/2007	0.12	1575	1.048	1.9	5.07	2.667	68.2	72.33	1.061

When K/S is low (OTM Puts & ITM Calls), IV is greater.

Volatility Smirk/Skew

- Note that the BSM assumes that volatility is constant.
- Volatility Smirk/Skew: The volatility used to price a low-strike-price option (i.e., a OTM put or a ITM call) is significantly higher than that used to price a high-strike-price option (i.e., a ITM put or a OTM call).

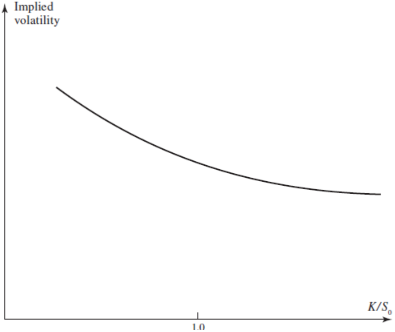

Volatility Skew/Smirk: Why?

- **1. Negative correlation between equity prices and volatility (the leverage effect).** When stock prices fall, volatility tends to rise $\rightarrow$ OTM puts become more valuable $\rightarrow$ higher IV. When prices rise, volatility tends to fall $\rightarrow$ OTM calls become less valuable $\rightarrow$ lower IV.
- **2. Market crash of 1987 fundamentally changed option pricing.** Before 1987, equity options showed little smile or skew. After Black Monday (a 22% one-day drop), traders demanded:
  - more protection against extreme downside risk,
  - more compensation for selling crash insurance.

  This permanently steepened the volatility skew.
- **3. Why don't ITM calls show similar high volatility?** ITM calls are often used as low-capital substitutes for stock holdings:
  - $\Delta \approx 1 \rightarrow$ behaves like stock,
  - lower cash outlay $\rightarrow$ leveraged exposure.

Volatility Smile

- **FX options typically exhibit a symmetric volatility smile** rather than the steep downside skew seen in equity options.
- **Why a smile in FX markets?**
  - **FX rates do not have a natural "crash direction."** Equity indices have a built-in downside risk (stocks can collapse). Currency pairs can move sharply in *either* direction.

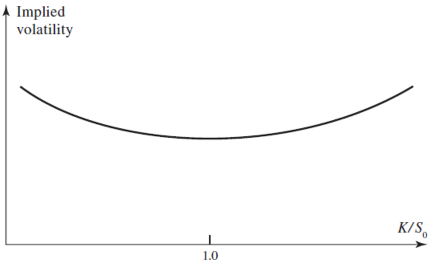

Volatility Prediction: Implied Volatility

- **The volatility smile reveals a limitation of the Black–Scholes model.** If BSM were correct, all options would share the *same* volatility. In reality, implied volatility varies by strike and maturity.
  - This means implied volatility contains market information about future uncertainty, but it is not an *unbiased* forecast of future realized volatility.
  - IV reflects both expectations *and* risk premia demanded by investors.
- **More sophisticated models can generate a smile.** Examples include: stochastic volatility models (Heston), local volatility models (Dupire), models with jumps or stochastic interest rates (see Appendix for discussion).
- **In practice: market makers use "practitioner Black–Scholes."**
  - They keep the BSM formula but replace the constant volatility assumption.
  - Each option gets its own volatility input taken from the **implied volatility surface**.
  - This ensures that model prices match observable market prices.

A Simple Practitioners' Approach to Pricing with the Smile

1. **Record the implied volatility smile.** Traders observe recent IVs across strikes and construct the current **implied volatility curve** (smile or skew).
2. **Assume the smile shape will persist.**
3. **Estimate the at-the-money (ATM) IV.** ATM implied volatility is typically the most reliable and liquidly observed.
4. **Adjust IV for each strike using the smile.** For a given moneyness K/S_0, use the smile curve to obtain an appropriate IV instead of using a single "constant" volatility.
5. **Plug this adjusted IV into the Black–Scholes formula.** This is known as **"practitioner Black–Scholes"**: BSM model + strike-specific IV.

Example: Practitioner Smile-Based Pricing

- Current stock price: $S_0 = 100$
- Time to maturity: $T = 0.25$ years (3 months)
- Risk-free rate: $r = 2\%$
- ATM implied volatility observed from the market: $\sigma_{ATM} = 20\%$
- Recent smile shows:
  - IV for 90 strike $= 25\%$
  - IV for 100 strike $= 20\%$
  - IV for 110 strike $= 23\%$

Goal: Price a 110-strike call.

1. ATM IV is 20%, but smile shows OTM calls use IV $= 23\%$.
2. Use $\sigma = 0.23$ in Black–Scholes:

$$C_{110} = 2.02 \quad \text{(using IV} = 23\%\text{, not } 20\%\text{)}$$

3. BSM with constant volatility (20%) would give:

$$C_{110}^{BSM} = 1.40 \quad \text{(underpricing the call)}$$

A Simple Practitioners' Approach to Pricing with the Smile

- **Use the adjusted IV surface to compute Greeks** (Δ, Γ, Θ, vega, etc.) that match market pricing behavior.
- **Market making: fit a smooth curve through the smile.**
  - Often fit a quadratic, spline, or exponential curve to IV vs. strike.
  - Options with IV above (below) the curve $\rightarrow$ **overpriced** (**underpriced**).
  - Market makers quote tighter or wider spreads based on this deviation.

A More Sophisticated Practitioners' Approach

- **Goal:** Compare the relative value of many options at once (different strikes and maturities).
- **Problem:** Raw option prices cannot be compared directly.
  - Different strikes $\rightarrow$ different intrinsic values.
  - Different maturities $\rightarrow$ different time value, interest rates, and uncertainty.
  - Price differences alone do not indicate whether one option is "expensive" or "cheap."
- **Solution:** Convert all option prices into *implied volatilities*. IV normalizes price differences by expressing everything in volatility units:

$$\text{IV} = \text{"how much volatility is the market implying?"}$$

This allows apples-to-apples comparison across all strikes and maturities.

A More Sophisticated Practitioners' Approach

Volatility Surface

- Once IVs are computed across all strikes and maturities, practitioners arrange them into a **volatility surface**:

$$\sigma = \sigma\left(\frac{K}{S_0}, T\right)$$

- This surface summarizes:
  - the skew/smile (variation across strike), and
  - the term structure (variation across maturity).

- The volatility surface becomes the central tool for:
  - pricing OTC and exotic options,
  - computing Greeks,
  - fitting models (local vol, stochastic vol),
  - relative value trading (identifying cheap/expensive options).

A More Sophisticated Practitioners' Approach

Volatility Surface

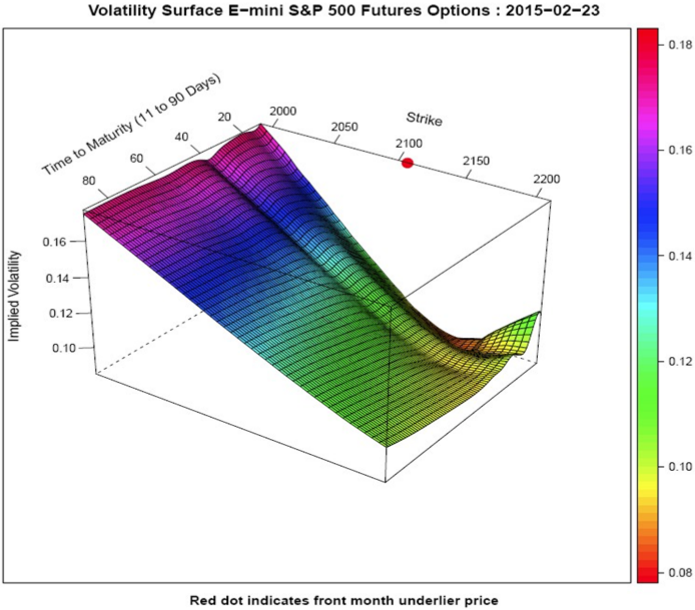

A More Sophisticated Practitioners' Approach

Examples

- **To value a new option, locate its position on the volatility surface.**

$$\sigma^* = \sigma\left(\frac{K}{S_0}, T\right)$$

Table 20.2 Volatility surface.

	K/S_0				
	0.90	*0.95*	*1.00*	*1.05*	*1.10*
1 month	14.2	13.0	12.0	13.1	14.5
3 month	14.0	13.0	12.0	13.1	14.2
6 month	14.1	13.3	12.5	13.4	14.3
1 year	14.7	14.0	13.5	14.0	14.8
2 year	15.0	14.4	14.0	14.5	15.1
5 year	14.8	14.6	14.4	14.7	15.0

A More Sophisticated Practitioners' Approach

Examples

- **Example 1: 9-month, $K/S_0 = 1.05$ (OTM call).**
  - Table gives IVs: 13.4% (0.5 year) and 14.0% (1.0 year).
  - Interpolate along maturity to get: $\sigma^* \approx 13.7\%$
  - Use σ^* in BSM or a binomial tree to compute the option value.
- **Example 2: 1.5-year, $K/S_0 = 0.925$ (ITM put).**
  - Strike dimension $\rightarrow$ interpolate across K/S_0.
  - Maturity dimension $\rightarrow$ interpolate across T.
  - Apply **bilinear interpolation** (2D). Result:$\sigma^* \approx 14.525\%$
  - Use this IV to price the option.
- **Alternative approach: fit a regression-based IV surface.**
  - Estimate:

$$\hat{\sigma} = a + b\left(\frac{K}{S_0}\right) + cT + d\left(\frac{K}{S_0}\right)^2 + e\left(\frac{K}{S_0}T\right) + fT^2.$$

  - Smooths noisy market IVs and avoids jumps in the surface.

Summary

- Despite its inaccuracies BSM serves as a useful benchmark.
  - Gives decent approximation to prices close to the money.
- It also works reasonably well to hedge options positions against changes in stock prices using delta or delta-gamma hedging.
- Models have been proposed to correct some of the shortcomings.
  - Stochastic volatility
  - Jumps
  - Fat tails
- All of these models are consistent with the idea that OTM puts are expensive relative to BSM prices because investors seeking protection from large losses (e.g., jumps down) must pay a higher (insurance) premium

Appendix: Volatility Estimation

Numerical Example: Prices and Log Returns

Suppose we observe the following daily closing prices for a stock:

Day	0	1	2	3	4	5
S_t	100	102	101	105	103	106

Daily log returns are

$$R_t = \ln\left(\frac{S_t}{S_{t-1}}\right)$$

t	1	2	3	4	5
R_t	0.01980	-0.00985	0.03884	-0.01923	0.02871
R_t^2	0.000392	0.000097	0.001509	0.000370	0.000824

Numerical Example: Historical Volatility

Using the 5 daily log returns

$$R_1, \ldots, R_5,$$

the (daily) historical variance estimator is

$$\hat{\sigma}^2_{\text{daily}} = \frac{1}{K} \sum_{t=1}^{K} R_t^2, \quad K = 5.$$

$$\sum_{t=1}^{5} R_t^2 \approx 0.003192 \quad \Rightarrow \quad \hat{\sigma}^2_{\text{daily}} = \frac{0.003192}{5} \approx 0.000638$$

Daily historical volatility:

$$\hat{\sigma}_{\text{daily}} = \sqrt{0.000638} \approx 0.02527 \quad (2.53\% \text{ per day})$$

Annualizing with $h = 252$ trading days:

$$\hat{\sigma}_{\text{annual}} = \sqrt{252}\, \hat{\sigma}_{\text{daily}} \approx 0.4011 \quad (40.1\% \text{ per year})$$

Numerical Example: EWMA Volatility

EWMA gives more weight to recent returns. Let $w = 0.94$ and use the same 5 returns.

Index the returns from most recent to oldest:

k	0	1	2	3	4
return	R_5	R_4	R_3	R_2	R_1
value	0.02871	-0.01923	0.03884	-0.00985	0.01980

Weights:

$$w^0 = 1.0000,\ w^1 = 0.94,\ w^2 \approx 0.8836,\ w^3 \approx 0.8306,\ w^4 \approx 0.7807$$

$$\sum_{k=0}^{4} w^k \approx 4.4349$$

EWMA daily variance:

$$\hat{\sigma}^2_{\text{daily, EWMA}} = \frac{\sum_{k=0}^{4} w^k R_{t-k}^2}{\sum_{k=0}^{4} w^k} \approx \frac{0.002892}{4.4349} \approx 0.000652$$

Daily EWMA volatility:

$$\hat{\sigma}_{\text{daily, EWMA}} = \sqrt{0.000652} \approx 0.02553 \quad (2.55\% \text{ per day})$$

Annualized:

$$\hat{\sigma}_{\text{annual, EWMA}} = \sqrt{252}\,\hat{\sigma}_{\text{daily, EWMA}} \approx 0.4053 \quad (40.5\% \text{ per year})$$

Numerical Example: GARCH(1,1) Volatility

Consider a simple GARCH(1,1) model:

$$r_t = \mu + \epsilon_t, \quad \epsilon_t \sim (0, \sigma_t^2),$$

$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2.$$

For this example, set:

$$\mu = 0, \quad \alpha = 0.05, \quad \beta = 0.94, \quad \alpha + \beta = 0.99 < 1.$$

We choose ω so that the long-run (unconditional) variance matches our historical estimate $\hat{\sigma}^2_{\text{daily}} \approx 0.000638$:

$$\sigma^2_\infty = \frac{\omega}{1 - \alpha - \beta} \quad \Rightarrow \quad \omega \approx (1 - 0.99) \times 0.000638 \approx 6.38 \times 10^{-6}.$$

Numerical Example: GARCH(1,1) Volatility (cont'd)

Start with the long-run variance as the initial value:

$$\sigma_1^2 = \sigma_\infty^2 \approx 0.000638.$$

Shocks are just returns (since $\mu = 0$):

$$\epsilon_t = R_t.$$

Step 1: Compute σ_2^2

$$\epsilon_1 = R_1 = 0.01980, \quad \epsilon_1^2 \approx 0.000392$$

$$\sigma_2^2 = \omega + \alpha\epsilon_1^2 + \beta\sigma_1^2 \approx 0.00000638 + 0.05(0.000392) + 0.94(0.000638) \approx 0.000626$$

$$\sigma_2 = \sqrt{0.000626} \approx 0.02502 \quad \text{(2.50\% per day)}$$

Numerical Example: GARCH(1,1) Volatility (cont'd)

Step 2: Compute σ_3^2

$$\epsilon_2 = R_2 = -0.00985, \quad \epsilon_2^2 \approx 0.000097$$

$$\sigma_3^2 = \omega + \alpha\epsilon_2^2 + \beta\sigma_2^2 \approx 0.00000638 + 0.05(0.000097) + 0.94(0.000626) \approx 0.000600$$

$$\sigma_3 = \sqrt{0.000600} \approx 0.02449 \quad (2.45\% \text{ per day})$$

Interpretation:

- A large shock today (ϵ_t^2 big) raises tomorrow's variance σ_{t+1}^2.
- In the absence of large shocks, σ_t^2 slowly drifts back toward the long-run level.

Appendix: Volatility Models
(Chapter 27)

Local Volatility Models

- **Idea:** The Black–Scholes model assumes volatility is constant. Local volatility models relax this by allowing volatility to depend on the **current stock price** and **time**:

$$dS_t = \mu S_t dt \; + \; \sigma(S_t, t)\, S_t dW_t.$$

- One popular example is the **Constant Elasticity of Variance (CEV)** model:

$$dS_t = \mu S_t dt + (\sigma S_t^{\gamma-1}) S_t dW_t.$$

- **Interpretation of γ:**
  - $\gamma = 1$: reduces to the Black–Scholes model (constant volatility).
  - $\gamma < 1$:
    - When S_t falls, S_t^{γ} increases relative to S_t,
    - $\Rightarrow$ volatility rises at low prices,
    - $\Rightarrow$ OTM puts become more valuable,
    - $\Rightarrow$ **volatility smirk (equity skew)** appears.
  - $\gamma > 1$: volatility rises with price (sometimes seen in futures/options on commodities).
- **Takeaway:** Local volatility models can generate a smile or smirk because volatility changes depending on where the stock price is.

Stochastic Volatility Models

- **Idea:** Volatility itself moves randomly over time, rather than being constant. This helps explain why options imply different volatilities at different strikes.
- **The Heston Model** is a leading example:

$$dS_t = \mu S_t dt + \sqrt{v_t}\, S_t dW_t,$$
$$dv_t = \theta(\omega - v_t)dt + \xi\sqrt{v_t}\, dB_t.$$

- **Meaning of parameters:**
  - ω: long-run average variance.
  - θ: speed of mean reversion (how fast variance returns to ω).
  - ξ: "vol of vol" – how much variance itself fluctuates.
  - ρ: correlation between stock returns (dW_t) and volatility shocks (dB_t).
- **Why Heston explains the smirk:**
  - If $\rho < 0$ (empirically true for equities): bad market moves $\rightarrow$ higher volatility.
  - Higher volatility $\rightarrow$ higher crash probability.
  - This makes OTM put options relatively expensive.
  - $\Rightarrow$ **downward-sloping volatility skew**.
- **Intuition:** A drop in price increases volatility, which increases the chance of an even bigger drop.

Jumps in Stock Prices

- **Idea:** Stock prices occasionally experience sudden large jumps (e.g., 1987 crash, 2020 COVID crash). The Black–Scholes model cannot capture this.
- A jump–diffusion model adds a jump component to price movements:

$$dS_t = \mu S_t dt + \sigma S_t dW_t + J(Q)\, S_t\, dP_t.$$

- **Interpretation:**
  - dP_t is usually 0, but equals 1 with a small probability $\rightarrow$ a jump occurs.
  - $J(Q)$ is the jump size (can be random or fixed).
  - If $J(Q) < 0$, the jump is downward $\rightarrow$ a crash.
- **Impact on options:**
  - If downward jumps are possible, tail risk increases.
  - OTM puts become much more valuable.
  - $\Rightarrow$ **steeper volatility smirk**.
- **Downside jump risk = expensive insurance.**
- **Pricing jumps is harder**, because Black–Scholes cannot be used directly. More advanced numerical methods are required (Monte Carlo, Fourier transforms, etc).

Implied Tree Models

- **Idea:** Normally, we start with a model (e.g., binomial tree) $\rightarrow$ compute option prices. With **implied trees**, we reverse the process:

$$\text{Use observed option prices} \;\Rightarrow\; \text{infer the stock price tree.}$$

- Once calibrated, the implied tree:
  - matches market option prices exactly,
  - produces a local volatility surface implicitly,
  - can be used to price other options consistently.

Implied Tree Models

- **Example:** Given:

$$S_0 = 1502.39,\ K = 1500,\ \sigma = 12.36\%,\ r = 4.713\%,\ \delta = 1.91\%,\ T = 0.12.$$

  Compute initial binomial parameters:

  - $u = e^{\sigma\sqrt{T}} = 1.0437, \quad d = 1/u = 0.9581.$
  - Risk-neutral $p = \dfrac{e^{(r-\delta)T} - d}{u - d} = 0.5286.$
  - BSM/binomial price: $c = 28.394.$
  - Market price: $c^{mkt} = 20.35$ (model overprices).

- **Implied-tree step:** Adjust σ (and therefore u and d) until model $=$ market:

$$\sigma = 8.24\% \quad (\text{new } p = 0.5446)$$

  Now the tree is calibrated $\rightarrow$ can be used to price other options.

Other Modern Volatility Models

- **Beyond local volatility, stochastic volatility, and jump models,** modern quantitative finance uses several advanced models to better match the observed volatility smile and surface.
- **SABR Model (Hagan, Kumar, Lesniewski, Woodward, 2002)**
  - Widely used in interest-rate and FX markets.
  - Models both the asset price and its volatility as stochastic processes.
  - Flexible enough to generate smiles, skews, and term structure effects.
  - Provides simple formulas that traders can implement quickly.
- **Rough Volatility Models (Gatheral, Jaisson, Rosenbaum, 2018)**
  - Based on the empirical finding that volatility moves "roughly," showing long memory and very jagged paths.
  - Captures the fine structure of volatility better than classical models.
  - Produces realistic short-term smiles and accurate VIX dynamics.
- **Why these models matter in practice:**
  - They fit market implied volatility surfaces more accurately.
  - They improve pricing of exotic options and risk management.
  - They help traders understand how the smile evolves over time.