Volatility Smiles, Skews, and Surfaces

Athanasios Sakkas, AUEB

Learning objectives

By the end of the lecture, students should be able to:

1. Define implied volatility, volatility smile, volatility skew, and volatility surface.
2. Explain why put-call parity implies equal implied volatility for matched European calls and puts.
3. Connect the shape of a smile/skew to the market-implied risk-neutral distribution.
4. Contrast typical FX and equity option volatility patterns.
5. Interpret a volatility surface across strike and maturity.

Black–Scholes Assumptions

- Stock prices follow a geometric Brownian motion:

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

- Constant volatility and interest rates
- No arbitrage
- Continuous trading
- Lognormal stock prices
- European exercise

Black–Scholes Pricing Formula

European call:

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

European put:

$$P = K e^{-rT} N(-d_2) - S_0 N(-d_1)$$

where

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

$$d_2 = d_1 - \sigma\sqrt{T}$$

Put–Call Parity

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

Applications:

- Detect arbitrage opportunities
- Construct synthetic positions
- Infer forward prices

Option Greeks

Delta

$$\Delta_{call} = N(d_1)$$

Gamma

$$\Gamma = \frac{N'(d_1)}{S_0\sigma\sqrt{T}}$$

Vega

$$\text{Vega} = S_0\sqrt{T}N'(d_1)$$

Theta

$$\Theta_{call} = -\frac{S_0N'(d_1)\sigma}{2\sqrt{T}} - rKe^{-rT}N(d_2)$$

Risk-Neutral Valuation

Under the risk-neutral measure:

$$\frac{dS_t}{S_t} = r\,dt + \sigma\,dW_t^{\mathbb{Q}}$$

Option value:

$$V_0 = e^{-rT}\mathbb{E}^{\mathbb{Q}}[\text{Payoff}]$$

Call payoff:

$$C_0 = e^{-rT}\mathbb{E}^{\mathbb{Q}}[\max(S_T - K, 0)]$$

Implied Volatility

Implied volatility solves:

$$C_{\mathrm{mkt}} = C_{\mathrm{BSM}}(\sigma_{\mathrm{imp}}(K, T))$$

Important implications:

- Volatility smiles
- Volatility skews
- Volatility surfaces

Observed smiles imply:

- Non-normal returns
- Fat tails
- Stochastic volatility
- Jump risk

Motivation: what Black–Scholes predicts

In the Black–Scholes–Merton framework, volatility is a single constant input.

Model implication
For the same underlying asset and maturity, every option should imply the same volatility, regardless of strike.

In real markets, options with different strikes usually imply different volatilities. This difference is the empirical starting point for volatility smiles and skews.

Implied volatility

Definition
The implied volatility of an option is the volatility input that makes the model price equal to the observed market price.

$$V_{\mathrm{BSM}}(S_0, K, T, r, q, \sigma_{\mathrm{imp}}(K, T)) = V_{\mathrm{mkt}}.$$

Implied volatility is not directly observed; it is backed out from option prices.

Volatility smile and volatility skew

- A **volatility smile** plots implied volatility against strike or moneyness for a fixed maturity.
- A **smile** is often U-shaped.
- A **skew** or **smirk** is asymmetric, usually with higher implied volatility on one side.

Put-call parity

For European options on a dividend-yield-paying asset:

$$p + S_0 e^{-qT} = c + Ke^{-rT}.$$

- c: call price; p: put price
- S_0: current underlying price
- K: strike; T: maturity
- r: domestic risk-free rate; q: dividend yield or foreign risk-free rate in FX

Key point
Put-call parity is a no-arbitrage relationship. It does not depend on the asset-price distribution being lognormal.

Why matched calls and puts have the same implied volatility

Put-call parity holds for both market prices and Black–Scholes–Merton prices:

$$p_{\text{mkt}} + S_0 e^{-qT} = c_{\text{mkt}} + Ke^{-rT},$$

$$p_{\text{BSM}} + S_0 e^{-qT} = c_{\text{BSM}} + Ke^{-rT}.$$

Subtracting gives:

$$p_{\text{mkt}} - p_{\text{BSM}} = c_{\text{mkt}} - c_{\text{BSM}}.$$

If the volatility input makes the put model price equal to the put market price, it also makes the matched call model price equal to the call market price. Therefore:

$$\sigma_{\mathrm{imp}}^{put}(K,T) = \sigma_{\mathrm{imp}}^{call}(K,T).$$

Important implication

Same smile for matched European calls and puts
For the same strike and maturity, European calls and puts imply the same volatility.

Therefore, the volatility smile is a property of the market's pricing of the underlying risk, not a separate call-only or put-only phenomenon. For American

options, early exercise features complicate the exact parity argument, but the approximation is often useful.

From smile to implied distribution

The volatility smile tells us how market option prices differ from the Black–Scholes lognormal benchmark.

- High implied volatility for low strikes means relatively expensive downside protection.
- High implied volatility for high strikes means relatively expensive upside exposure.
- A symmetric smile suggests both tails are priced as heavier than lognormal.
- An equity skew often suggests a heavier left tail than the lognormal benchmark.

FX volatility smile

Typical FX smiles are relatively symmetric:

- At-the-money options often have lower implied volatility.
- Deep in-the-money and out-of-the-money options often have higher implied volatility.
- Interpretation: market-implied exchange-rate distribution has fatter tails than lognormal.

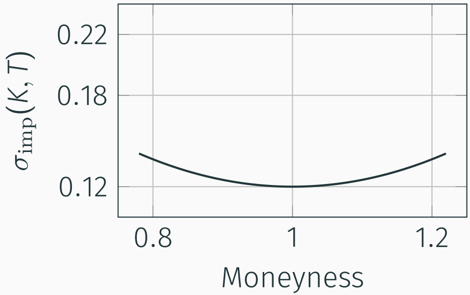

Equity index options often display a downward-sloping skew:

- Low-strike puts are expensive.
- Downside crash insurance commands a premium.
- The implied distribution has a heavier left tail than the lognormal benchmark.

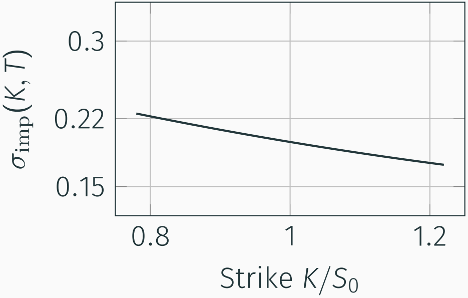

Why do smiles exist?

Possible explanations include:

1. **Jumps**: asset prices may jump rather than follow continuous diffusion.
2. **Stochastic volatility**: volatility itself varies over time.
3. **Leverage effect**: for equities, falling prices often coincide with rising volatility.
4. **Crash risk and demand pressure**: investors buy downside protection, raising put prices.
5. **Market frictions**: liquidity, margin, and funding constraints can affect option prices.

Volatility term structure

A smile is a cross-section across strikes for one maturity. A term structure looks across maturities.

Empirical tendency
When current volatility is high, the volatility term structure is often downward sloping. When current volatility is low, it is often upward sloping.

This reflects mean reversion in volatility expectations.

Volatility surface

A volatility surface combines strike and maturity:

$$\sigma_{\mathrm{imp}}(K, T) = f(K, T).$$

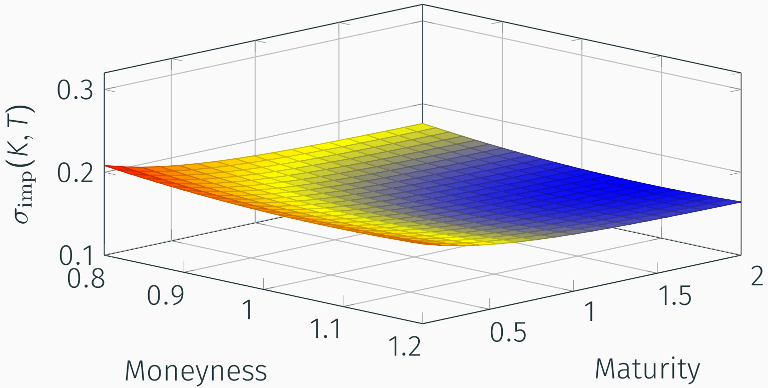

Interpreting a volatility surface table

	$K/S_0 = 0.90$	0.95	1.00	1.05	1.10
1 month	22.0%	19.0%	17.0%	16.5%	16.3%
3 months	21.0%	18.5%	17.2%	16.8%	16.7%
1 year	19.0%	18.0%	17.5%	17.2%	17.0%

Questions to ask:

- Is skew stronger for short or long maturities?
- Is the surface smooth across strikes and maturities?

Alternative smile characterizations

Traders often quote smiles using market conventions rather than raw strike:

- **Moneyness**: K/S_0 or Δ.
- **At-the-money volatility**: the center of the smile.
- **Risk reversal**: difference between call and put implied volatilities at comparable deltas.

Summary

- A volatility smile/skew shows how implied volatility varies with strike for a fixed maturity.
- Put-call parity implies matched European calls and puts have the same implied volatility.
- Smile shape is linked to the market-implied distribution of the underlying asset.
- FX smiles are often more symmetric; equity smiles are often left-skewed.
- A volatility surface extends the smile across maturities and is central to option pricing and risk management.

Lognormal Benchmark vs Implied Distribution

Black–Scholes lognormal terminal distribution

In Black–Scholes–Merton, the underlying follows geometric Brownian motion:

$$\frac{dS_t}{S_t} = (r - q)dt + \sigma dW_t^{\mathbb{Q}}.$$

Therefore the risk-neutral terminal price is lognormally distributed:

$$\ln(S_T) \sim N\left(\ln(S_0) + (r - q - \tfrac{1}{2}\sigma^2)T,\ \sigma^2 T\right).$$

Implication
A single constant volatility produces one lognormal distribution and a flat implied-volatility curve:

$$\sigma_{\mathrm{imp}}(K, T) = \sigma \quad \text{for all strikes } K.$$

Market-implied distribution

Market option prices reveal a different distribution from the Black–Scholes lognormal benchmark.

- If low-strike puts are expensive, the market assigns more value to downside states.
- If deep OTM calls and puts are expensive, the market assigns more value to tail states.
- Hence a smile or skew indicates that the implied distribution is not lognormal.

For a continuum of European call prices, the risk-neutral density can be inferred from strike curvature:

$$f_Q(K) = e^{rT}\frac{\partial^2 C(K,T)}{\partial K^2}.$$

This is the Breeden–Litzenberger relation.

Graph: lognormal vs implied distribution

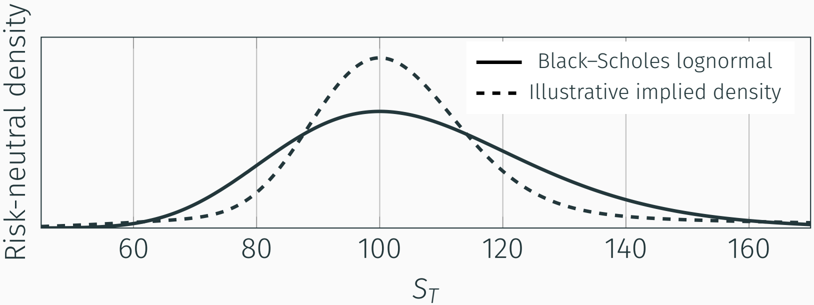

- Black–Scholes: constant σ implies a lognormal terminal density and flat $\sigma_{\mathrm{imp}}(K, T)$.
- Typical smile intuition: the implied density can be more peaked near the same centre and fatter in the tails.
- Equity-index skew adds relatively more probability/state-price weight to the left tail; in this simplified graph we keep the modal point fixed.

How the smile changes the distribution

Flat volatility

- One volatility input
- Lognormal terminal price
- Thin tails relative to crisis data
- No strike dependence

Smile or skew

- Volatility depends on strike and maturity
- Non-lognormal implied distribution
- Fatter tails and/or asymmetry
- Market prices downside and upside states differently

$$\text{Smile/skew in } \sigma_{\mathrm{imp}}(K,T) \quad \Longleftrightarrow \quad \text{non-lognormal implied } f_Q(S_T).$$

Equity-index implied distributions

For equity indices, the smile is usually a negative skew:

$$\sigma_{\mathrm{imp}}(K_{low}, T) > \sigma_{\mathrm{imp}}(K_{ATM}, T) > \sigma_{\mathrm{imp}}(K_{high}, T).$$

Interpretation:

- Investors demand crash insurance.
- OTM puts become expensive.
- The implied distribution has a heavier left tail than the lognormal benchmark.
- The distribution is negatively skewed under the risk-neutral measure.

FX implied distributions

FX options often display a more symmetric smile:

$$\sigma_{\mathrm{imp}}(K_{low}, T) > \sigma_{\mathrm{imp}}(K_{ATM}, T), \qquad \sigma_{\mathrm{imp}}(K_{high}, T) > \sigma_{\mathrm{imp}}(K_{ATM}, T).$$

Interpretation:

- Large appreciations and depreciations are both priced as possible.
- The implied distribution may have two fat tails relative to lognormal.
- The direction of skew depends on the currency pair and quotation convention.

Important warning: physical vs risk-neutral distributions

The implied distribution is not the same as the real-world forecasting distribution.

- **Physical distribution** P: actual probabilities investors expect.
- **Risk-neutral distribution** Q: probabilities adjusted for risk premia and state prices.
- Option prices identify Q, not directly P.

$$C_0 = e^{-rT}\mathbb{E}^{\mathbb{Q}}[(S_T - K)^+].$$

A high implied probability of extreme negative outcomes may reflect not only expectations of a market crash, but also the fact that protection against crashes becomes especially valuable during bad times.