Department of Economics & Finance

Chair of Mathematical Finance

Volatility in Financial Markets: From Black-Scholes to the Volatility Surface

SUPERVISOR:

Prof. Sara Biagini

CANDIDATE:

Lorenzo Damiano

Identification Number 281991

Academic Year 2024-2025

Table of Contents

Introduction

Volatility plays a fundamental role in options markets, given its role as a driver of option prices. Yet the assumption of constant volatility in the Black-Scholes model has been found to be insufficient to account for the nuances of real-world data, including implied volatility smile and skew. This mismatch has motivated the creation of increasingly sophisticated and realistic models for explaining the dynamics of volatility in time, and across strikes and maturities. More specifically, the investigation of local volatility, which redefines volatility as a function of stock price and time-to-maturity, has become a crucial element in the analysis of vanilla and exotic options.

This thesis consists of three chapters and attempts to give a general yet focused exposition of the role of volatility. Chapter 1 lays the theoretical groundwork for the Black-Scholes model by introducing the necessary mathematical and probability concepts, namely stochastic processes, Brownian motion, and Ito calculus. A good understanding of these notions is needed to appropriately appreciate both the strengths and weaknesses of Black-Scholes.

Chapter 2 deals with the analysis of the patterns of implied volatility observed in financial markets, with particular emphasis placed on the volatility smile, skew, and construction of a complete volatility surface. Starting with empirical imperfections of the constant volatility assumption made by Black-Scholes, this chapter explores how implied volatility varies with strike price as well as time to maturity. The conversation has specific emphasis on the mathematical constraints required to guarantee a well-functioning, arbitrage-free volatility surface.

Chapter 3 presents local volatility models as a possible way to align market-observed option prices with no-arbitrage dynamics. The presentation starts with a discrete setup based on binomial trees, that are used to introduce the Derman-Kani model, which enhances the basic binomial model by employing a volatility that is both strike and time-dependent. Then the chapter moves to a continuous setup with Dupire's forward partial differential equation, enabling the derivation of the local volatility surface from market vanilla option prices. The aim for such models is to replicate the observed volatility surface with analytical tractability and internal consistency.

1. The Mathematics of the Black-Scholes Model

This chapter presents the mathematical underpinnings of the Black-Scholes model, a fundamental component of contemporary quantitative finance. Its elegance and ability to be analytically solvable have made it an effective and widely used instrument for option pricing. Meanwhile, some of its simplifying assumptions, notably the constant volatility assumption, instigate the need for more adaptable methods. Knowing the fundamentals of this model is an essential step in learning about the volatility surface and the advantages of local volatility models.

1.1 Probability Spaces

A fundamental concept in the theory of stochastic processes is the one of probability space. It is defined as a triple $(\Omega, \mathcal{F}, \mathbb{P})$, where:

- Ω is a nonempty set, called the *sample space*. It is useful to imagine it as the set of all possible outcomes of the random experiment we are modelling.
- $\mathcal{F}$ is a collection of subsets of Ω, called *σ-algebra*, with the following properties:
    1. The empty set $\emptyset$ is in $\mathcal{F}$
    2. If a set A belongs to $\mathcal{F}$, its complement A^c also belongs to $\mathcal{F}$ (closure under complementation)
    3. If a sequence of sets $A_1, A_2, \ldots, A_n$ belongs to $\mathcal{F}$, their union $\bigcup_{n=1}^{\infty} A_n$ also belongs to $\mathcal{F}$ (closure under countable unions)
- $\mathbb{P}$ is a function $\mathbb{P}: \mathcal{F} \to [0,1]$, called the probability measure. It assigns to every set $A \in \mathcal{F}$ a number between zero and one, defined as the probability of A, or $\mathbb{P}(A)$. Additionally, its properties require that:
    1. $P(\Omega) = 1$
    2. For any sequence of disjointed sets $A_1, A_2, \ldots, A_n$ in $\mathcal{F}$, $\mathbb{P}\left(\bigcup_{n=1}^{\infty} A_n\right) = \sum_{n=1}^{\infty} \mathbb{P}(A_n)$ (countable additivity)

This structure provides a very useful framework on which to base the Brownian motion that stocks are assumed to have in the Black-Scholes model. For now, however, it is useful to visualize the probability space in the context of a simpler event, such as the roll of a die. For a fair six-sided die, the sample space is Ω = {1, 2, 3, 4, 5, 6}, representing all possible outcomes of a single roll.

The σ-algebra $\mathcal{F}$ consists of subsets of Ω that satisfy the aforementioned properties. For instance, we may define the subsets A_1= {1}, meaning "the die lands on 1", and A_2 = {2, 3, 4, 5, 6}, meaning "the die lands on any number but 1". We can then construct the $\mathcal{F}_1$ = $\{\emptyset, \Omega, A_1, A_2\}$, for which it is trivial to prove that it is a σ-algebra.

The probability measure ℙ can then be used to assign probabilities to the events in $\mathcal{F}_1$. For a fair die, these would be $\mathbb{P}(\emptyset) = 0, \mathbb{P}(\Omega) = 1,\ \mathbb{P}(A_1) = \frac{1}{6}, \mathbb{P}(A_2) = \frac{5}{6}$

1.2 Random Variables

Another fundamental concept in the theory of stochastic processes is the random variable. Given a probability space (Ω, $\mathcal{F}$, ℙ), we call random variable a measurable function X: Ω ⟶ ℝ, with the property that, for every Borel set $B \in \mathbb{R}$, the pre-image $X^{-1}(B) = \{\omega \in \Omega \mid X(\omega) \in B\}$ belongs to the σ-algebra $\mathcal{F}$. A Borel set on the real line is any set that belongs to the Borel σ-algebra B(ℝ), which is the smallest σ-algebra containing all open intervals in ℝ. This means it includes all open and closed sets, countable unions and intersections of them, and more complex sets built through these operations. By requiring the random variable to be measurable with respect to the Borel σ-algebra, we ensure that we can assign well-defined probabilities to events involving the values of X, such as $\mathbb{P}\{a < X \leq b\}$. Continuing from the dice roll example, we could imagine a random value that takes the following values: X(ω)= $\begin{cases} 1\ if\ \omega\ is\ odd \\ 0\ if\ \omega\ is\ even \end{cases}$. It is also possible to obtain the probabilities associated to the values the function can take, for example: $\mathbb{P}(X = 1) = \mathbb{P}(A_{\omega=1} \cup A_{\omega=3} \cup A_{\omega=5}) = \frac{1}{2}$

1.3 Stochastic Processes

A stochastic process is a collection of random variables indexed by time, typically written as $\{X_t\}_{t \in T}$, where each X_t is a random variable defined on a common probability space (Ω, $\mathcal{F}$, ℙ). The index T usually represents time and can be discrete (T = ℕ) or continuous (T=[0,∞)). Each random variable X_t:

$\Omega \longrightarrow \mathbb{R}$ (or possibly R^d) represents the state of the process at time t, meaning that a stochastic process describes how a random system evolves over time. Stochastic processes are used to model phenomena that evolve unpredictably over time, from stock prices to particle diffusion.

When observing the results of a stochastic process, it's often necessary to keep track of the information available up to a certain time. In the case of dice rolls, we may for example only be aware of the outcome of one out of a number of rolls. In this case, although we are not aware of the real path ω, we can be sure that a number of sets either could or could not contain it. This leads to the concept of filtration, a family of sub σ -algebras of $\mathcal{F}$ such that $\mathcal{F}_s \subseteq \mathcal{F}_t$, whenever $s \leq t$. The filtration represents the accumulation of information over time: $\mathcal{F}_t$contains all the events observable by time t.

1.4 Random Walk

Before introducing Brownian motion, it is useful to present one of its discrete-time counterparts: the random walk. The random walk is a discrete model of stochastic motion. It is also often considered the building block of more complex stochastic processes, including Brownian motion. A simple and intuitive way to visualize the random walk is by imagining a person taking steps along a straight line and, at each time step, tossing a coin. If it lands on heads, they step to the left, if it lands on tails, they step to the right. Over time, their position evolves in a random manner, depending on the sequence of coin tosses.

We can give a more formal definition as follows. Let $(\xi_i)_{i\in\mathbb{N}}$ be a sequence of independent and identically distributed (i.i.d.) random variables such that: $\xi_i = \begin{cases} 1 \; with \; probability \; p \\ -1 \; with \; probability \; 1-p \end{cases}$

We also define the position of the walk at time n, denoted by S_n, as $\sum_{i=1}^{n} \xi_i$. It is easy to see that, in the case where the coin is fair (p = 1-p = 0.5), each step has expected value $\mathbb{E}[\xi_i] = 0$, and the walk is centered, meaning there is no systematic drift in either direction. The variance of each step is $\mathbb{V}ar[\xi_i] = 1$, and the central limit theorem tells us that, as the number of steps increases, the distribution of approaches a normal distribution.

1.5 Brownian Motion

If we rescale both time and space appropriately, the random walk can be used to approximate a Brownian motion. More precisely, we define a scaled random walk: $W_n(t) = \frac{1}{\sqrt{n}} S_{nt}$, where nt is an integer. We can then obtain a Brownian motion in the limit as $n \to \infty$. Using the framework developed earlier, we are now ready to define Brownian motion rigorously.

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. A standard Brownian motion or Wiener process is a family of random variables $\{W_t\}_{t\geq 0}$ defined on this space with the following properties:

1. $W_0 = 0$
2. For any $0 \leq t_1 \leq t_2 \leq \ldots \leq t_n$ the random variables $W_{t_2} - W_{t_1}, W_{t_3} - W_{t_2}, \ldots, W_{t_n} - W_{t_{n-1}}$ are independent (independent increments)
3. For $s < t$, the increment $W_t - W_s$ is normally distributed with mean 0 and variance $t - s$: $W_t - W_s \sim \mathcal{N}(0, t - s)$
4. Almost surely (meaning with probability 1), the function $t \to W(\omega)$ is continuous in t for every $\omega \in \Omega$

We can also give a definition of filtration for Brownian motion: Given a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ on which a Brownian motion $\{W_t\}_{t\geq 0}$ is defined. A filtration for the Brownian motion is a collection of σ-algebras $\{\mathcal{F}_t\}_{t\geq 0}$ such that:

1. For every $0 \leq s < t$, $\mathcal{F}(s) \subset \mathcal{F}(t)$. That is, the information available at time t includes at least everything that was known at time s.
2. For each $t \geq 0$, the Brownian motion W(t) at time t is $\mathcal{F}(t)$-measurable. In other words, the σ-algebra $\mathcal{F}(t)$ contains all the information necessary to determine the value of $W(t)$ at time t.
3. For $0 \leq t < u$, the increment W(u) - W(t) is independent of $\mathcal{F}(t)$. In other words, the future increments of the Brownian motion after time t are independent of the information contained in the σ-algebra $\mathcal{F}(t)$.

1.6 Ito Calculus

Classical calculus, built on the notion of differentiable functions, falls short when applied to stochastic processes like Brownian motion, whose sample paths are continuous but nowhere differentiable. This limitation led to the development of Ito calculus, a branch of stochastic calculus tailored to the analysis of random processes. Ito calculus provides the mathematical tools necessary to define and work with integrals involving Brownian motion. These integrals form the basis for solving stochastic differential equations (SDEs), which are essential for modeling dynamic systems under uncertainty.

1.6.1 Ito's Lemma

At the heart of Ito calculus is Ito's Lemma, a stochastic version of the chain rule. It allows us to compute the differential of a function of a stochastic process, taking into account the random fluctuations introduced by Brownian motion. This is particularly important because, unlike in classical calculus, second-order terms in stochastic differentials do not disappear and must be carefully handled.

Suppose a stochastic process X_t satisfies the stochastic differential equation:

$$dX_t = \mu_t dt + \sigma_t dW_t \tag{1.6.1.1}$$

Where μ_t and σ_t are deterministic functions of time. Let $f(t, X_t)$ be a function that is once continuously differentiable in t and twice continuously differentiable in x. Then Itô's Lemma states:

$$d\, f(t, X_t) = \left(\frac{\partial f}{\partial t} + \mu_t \frac{\partial f}{\partial x} + \frac{1}{2}\sigma_t^2 \frac{\partial^2 f}{\partial x^2}\right)\mathrm{dt} + \sigma_t \frac{\partial f}{\partial x}\, dW_t \tag{1.6.1.2}$$

This formula is widely used in mathematical finance. For example, applying Itô's Lemma to the logarithm of a geometric Brownian motion allows one to derive the lognormal distribution for the asset price:

$$S_t = S_0 exp\left(\left(\mu - \frac{1}{2}\sigma^2\right) t + \sigma W_t\right) \tag{1.6.1.3}$$

This result is crucial in the derivation of the Black-Scholes formula for option pricing

1.6.2 Ito Integral

The Ito integral formalizes the concept of integrating stochastic processes with respect to Brownian motion. Let $\{W_t\}_{t\geq 0}$ be a standard Brownian motion defined on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$. Suppose $\{\xi_t\}_{t\in[0,T]}$ s an adapted stochastic process that satisfies the integrability condition:

$$\mathbb{E}\left[\int_0^T \xi_t^2 \, dt]\right] < \infty \quad \textbf{(1.6.2.1)}$$

We then call the Ito integral of ξ_t with respects to W_t over the interval $[0, T]$:

$$\int_0^T \xi_t \ dW_t \quad \textbf{(1.6.2.2)}$$

Unlike classical Riemann or Lebesgue integrals, the Ito integral accounts for the randomness of the integrator W_t and is constructed as the L^2-limit of discrete approximations. For a partition $0 = t_0 \leq t_1 \leq t_2 \leq \ldots \leq t_n = T$, the Ito integral is approximated by:

$$\int_0^T \xi_t \ dW_t \ \approx \sum_{i=1}^{n} \xi_{t_{i-1}}(W_{t_i} - W_{t_{i-1}}) \quad \textbf{(1.6.2.3)}$$

where the integrand is evaluated at the left endpoint of each interval to ensure it is $\mathcal{F}_{ti} - 1$ measurable and independent of future Brownian increments.

It can be shown that the Ito integral has the Ito isometry property, which relates the square of the Ito integral to the quadratic variation of W_t:

$$\mathbb{E}\left[\left(\int_0^T \xi_t \ dW_t\right)^2\right] = \mathbb{E}\left[\int_0^T {\xi_t}^2 dt\right] \quad \textbf{(1.6.2.4)}$$

1.7 The Black-Scholes Model

The Black-Scholes model, first developed by Fischer Black and Myron Scholes in 1973 and later developed further by Robert Merton, was a major development in the mathematics of finance. The model built the first theory widely recognized to justify the pricing of European options and thus forms a pillar of modern-day quantitative finance. As we have already noted, the model assumes the price movements of the underlying security follow a geometric Brownian motion and uses the tools of stochastic calculus, specifically the framework based on Ito calculus, to obtain a partial differential equation describing the behavior of option pricing. By solving the equation in the context of appropriate boundary conditions, it becomes possible to obtain a closed-form solution to the valuations of European put and call options.

1.7.1 Assumptions of the Black-Scholes Model

The Black-Scholes model makes a number of assumptions, the most basic of which is the need for the existence of at least one risky asset (the stock) and one risk-free asset (the bond). It is also necessary that:

1. Markets are frictionless (meaning there are no transaction costs for buying or selling)
2. There are no arbitrage opportunities in the market (no opportunities for riskless profit)
3. The risk-free interest rate remains constant over the option's life
4. Agents in the market can buy and sell any amount of the stock or the bond, even fractional
5. The stock does not pay a dividend (although later formulations incorporate this possibility)
6. The options are European (meaning they can only be exercised at expiration)
7. The stock follows a Geometric Brownian motion, with lognormal returns

This factors, if respected, make it possible to perfectly hedge a long position in a stock by selling an option, so that the combined value of the position does not depend on the stock's price.

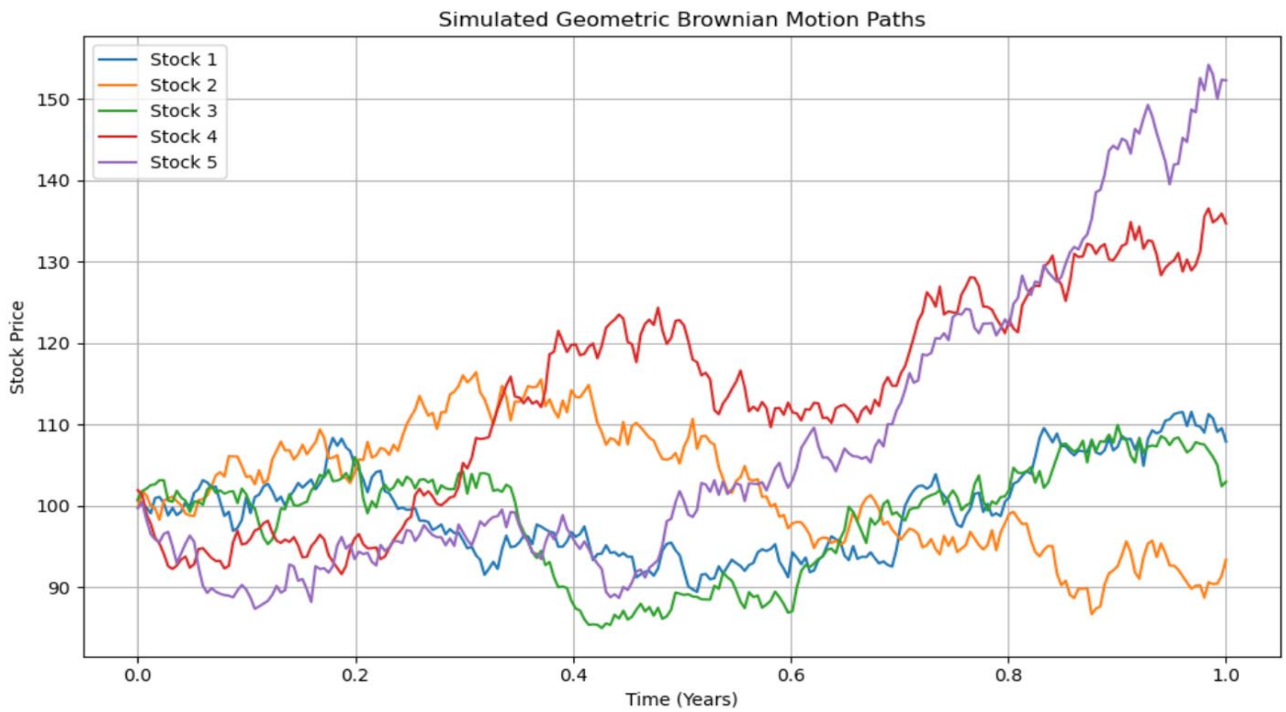

Fig.1: Simulated stock paths following Geometric Brownian motion

1.7.2 Derivation of the Black-Scholes Differential Equation

Let the stock price process be the one we described in **1.6.1.1**: $dS_t = \mu S_t dt + \sigma S_t dW_t$ (for brevity from now on we are going to omit the index t). Let V(t, S) be the price of a generic European option, that depends on the time t and the price of the underlying asset S. In **1.6.1.2**, we already described how we can use Ito's lemma on this type of functions to obtain:

$$dV = \left(\frac{\partial V}{\partial t} + \mu S \frac{\partial V}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2}\right) d\text{t} + \mu S \frac{\partial V}{\partial S} dW \quad \textbf{(1.7.2.1)}$$

Now, we image creating a hedged portfolio Π consisting of one derivative contract -V and a quantity Δ of the underlying asset. The value of this portfolio is

$$\Pi = -V + \Delta S \quad \textbf{(1.7.2.2)}$$

and its change can be expressed as $\mathrm{d}\Pi = -dV + \Delta dS$. We now want to eliminate the Brownian motion component, and to do so first we can choose $\Delta = \frac{\partial V}{\partial S}$, in order to obtain:

$$\mathrm{d}\Pi = -dV + \frac{\partial V}{\partial S} dS. \tag{1.7.2.3}$$

We can then substitute **1.6.1.1** and **1.7.2.1** inside **1.7.2.3** and obtain:

$$\mathrm{d}\Pi = \left(-\frac{\partial V}{\partial S} - \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2}\right) dt \tag{1.7.2.4}$$

We can clearly see how, having lost the stochastic dW component, the portfolio is now by definition considered risk-free. By the model's own assumptions, specifically for arbitrage to be absent, this must mean it earns the same return as other riskless securities. If we set the risk-free interest rate as r, we can represent this as:

$$\mathrm{d}\Pi = \mathrm{r}\Pi\mathrm{dt}. \tag{1.7.2.5}$$

Finally, by substituting **1.7.2.2** and **1.7.2.3** into **1.7.2.5**, we obtain:

$$\left(\frac{\partial V}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2}\right) dt = r\left(V - \frac{\partial V}{\partial S} S\right) dt \tag{1.7.2.6}$$

From which we get the Black-Scholes differential equation:

$$\frac{\partial V}{\partial t} + rS\frac{\partial V}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} - rV = 0 \tag{1.7.2.7}$$

The solutions of this PDE, which give possible arbitrage free prices for an option, will then depend on the boundary conditions set. For a European call option, the key boundary condition is: $V = \max(S - K, 0) \;\; when\ t = T$, while for a European put option it is $V = \max(K - S, 0) \;\; when\ t = T$.

1.7.3 The Black-Scholes Formula

After having outlined the Black-Scholes differential equation and its boundary conditions, we can then proceed to solve it and obtain the Black-Scholes pricing formulas for both put and call options. These are:

$$C = S_0 N(d_1) - Ke^{-rT} N(d_2) \tag{1.7.3.1}$$

For the price of a European call option and

$$P = Ke^{-rT} N(-d_2) - S_0 N(-d_1) \tag{1.7.3.2}$$

For the price of a European put option. In both formulas above, $N(x)$ refers to the cumulative distribution function of the standard normal distribution (meaning the probability that a normally distributed variable is less than or equal to x), while:

$$d_1 = \frac{ln\,(S_0/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}} \tag{1.7.3.3}$$

$$d_2 = d_1 - \sigma\sqrt{T} = \frac{ln\,(S_0/K) + (r - \sigma^2/2)T}{\sigma\sqrt{T}} \tag{1.7.3.4}$$

An important assumption made in the derivation of the Black-Scholes formulas is that all investors are risk-neutral, that is, they do not require additional compensation to bear the risk associated with volatility. This assumption is reflected in the fact that the expected return μ of the underlying asset does not appear in the Black-Scholes partial differential equation. Under the risk-neutral measure, all assets are assumed to grow at the risk-free rate r. As a result, the pricing of derivatives is based on expected payoffs discounted at the risk-free rate, rather than on their real-world expected returns. For example, the price of a European call option, whose expected payoff at maturity is $\hat{\mathbb{E}}[max(S_T - K,\ 0)]$, can be computed as:

$$C = e^{-rT}\hat{\mathbb{E}}[max(S_T - K, 0)] \tag{1.7.3.5}$$

where the expectation is taken under the risk-neutral measure $\hat{\mathbb{E}}$.

With this context, it is also easier to explain the meaning of $N(d_1), N(d_2)$. The term d_2 represents the standardized distance between the logarithmic forward price of the asset and the strike price, adjusted for time and volatility. As such, $N(d_2)$ can be interpreted as the risk-neutral probability that the option will expire in the money.

The meaning of d_1 is a bit more obscure. It can be viewed as capturing the risk-adjusted moneyness of the option under the risk-neutral measure. The function $N(d_1)$ therefore reflects the option's delta, the sensitivity of the option price to small changes in the underlying asset. In this way, $N(d_1)$ can also be loosely interpreted as the risk-neutral probability that the option will be exercised, weighted by the expected exposure to the asset.

1.7.4 The Greeks

Having derived the Black-Scholes pricing formula for European options, it becomes important to understand how option prices respond to changes in their underlying parameters. This is where the Greeks come into play. The Greeks are partial derivatives of the option price with respect to key inputs, which can be calculated from put-call parity.

We have already mentioned *delta*, which expresses the degree to which the option price is affected by a change in the underlying asset price. For a European call option:

$$\Delta = N(d_1) = \frac{\partial C}{\partial S} \qquad \textbf{(1.7.4.1)}$$

Gamma is a second order derivative, expressing the sensitivity of delta to changes in the underlying price:

$$\Gamma = \frac{\partial^2 C}{\partial S^2} \qquad \textbf{(1.7.4.2)}$$

Gamma captures the curvature of the option price with respect to the underlying and is highest for at-the-money options close to expiration. In the Black-Scholes model:

$$\Gamma = \frac{N'(d_1)}{S\sigma\sqrt{T}} \tag{1.7.4.3}$$

A portfolio with high gamma requires more frequent delta adjustments.

Theta captures the sensitivity of the option price to the passage of time:

$$\Theta = \frac{\partial C}{\partial t} \tag{1.7.4.4}$$

It usually takes a negative value, indicating that the value of an option generally decreases as time passes, all else being equal. This phenomenon is referred to as time decay.

Vega measures the sensitivity of the option price to changes in volatility:

$$\upsilon = \frac{\partial C}{\partial \sigma} \tag{1.7.4.5}$$

The mere existence of vega gives us insights into the fact that constant volatility is an unrealistic assumption, something we will explore in further detail in the coming chapters. For now, we can say that vega is always positive for both calls and puts, reflecting the principle that options become more valuable when volatility increases.

2. Volatility and the Volatility Surface

As we have mentioned before, in the Black-Scholes framework, volatility is assumed to be constant over time and across all strike prices. However, empirical observations reveal that implied volatility extracted from market prices varies systematically with both the strike price and the time to maturity. This variation is manifested in phenomena known as volatility smile and skew, which have significant implications for option pricing, hedging, and risk management. In this chapter, we describe these phenomena, with the aim of shifting our focus from the rigid assumptions of Black-Scholes to better capture the dynamic behavior of volatility in financial markets.

2.1 The Volatility Smile and Skew

A major departure from the constant volatility assumption is observed in the phenomenon commonly known as the “volatility smile”. When plotting the implied volatility against the strike price for options with a fixed maturity, one often observes a U-shaped curve rather than the expected flat line. In many cases, on the other hand, the curve is not symmetric, resulting instead in a pronounced downward volatility skew.

Both phenomena capture an important insight into how market participants perceive risk. The volatility smile, mostly seen for currency exchanges, reflects the market's recognition of the fat tails in the underlying’s distribution. That is to say, extreme movements in the underlying asset, whether upward or downward, are of higher magnitude than what is implied by a lognormal distribution. This results in options with strike prices more distant from the underlying’s being assigned higher implied volatilities.

The volatility skew, on the other hand, introduces asymmetry into the smile and is mainly found in equity markets. It typically manifests as higher implied volatilities for options with lower strike prices, compared to those with higher strikes. This phenomenon can be linked to several factors. Historically, when the price of an asset falls, its volatility tends to rise more compared to the inverse, as investors are generally more concerned about large downward moves than upward ones. This is also linked to the fact that jumps, one of the main factors not accounted for when modeling stocks as geometric Brownian motions, are more frequent and more significant to the downside. As a result, out-of-the-money puts tend to be more expensive (in volatility terms) than equivalent out-of-the-money calls.

An additional explanation links this phenomenon with the so-called "leverage effect", which is based on the fundamental accounting identity EV = E + D, where EV is the enterprise value, E is the value of the company's equity and D is the value of its debt. Based on this, Robert Merton's structural credit risk model reimagined equity as the value of a call option on the company's assets, with strike price equal to the face value of debt D. Given that equity holders are residual claimants, at time T, $E = \max\{EV - D, 0\}$. Assuming the risk of debt is negligible, we can approximate the equity's volatility as $\sigma_E \approx \frac{EV}{E}\sigma_{EV}$. By construction $\frac{EV}{E} \geq 1$, so this relationship shows that equity volatility increases as the value of equity falls, holding asset volatility constant. This inverse relationship between stock price and volatility is the essence of the leverage effect: as a firm's value drops, equity becomes more leveraged and more sensitive to fluctuations in enterprise value, leading to rising observed volatility in the stock.

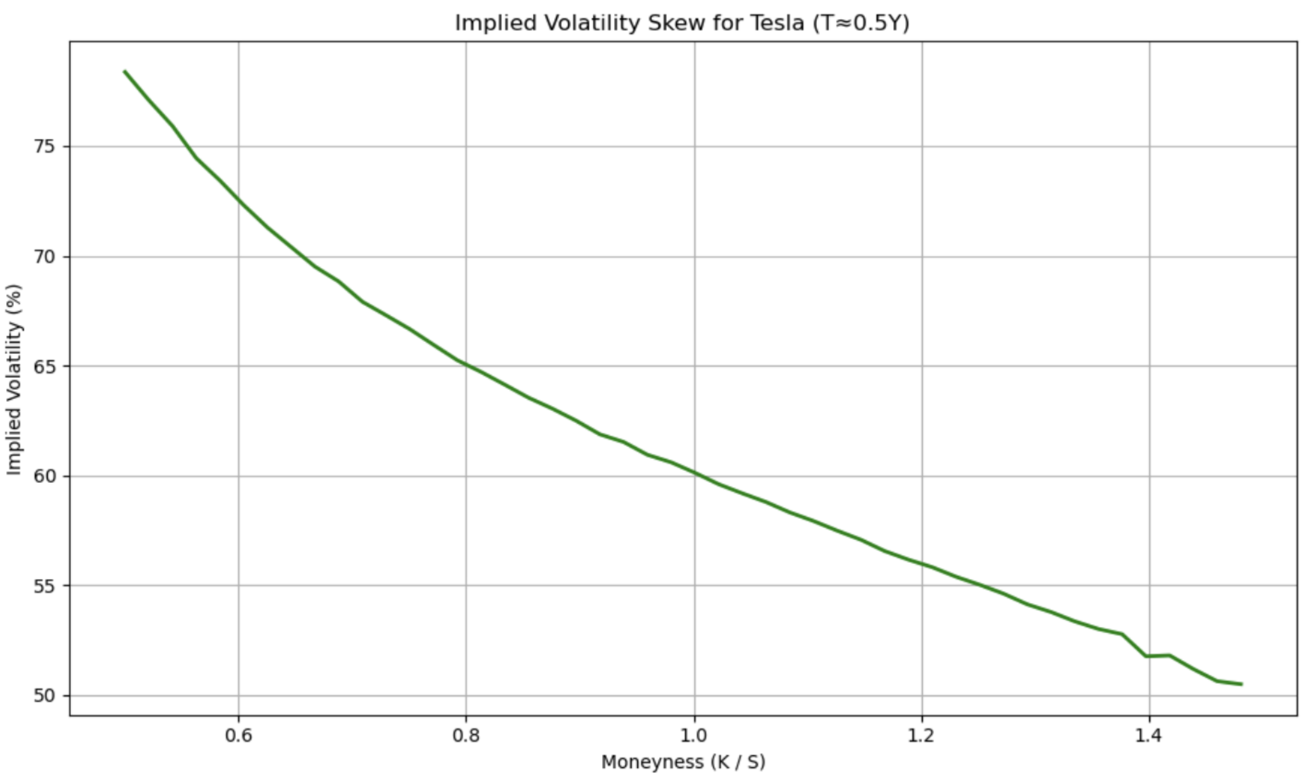

Fig.2: Volatility skew for Tesla calls 6 months from expiration, as of 18/04/2025

2.2 Term Structure of Implied Volatility

In the previous section we have established that, holding time to maturity constant, the implied volatility of options varies widely with changes in the strike price, in contradiction with the assumptions of the Black-Scholes model. We can show that a similar phenomenon shows up when holding the strike price constant and plotting implied volatility for various times to maturity. This second type of variation in implied volatility is known as the term structure of implied volatility. We typically observe that implied volatility increases or decreases systematically with time to maturity, depending on market conditions and the underlying asset. This behavior reflects the market's evolving expectations of future volatility.

For instance, options on equities often display an upward-sloping term structure in times of market calm, implying that investors expect volatility to rise in the future. Conversely, during periods of market stress, the term structure may invert, reflecting short-term spikes in uncertainty that are expected to subside over time. The latter case can be observed in the graph below, as a result of the uncertainty brought about by U.S. President Donald Trumps' Tariffs.

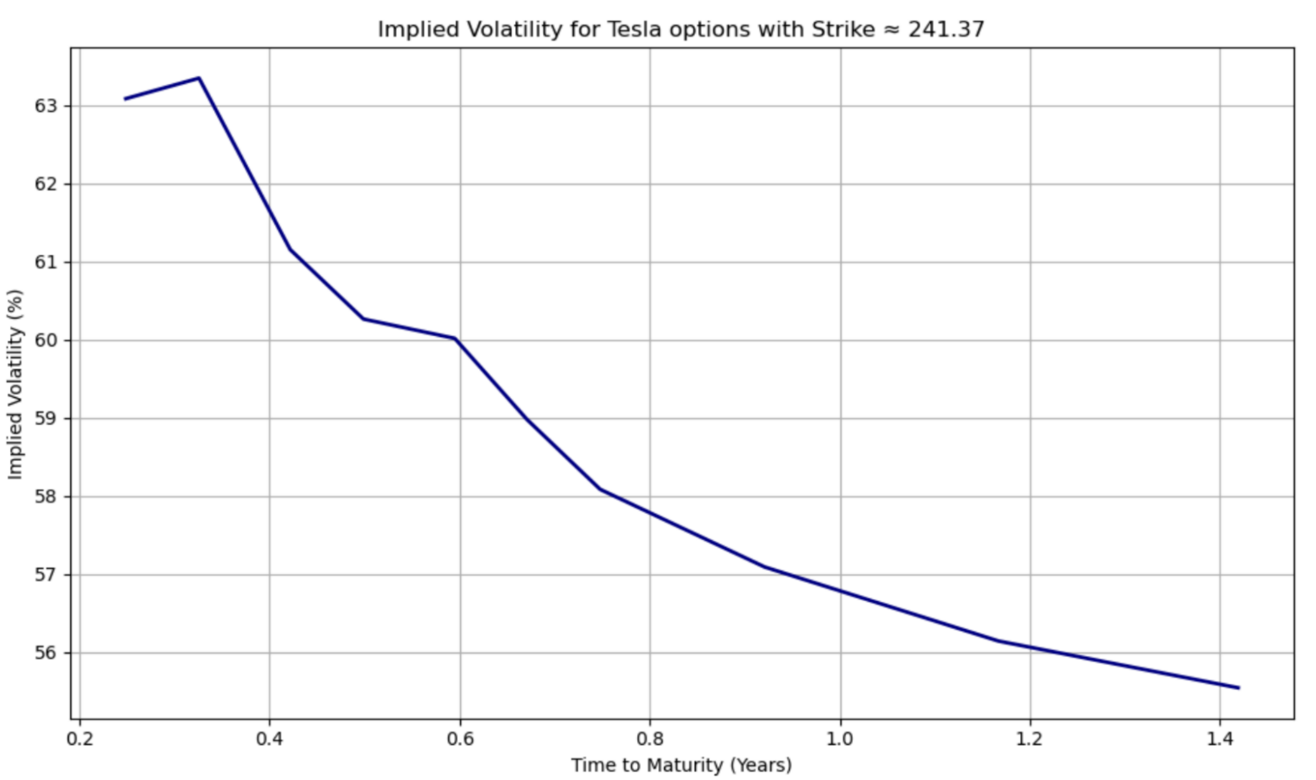

Fig. 3: Implied volatility for at-the-money Tesla calls, as of 18/04/2025

2.3 The Volatility Surface

To get a complete picture of the market's implied volatility, practitioners often work with a more comprehensive object that combines volatility skew and term structure: the volatility surface. The volatility surface is a function that assigns an implied volatility $\sigma(K,T)$ to each combination of strike price K and time to maturity T. Formally, for each pair (K,T), the implied volatility is defined as the unique value σ such that the Black-Scholes price of the option equals its observed market price. That is,

$$C(S,K,T) = BS(S,T,r,K,\sigma(K,T)) \qquad \textbf{(2.3.1)}$$

Where $C(K,S,T)$ is the defined as the market price of a call option with strike price K and time to maturity T, while $BS(\cdot)$ is the Black-Scholes formula for the price of a call option.

2.3.1 Constructing the Volatility Surface

When deciding how to graph the implied volatility surface, perhaps the most obvious approach would be to express the surface in terms of the strike price K and time to maturity T, given that both values can be easily extracted from market data. However, plotting implied volatility directly against strike may produce surfaces that are not easily comparable across different underlyings or time periods, since the range of relevant strikes depends on the spot price S, which obviously can vary widely. To address this, many practitioners instead use moneyness, defined as the ratio between strike and spot price K/S. This allows us to normalize the strike relative to the underlying and provides us with more stable visualizations, while retaining the aforementioned advantages. This is why it was chosen for the figure below.

Another widely adopted parameterization uses delta as the horizontal axis, something especially common in FX and Equity Derivatives markets. This approach has several practical benefits. First, it aligns with how traders hedge positions, since delta is directly related to trading risk. Second, delta-bucketed surfaces tend to remain more stable under changes in the underlying asset price, making them more suitable for quoting and interpolation. However, using delta comes with some caveats. Since delta depends on the very volatility being modeled, constructing a delta-based surface requires an iterative or inverse mapping between implied volatility and delta. Additionally, different conventions exist for computing delta (e.g., spot delta vs. forward delta), which must be clearly specified to avoid ambiguity.

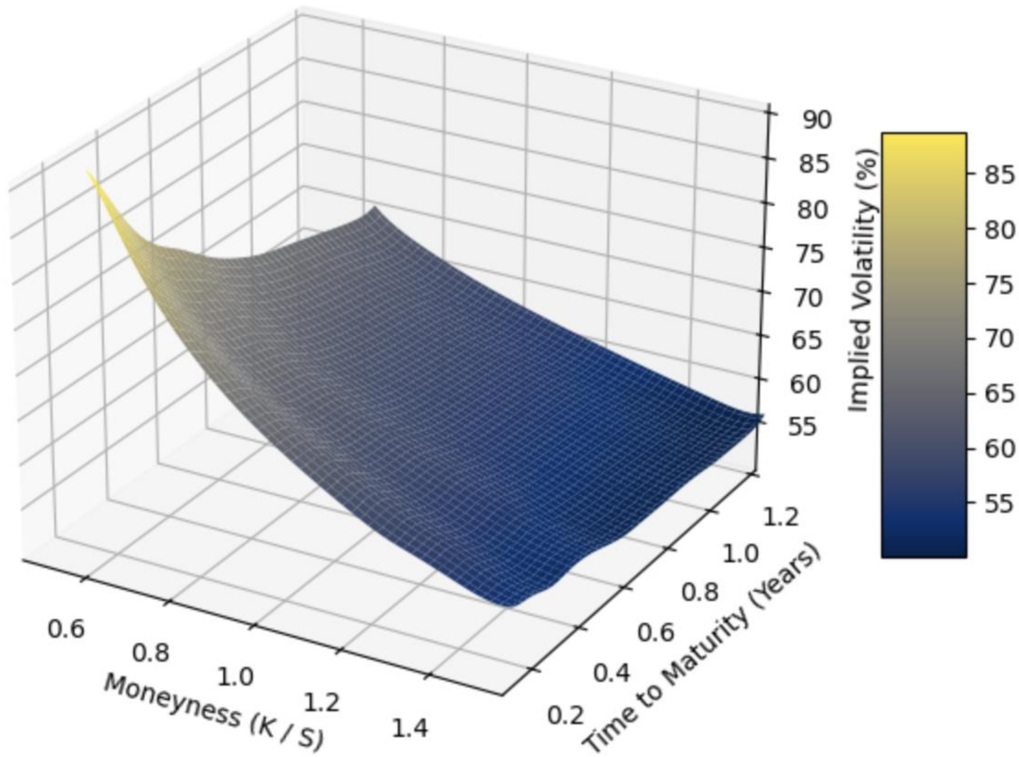

Fig. 4: Volatility surface for Tesla, as of 18/04/2025

2.3.2 No-Arbitrage Conditions

To be financially meaningful, the volatility surface must be free of static arbitrage. Static arbitrage refers to arbitrage opportunities that can be exploited by constructing simple portfolios of vanilla options without any need for rebalancing. Two major forms of static arbitrage are calendar spread arbitrage and butterfly arbitrage, each of which imposes important mathematical constraints on the shape of the implied volatility surface.

Calendar spread arbitrage arises when the price of a longer-dated option is lower than the price of an otherwise identical shorter-dated one. To better explain this phenomenon, we first need to introduce the concept of time value, which can be viewed as the premium above the intrinsic value arising from the uncertainty about future movements in the underlying asset before expiration. As time to maturity increases, the time value of an option also increases, because there is more opportunity for the underlying asset to move into a favorable position. Even if an option is currently out of the money, more time implies a higher probability of becoming profitable, thereby increasing its value. Consequently, holding all else equal, an option with longer maturity should never be cheaper than an identical one with shorter maturity. This means that, to prevent calendar arbitrage, the price of a European call option must always satisfy:

$$\frac{\partial C}{\partial T} \geq 0$$

If this inequality is violated, it would imply that a trader could buy the underpriced long-dated option and sell the overpriced short-dated one, both with the same strike, locking in a profit without risk, hence violating the no-arbitrage principle.

On the other hand, butterfly arbitrage arises when the option price is not sufficiently convex in the strike price. In a well-functioning market, the price of an option should increase more slowly as the strike moves away from the money, but always in a smooth, convex manner. If this curvature becomes negative, meaning the option price curve bends downward, then arbitrage opportunities can be constructed using a butterfly spread strategy.

A butterfly spread involves:

- Buying one call option with a low strike K_1,
- Selling two call options with a middle strike K_2,
- Buying one call option with a high strike K_3, where $K_1 < K_2 < K_3$ and the strikes are equally spaced.

The total value of this strategy, V_0, at $t = 0$, is equal to:

$$V_0 = C(K_1, T) - 2C(K_2, T) + C(K_3, T)$$

Which can be approximately interpreted as its discounted value at expiration, or:

$$e^{-rT} f(K_2, T) * (K_3 - K_1)^2$$

Where $f(K_2, T)$ represents the probability distribution function of S_T around K_2, meaning:

$$f(K_2, T) \approx \frac{C(K_1, T) - 2C(K_2, T) + C(K_3, T)}{(K_3 - K_1)^2} \qquad \textbf{(2.1.2)}$$

Then, for $(K_3 - K_1) \longrightarrow 0$, we get:

$$f(K_2, T) = e^{rT} \frac{\partial^2 C}{\partial {K_2}^2} \qquad \textbf{(2.1.3)}$$

This result, commonly called the *Breeden-Litzenberg formula*, reveals a link between the curvature of the call price function and the implied probability distribution of the asset's future value under the risk-neutral measure. Specifically, it shows that the second derivative of the call price with respect to strike reflects the market-implied probability density function of the asset price at maturity. Given that any PDF is, by definition, non-negative, it also implies that, if the condition

$$\frac{\partial^2 C}{\partial K^2} \geq 0$$

is violated at any point along the strike axis, the implied probability density becomes not meaningful from a probabilistic perspective. Such a scenario would indicate a violation of no-arbitrage conditions and would allow for the construction of butterfly spreads with negative cost and non-negative payoff, giving rise to a clear arbitrage opportunity.

3. Local Volatility Models

Following the 1987 market crash, the limitations of the Black-Scholes model became increasingly evident, particularly in its inability to capture the volatility smile and skew observed in option markets. These discrepancies prompted both academics and practitioners to seek more refined models capable of reproducing such market phenomena while preserving the no-arbitrage framework underpinning Black-Scholes. One notable approach is to model volatility as a deterministic function of both the underlying asset price and time, leading to the concept of local volatility. In this setting, volatility is described by a smooth function $\sigma(S, t)$ that varies across the state space, allowing the model to match the entire implied volatility surface observed in the market. Local volatility models retain analytical tractability and are consistent with arbitrage-free pricing. This chapter introduces the theoretical foundations of local volatility, beginning with the discrete-time solution given by Derman and Kani through the use of binomial trees. Then, still with the help of the binomial model, Dupire's Forward Equation is introduced and derived from two of the most common option strategies.

3.1 Derman-Kani Implied Volatility Tree

In the quest for local volatility models, among the earliest and most influential contributions to the effort was the implied volatility tree proposed by Derman and Kani in 1994. Rather than immediately resorting to complex continuous-time stochastic models, they chose to work within the discrete setting of a binomial tree, which offered an intuitive and tractable structure for modeling uncertainty in asset prices. By allowing the tree's branching behavior to vary locally with the asset price and time, they were able to reproduce the shape of the implied volatility surface.

3.1.1 The Binomial Model

To lay the groundwork for constructing an implied volatility tree, it is first necessary to introduce the classical binomial model of asset price dynamics. In its basic form, the binomial model assumes that over a small discrete time interval dt, an asset priced S can move either up to Su or down to Sd, where $u >$

1 and $0 < d < 1$ are constant factors, with probability respectively p and $(1-p)$. To be consistent with the Black-Scholes framework, where stocks are assumed to follow $dS = \mu dt + \sigma dW$, we need the binomial parameters u, d, p to match the Black-Scholes parameters μ, σ. Namely, we require that

$$pu + (1-p)d = \mu dt \tag{3.1.1.1}$$

$$p(1-p)(u-d)^2 = \sigma^2 dt \tag{3.1.1.2}$$

A particularly useful version of the binomial model is the Cox-Ross-Rubinstein (CRR) parametrization, which imposes the condition $u + d = 0$. This choice ensures that the tree is recombining, meaning that an up move followed by a down move leads to the same price as a down move followed by an up move, and allows us to simplify the equations to

$$(2p-1)\,u = \mu dt \tag{3.1.1.3}$$

$$4p\,(1-p)\,u^2 = \sigma^2 dt \tag{3.1.1.4}$$

By squaring **3.1.1.3** and adding it to **3.1.1.4**, we obtain

$$u^2 = \mu^2 dt^2 + \sigma^2 dt \tag{3.1.1.5}$$

The limit of **3.1.1.5** as $dt \longrightarrow 0$ becomes

$$u = \sigma\sqrt{dt} \tag{3.1.1.6}$$

And, given that $u = -d$

$$d = -\sigma\sqrt{dt} \tag{3.1.1.7}$$

Then, to obtain p as an equation of μ and σ we can substitute **3.1.1.6** into **3.1.1.3**:

$$p = \frac{1}{2} + \frac{1}{2}\frac{\mu}{\sigma}\sqrt{dt} \tag{3.1.1.8}$$

3.1.2 Option Valuation in the Binomial Model

We can show that the binomial model can be used to obtain the price of an option starting from a risky stock and a riskless bond, much like we did in the chapter on Black-Scholes. We begin by defining two securities Π_u and Π_d. The former pays out 1 if, after dt time has passed, the stock is in the up state, and 0 if in the down state, while the latter does the exact opposite. We know that, given the stock S, the value of 1 invested in it after dt time has passed, which we are going to call 1_S, is either $U = Su/S = e^u$ or $D = Sd/S = e^d$, while the value of 1 invested in the bond B, which we call 1_B, will always be equal to e^{rdt}, where r is the risk free rate of interest. We then define Π_u in terms of 1_S and 1_B, $\Pi_u = \alpha 1_S + \beta 1_B$, both for the up and the down state.

Just as we did when deriving Black-Scholes, we can construct a riskless portfolio with a guaranteed payout of 1 by combining equal amounts of Π_u and Π_d. Because of the no-arbitrage principle, we know that this portfolio will return the same as the risk-free bond, so that its present value is equal to

$$\Pi_u + \Pi_d = e^{-rdt} \quad \textbf{(3.1.2.1)}$$

To determine α and β, we require that the portfolio replicates the correct payoffs in both possible future states. This leads to the system of linear equations:

$$\alpha U + \beta e^{rdt} = 1 \quad (payoff\ in\ the\ up\ state)$$

$$\alpha D + \beta e^{rdt} = 0 \quad (payoff\ in\ the\ down\ state)$$

From which we can find

$$\alpha = \frac{1}{U - D} \quad \textbf{(3.1.2.2)}$$

$$\beta = \frac{-e^{-rdt} D}{U - D} \quad \textbf{(3.1.2.3)}$$

3.1.2.2 and **3.1.2.3** can then be substituted back to find the present values of Π_u and Π_d, namely

$$\Pi_u = \frac{e^{rdt} - D}{e^{rdt}(U - D)} \quad \textbf{(3.1.2.4)}$$

$$\Pi_d = \frac{U - e^{rdt}}{e^{rdt}(U - D)} \tag{3.1.2.5}$$

We know that, under the risk-neutrality assumption, the value of these securities can also be defined as the discounted risk-neutral probability of obtaining their payoff, which we can call q for Π_u and (1-q) for Π_d. This means we can write $\Pi_u = qe^{-rdt}$ and $\Pi_d = (1 - q)e^{-rdt}$, where

$$q = \frac{e^{rdt} - D}{U - D} \tag{3.1.2.6}$$

$$1 - q = \frac{U - e^{rdt}}{U - D} \tag{3.1.2.7}$$

We can therefore rewrite the stock's price as a discounted weighted average of S_u and S_d:

$$S = e^{-rdt}(qS_u + (1 - q)S_d) \tag{3.1.2.8}$$

We could then go on to prove that the binomial model converges to Black-Scholes when $dt \longrightarrow 0$, although the proof is beyond the scope of this chapter.

3.1.3 Incorporating Term Structure

In our quest to properly account for the variations in implied volatility, we begin by imagining a volatility which varies with time, so that the stock is now modeled by

$$\frac{dS}{S} = \mu dt + \sigma(t)dW \tag{3.1.3.1}$$

This immediately presents a problem however, because having volatility vary means that an up move followed by a down move

$$Se^{u+d} = Se^{\sigma_1\sqrt{dt_1} - \sigma_2\sqrt{dt_2}} \tag{3.1.3.2}$$

will not necessary be equal to a down move followed by an up move

$$Se^{d+u} = Se^{-\sigma_1\sqrt{dt_1} + \sigma_2\sqrt{dt_2}} \quad \textbf{(3.1.3.3)}$$

This makes the tree non-recombining and significantly complicates calculations. A potential solution to the issue is to require that

$$\sigma_i\sqrt{dt_i} = \sigma_{i+1}\sqrt{dt_{i+1}} \quad \forall i \in \mathbb{N}$$

This makes the tree recombining, but leaves us with an unknown number of steps to maturity. Given the unknown size of the time step, the value of q for a generic i therefore varies as well:

$$q_i = \frac{e^{rdt_i} - e^{-\sigma_i\sqrt{dt_i}}}{e^{\sigma_i\sqrt{dt_i}} - e^{-\sigma_i\sqrt{dt_i}}} \quad \textbf{(3.1.3.4)}$$

$$1 - q_i = \frac{e^{\sigma_i\sqrt{dt_i}} - e^{rdt_i}}{e^{\sigma_i\sqrt{dt_i}} - e^{-\sigma_i\sqrt{dt_i}}} \quad \textbf{(3.1.3.5)}$$

This apparent problem is however quickly solved by remembering that we are only concerned about the stock price distribution at expiration, meaning the total variance of the stock. We can therefore reconnect with the Black-Scholes implied volatility by expressing the implied variance $\Sigma^2(t,T)$ of a European option maturing at time T as the time-average of the instantaneous (or forward) variances $\sigma^2(t)$ over the interval from t to T:

$$\Sigma^2(t,T) = \frac{1}{T-t}\int_t^T \sigma^2(v)\,dv \quad \textbf{(3.1.3.6)}$$

3.1.4 Incorporating the Skew

Having incorporated the term structure in the stochastic differential equation, we know aim to include in it the volatility smile/skew:

$$\frac{dS}{S} = (r-b)dt + \sigma(S,t)dW \quad \textbf{(3.1.4.1)}$$

It must be noted that, in this formulation, the drift term *r - b* reflects that we are working under the risk-neutral measure. Here, *r* is the risk-free rate and *b* represents the continuous dividend yield of the asset, while $\sigma(S,t)$ represents local volatility. The instantaneous variance of S in dt is

$$(dS)^2 = S^2\sigma^2(S,t)dt \qquad \textbf{(3.1.4.2)}$$

While the expected value after dt is:

$$F = Se^{(r-b)dt} \qquad \textbf{(3.1.4.3)}$$

From **3.1.2.7**, we then derive the forward price and solve for q to obtain:

$$q = \frac{F - S_d}{S_u - S_d} \qquad \textbf{(3.1.4.4)}$$

Given that the variance of a random variable x is defined as $Var[x] = E[(\mathrm{x} - \mathrm{E[x]})^2]$, we can define

$$Var[dS] = q(S_u - \mathrm{F})^2 + (1 - q)(S_d - \mathrm{F})^2 \qquad \textbf{(3.1.4.5)}$$

But, if we take the limit $dt \longrightarrow 0$, **3.1.4.2** and **3.1.4.5** have to coincide, so we put

$$S^2\sigma^2(S,t)dt = (S_u - \mathrm{F})^2 + (1 - q)(S_d - \mathrm{F})^2 \qquad \textbf{(3.1.4.6)}$$

We then substitute **3.1.4.4** into **3.1.4.6**:

$$S^2\sigma^2(S,t)dt = (S_u - F)(F - S_d) \qquad \textbf{(3.1.4.7)}$$

And solve for S_u and S_d:

$$S_u = F + \frac{S^2\sigma^2(S,t)dt}{F - S_d} \qquad \textbf{(3.1.4.8)}$$

$$S_d = F - \frac{S^2\sigma^2(S,t)dt}{S_u - F} \qquad \textbf{(3.1.4.9)}$$

3.1.5 Building the Tree

We can now start to build our tree, beginning from S_0, the initial stock price at the root of the tree. Following the spine-based construction, we first determine the central path by alternating between odd and even time steps. At odd time levels, the central node remains fixed at S_0, ensuring that the tree is recombining. At even time levels, two central nodes are calculated using the CRR convention to generate upward and downward movements that follow **3.1.3.2** and **3.1.3.3** thus defining the core "spine" of the tree. Once the spine has been established, we recursively expand the tree outward by determining the remaining non-central nodes using the previously derived expressions **3.1.4.7** and **3.1.4.8**.

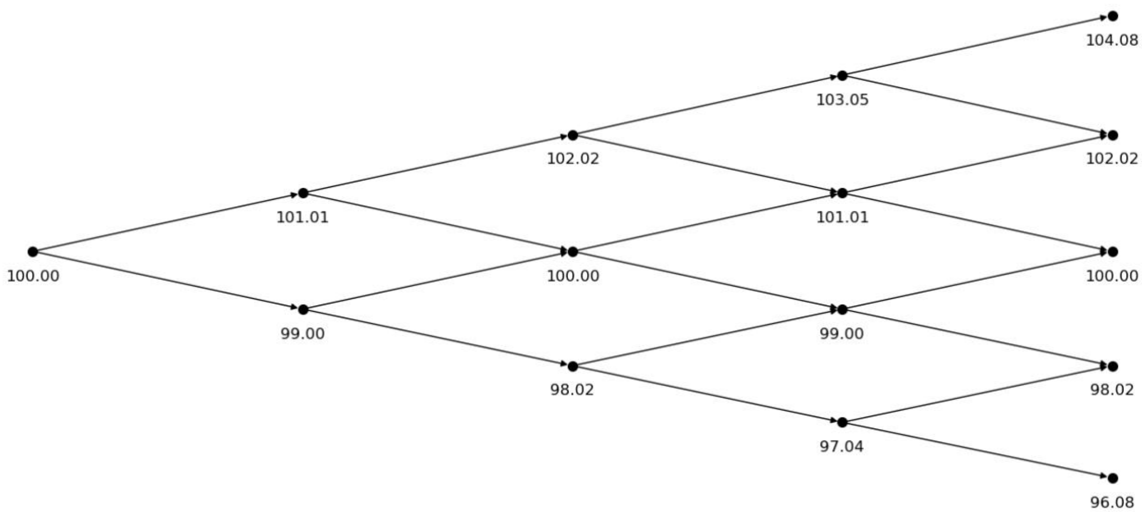

Fig. 5: Binomial tree of stock prices, constructed under the assumption of constant volatility

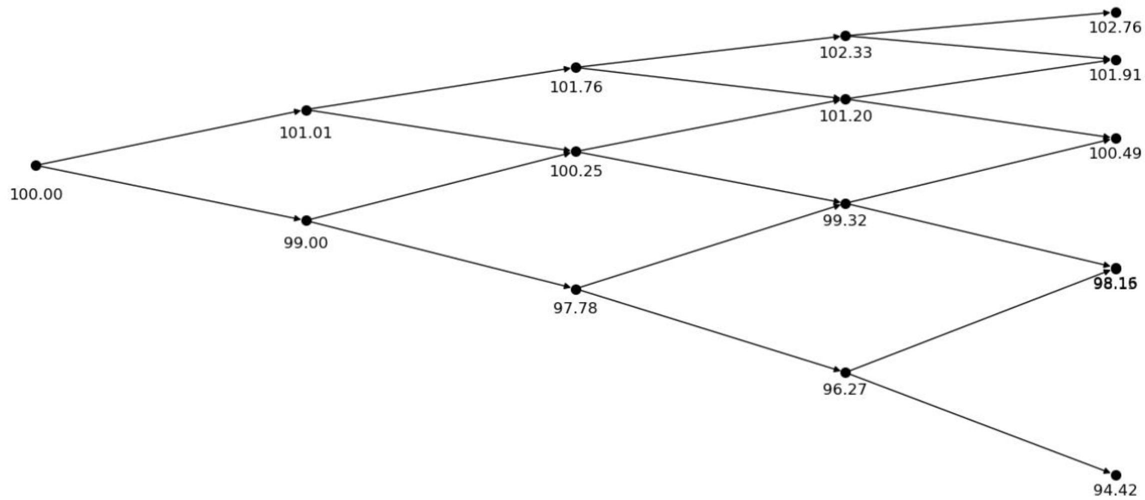

Fig. 6: Stock prices Derman-Kani implied volatility tree

A comparison of the two trees reveals how different modeling assumptions lead to significantly different outcomes in option valuation. The first tree is built under the assumption of constant local volatility, while the second incorporates a hypothetical volatility skew through a state-dependent local volatility function. As a result, the second tree captures asymmetries in the distribution of future asset prices that are consistent with observed market phenomena such as the implied volatility skew.

To illustrate this, consider the pricing of a European put option with a strike price of 95.5 and maturity at t = 4. In the constant volatility tree, none of the terminal nodes at t = 4 reach or fall below the strike price. Consequently, the model would assign a price of zero to the option, as it finishes out-of-the-money across all possible paths. However, in the local volatility tree, the lowest terminal node falls below the strike, and the option would expire in-the-money along at least one path. This implies that the option must have a positive value under this model. To compute the correct price, we must consider the risk-neutral probabilities implied by the structure of the tree. These determine the likelihood of each path under the risk-neutral measure and allow us to compute the expected discounted payoff of the option.

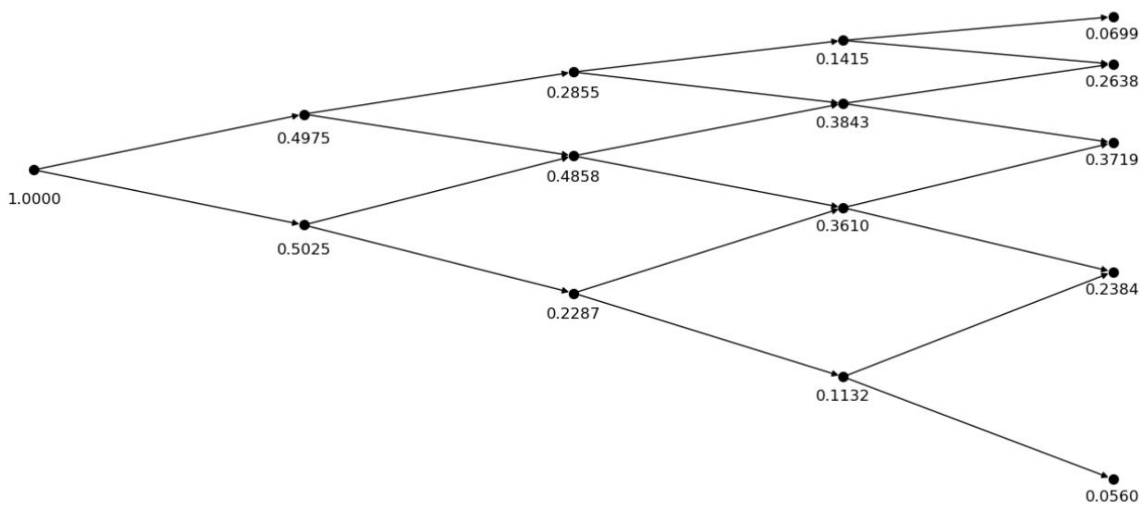

Fig. 7: Risk-neutral probabilities of the previous Derman-Kani implied volatility tree

To calculate the value of the aforementioned option, we calculate the profit if the stock lands on the bottom node at t = 4 and multiply it by the probability shown in the graph: $(95.5 - 94.2) * 0.0560 = 0.0728$. This means that, if we assume that the risk-free rate of interest r = 0, the price of the option at t = 0 according to this model is 0.0728, a small but significant difference.

Just as we define a Black–Scholes implied volatility for an option (the constant volatility input that reproduces the option's market price in the BS formula), we can similarly define a CRR implied volatility for a given option. The CRR implied volatility is the single volatility value such that a CRR binomial tree (with that constant σ) yields the option's correct market price. As we have already mentioned, the CRR model approximates Black-Scholes as time steps tend to zero, which also means that CRR implied volatility will approach the Black–Scholes implied volatility.

Empirically, we can find that the constant volatility needed in a CRR tree to match an option's price is approximately the linear average of the true local volatility values between the current stock price S_0 and the option's strike K. To justify this proposition, we first intuitively notice that any path that leads to a call option finishing in the money must see the stock move from S_0 to K (or beyond) by expiration. Thus the option's payoff "samples" the local volatility function over $S_0 \leq S_t \leq K$ before expiry. If local volatility is roughly constant throughout that region, clearly the implied vol will equal that value. If instead $\sigma(S)$ varies, the option's price will be influenced by the portions of the path where volatility is higher or lower. To a first approximation, these effects average out linearly. For example, an out-of-the-money call will derive most of its value from paths that climb upward into the $[S, K]$ range (since paths far below K contribute little to payoff). Conversely, an in-the-money put (strike $K > S$) depends on paths that move downward through $[K, S]$. Therefore, in both cases the implied volatility effectively "feels" the local volatility in that band.

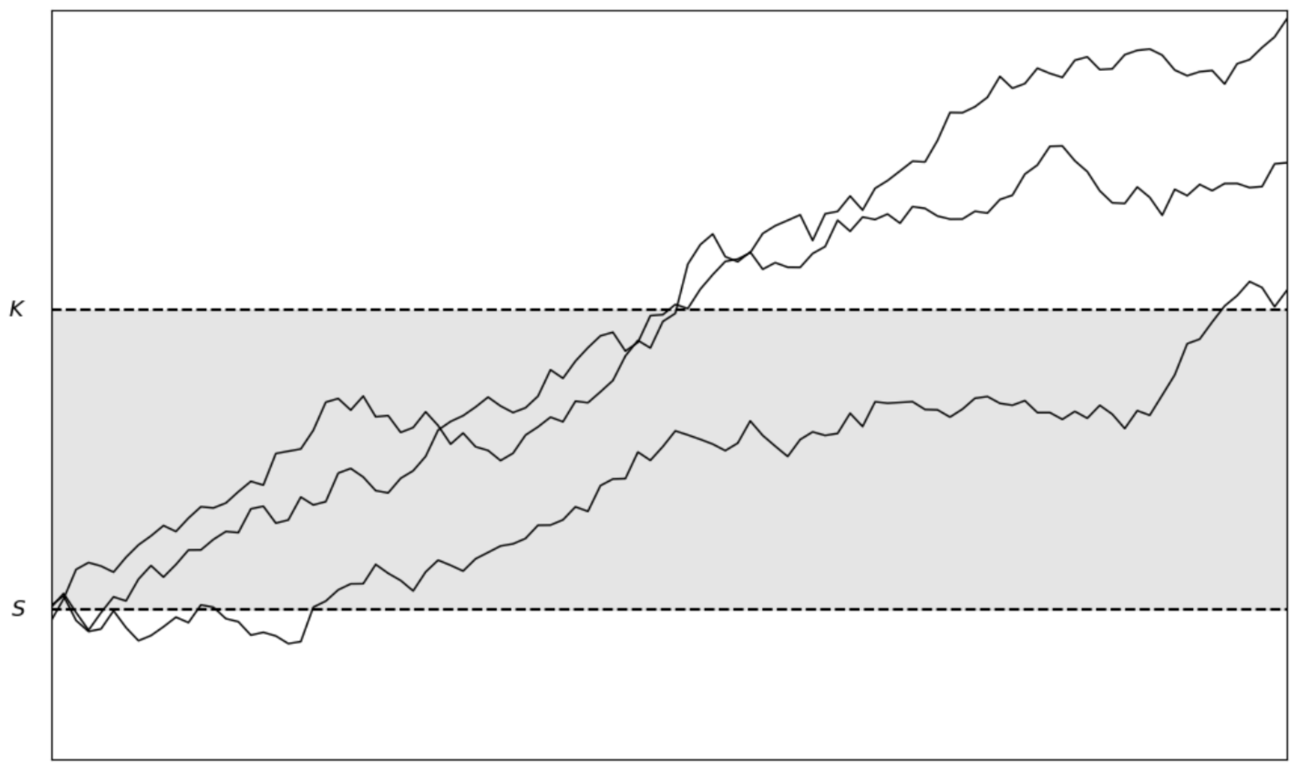

Fig. 7: All paths to expiration which end in-the-money must traverse [S,K]

3.2 Dupire Forward Equation

In the previous chapter, when detailing the no-arbitrage conditions of the volatility surface, we introduced the Breeden-Litzenberg formula (**2.1.3.**), which extract the risk-neutral probability distribution function from the market prices of European call options. Similarly, Dupire's Equation, which we will introduce in the coming section, provides a unique mapping from option prices to local volatility.

3.2.1 Derivation using Binomial Trees

In line with what we did for Derman-Kani, we can provide an intuitive way to derive Dupire's forward equation using binomial trees. This time, however, we are not using Cox-Ross-Rubenstein, but instead adopting the Jarrow-Rudd convention, under which $p = 1 - p = 1/2$. For future ease of calculation, we then imagine constructing three time steps of size $\frac{dT}{2}$, on which we implement a calendar spread, which, as we mentioned before, is a trading strategy made up of shorting a call option $C(K,T)$ and buying a call option $C(K,T + dT)$.

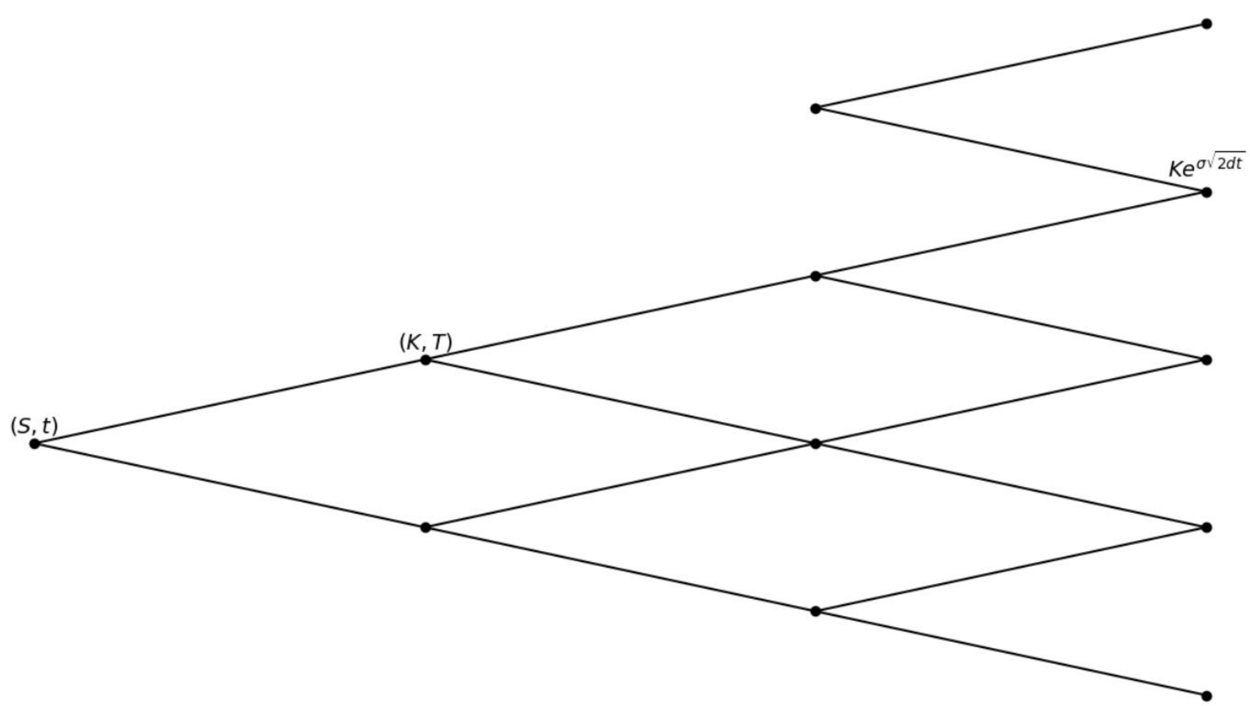

Fig. 8: Jarrow-Rudd Tree. The three most important nodes for our derivation are highlighted

We are going to call the risk-neutral probability of the stock price reaching K by time T $p_{K,T}$. Let's now analyze the effect of each node at time T, with the aid of the above figure as a point of reference.

Nodes with underlying price below K result in zero payoff for both options and thus contribute nothing to the spread. Similarly, nodes with prices above K at time T are in-the-money for the short call, but because of the martingale property of the tree under the risk-neutral measure, the expected payoff of the long call at $T + dT$ exactly matches that of the short call at T, again yielding zero net contribution. The only non-trivial contribution arises from the node where the stock price equals the strike K at time T. At this point, the shorter option expires worthless, but the longer-dated option may end up in-the-money through its possible transitions to future nodes. In particular, only the highest of the three child nodes at $T + dT$ contributes a positive payoff, and the probability of reaching it is $\frac{1}{4}p_{K,T}$ (remember that the central node has twice the probability of being reached than the outer ones, given there are two ways of getting to it). This makes the value of the spread equal to the probability weighed payoff of the node, or

$$C(K, T + dt) - C(K, T) = \frac{1}{4}p_{K,T}dK \tag{3.2.1.1}$$

Additionally, if we divide by dT, as $dT \rightarrow 0$ we obtain

$$\frac{\partial C(K, T)}{\partial T} \approx \frac{C(K, T + dt) - C(K, T)}{\partial T} = \frac{1}{4}p_{K,T}\frac{dK}{dT} \tag{3.2.1.2}$$

Which will be the numerator of the formula. For the next part of the derivation, we can borrow chapter two's infinitesimally small butterfly spread centered on K, which pays dK if $S = K$ at $t = T$. The probability $p_{K,T}$ of obtaining the payoff approaches the Breeden-Litzenberg formula we obtained in **2.1.3** as $dK \rightarrow 0$:

$$p_{K,T} \approx \frac{\partial^2 C(K, T)}{\partial K^2} dK \tag{3.2.1.3}$$

Then, plugging **3.2.1.3** into **3.2.1.2** we get

$$\frac{\partial C(K, T)}{\partial T} = \frac{1}{4}\frac{\partial^2 C(K, T)}{\partial K^2}\frac{dK^2}{dT} \tag{3.2.1.4}$$

We then remember that, following the notation we first introduced in **3.1.1.6**, the payoff of the calendar spread when the upper node is reached, which we have simply called dK so far, is equal to

$$dK = Ke^{2u} - K = Ke^{\sigma\sqrt{2dT}} - K \approx K\sigma\sqrt{2dT} \tag{3.2.1.5}$$

This means that finally, by substituting **3.2.1.5** into **3.2.1.4.** and rearranging the terms, we get Dupire's equation for risk-free rates equal to zero:

$$\frac{\sigma^2(K,T)}{2} = \frac{\frac{\partial C(K,T)}{\partial T}}{K^2 \frac{\partial^2 C(K,T)}{\partial K^2}} \tag{3.2.1.6}$$

Being able to derive this equation from two of the most common option strategies allows us to easily determine a unique local volatility function from market prices and therefore construct a local volatility surface, making Dupire's model quite attractive for practitioners.

Conclusion: An Example of Local Volatility for Exotic Option Pricing

Among the main applications for developing more sophisticated volatility models is the pricing of exotic derivatives whose value depends on the path of the underlying asset, not merely its terminal value. For this example, we will pick a lookback option, whose payoff is the value of the underlying at expiration minus the minimum of its value from emission to expiration. Its value at time τ is $C_{LB}(S_T, M_T, \tau)$, so that at expiration

$$C_{LB}(S_T, M_T, 0) = max(S_T - M_T, 0) \tag{4.1}$$

To understand how this type of option could be mispriced under the assumption of constant volatility, let's consider the scenario in which the current price of the underlying asset equals its minimum, $S_T = M_T$. Under such conditions, a small upward move in the asset price should logically have approximately the same effect on the option's value as a corresponding decrease in the running minimum. Mathematically, this leads to the identity:

$$\left.\frac{\partial C_{LB}}{\partial S}\right|_{S_T = M_T} \approx -\left.\frac{\partial C_{LB}}{\partial M}\right|_{S_T = M_T} \tag{4.2}$$

On a binomial model with the Jarrow-Rudd convention, over a small time step $dS = S\sigma\sqrt{\tau}$:

$$C_{LB}(S_t, S_t, \tau) = \frac{1}{2}C_{LB}(S_t + dS, S_t, \tau + dt) + \frac{1}{2}C_{LB}(S_t - dS, S_t - dS, \tau - dt) \tag{4.3}$$

Expanding the expectation of the future value of the lookback option using a Taylor approximation and assuming risk neutrality, we arrive at:

$$C_{LB}(S_t, S_t, \tau) \approx C_{LB}(S_t, S_t, \tau) - \frac{\partial C_{LB}}{\partial M}\frac{dS}{2} \tag{4.4}$$

Given that by definition $dS > 0$, it must mean that:

$$\frac{\partial C_{LB}}{\partial M} \approx 0$$

From **4.2**, this also means that for $S_T = M_T$, the delta of the option is zero.

What does this mean for option prices? From a reasoning similar to the one in chapter **3.1.5**, the vast majority of their value comes from a selected number of “dominant” paths. For lookback calls, that is paths which reach their minimum early on and then go on to rise, behaving similarly to vanilla in-the-money calls. Inversely, dominant paths for lookback puts first reach a high minimum, then decline and behave like standard puts. From what we learned in chapter 2, the presence of a volatility skew (*Fig.2*) means that, for lower levels of moneyness (K/S) volatility is higher than for higher levels of moneyness. Therefore, when accounting for the negative skew, lookback calls will be more expensive (and lookback puts less expensive) than their Black-Scholes counterparts.

Bibliography

Hull, J. C. (2021). *Options, futures, and other derivatives* (11th ed.). Pearson.

Shreve, S. E. (2004). *Stochastic calculus for finance II: Continuous-time models*. Springer.

Klebaner, F. C. (2012). *Introduction to stochastic calculus with applications* (3rd ed.). Imperial College Press.

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

Derman, E. (2016). *The Volatility Smile*. Wiley.

Derman, E. & Kani, I. (1994). *Riding on a Smile*. Risk, 7.

Derman, E. and Kani, I. (1994). *The volatility smile and its implied tree*. Quantitative Strategies Research Notes, Goldman Sachs.

Björk, T. (2020). *Arbitrage Theory in Continuous Time* (4th ed.). Oxford University Press.

Dupire, B. (1994). *Pricing with a Smile*. Risk, 7.

Black, F., & Scholes, M. (1973). *The Pricing of Options and Corporate Liabilities*. Journal of Political Economy.

Cox, J. C., Ross, S. A., & Rubinstein, M. (1979). *Option Pricing: A Simplified Approach*. Journal of Financial Economics

Black, F. (1975). *Fact and Fantasy in the Use of Options and Corporate Liabilities*. Financial Analysts Journal.

Merton, R. C. (1974). *On the Pricing of Corporate Debt: The Risk Structure of Interest Rates*. The Journal of Finance

Derman, E., Kani, I., & Zou, J. Z. (1996). *The local volatility surface: Unlocking the information in index option prices*. Financial Analysts Journal