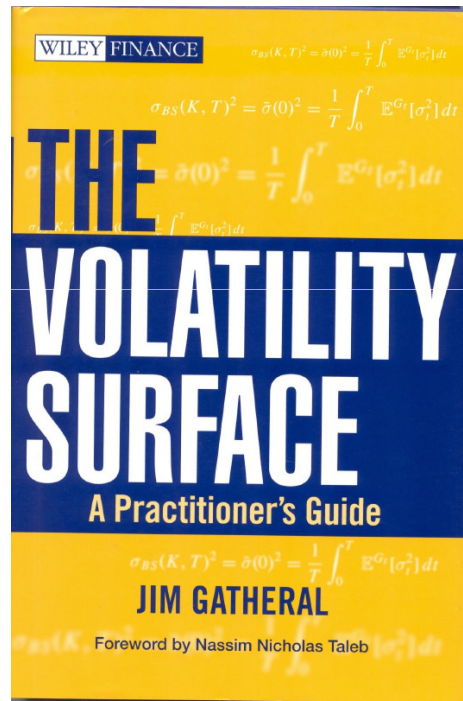


Book Review:
'The Volatility Surface. A Practitioner's
Guide'
by Jim Gatheral
(Wiley-Finance, 2006) *

Anatoliy Swishchuk
University of Calgary
'Lunch at the Lab' Talk
September, 2010

* *Chapter 1: 'Stochastic Volatility and Local Volatility'*
(*'Lunch at the Lab' Finance Seminar Reviewing Series*)

Book's Cover



Outline of Presentation


1. Introduction to the Book
2. Contents of the Book
3. Chapter 1: Stochastic Volatility and Local Volatility

Intro to the Book: About the Author

Jim Gatheral is a Managing Director at Merrill Lynch and also an Adjunct Professor at the Courant Institute of Mathematical Sciences, New York University.

Dr. Gatheral obtained a PhD in theoretical physics from Cambridge University in 1983. From 1997 to 2005, Dr. Gatheral headed the Equity Quantitative Analytics group at Merrill Lynch. His current research focus is equity market microstructure and algorithmic trading.

Jim Gatheral's Picture



Intro to the Book: Purpose

The purpose of 'The Volatility Surface' is not to just present results, but to provide you with ways of thinking about and solving practical problems that should have many other areas of application.

Intro to the Book: Purpose

Understanding the volatility surface is a key objective for both practitioners and academics in the field of finance.

Intro to the Book: Purpose

Implied volatilities evolve randomly and so models of the volatility surface-which is formed from implied volatilities of all strikes and expirations-need to explicitly reflect this randomness in order to accurately price, trade, and manage the risk of derivative products.

Intro to the Book: Focus

The first half of the book focuses on setting up the theoretical framework, while the later chapters are oriented towards practical applications.

Informative and accessible, 'The Volatility Surface':

Intro to the Book: Focus

- Contains a detailed derivation of the Heston model and explanations of many other popular models such as SVJ, SVJJ, SABR, and CreditGrades
- Discusses the characteristics of various types of exotic options from the humble barrier option to the super exotic Napoleon

Intro to the Book: Focus

- Exhaustively covers volatility derivatives with elegant and robust presentations of the latest research
- Examines performance of exotic cliquet contracts through in-depth case studies of actual bonds that have already matured

Contents of the Book

Book contains:

11 Chapters

Bibliography (74 sources)

Index

Many figures (47)

12 Tables

Contents of the Book: Chapters' Short Descriptions

Chapter 1 'Stochastic Volatility and Local Volatility': Provides an explanation of stochastic (SV) and local volatility (LV); local variance is shown to be the risk-neutral expectation of instantaneous variance, a result that is applied repeatedly in later chapters.

Chapter 2 'The Heston Model': Presents the still supremely popular Heston model and derive the Heston European option pricing formula. It's also shown how to simulate the Heston model.

Contents of the Book: Chapters' Short Descriptions

Chapter 3 'The Implied Volatility Surface': Derives a powerful representation for implied volatility (IV) in terms of local volatility and applies this to build intuition and derive some properties of the implied volatility surface (VS) generated by the Heston model and compare with the empirically observed SPX surface; deduces that SV cannot be the whole story.

Chapter 4 'The Heston-Nandi Model': Chooses specific numerical values for the parameters of the Heston model, $\rho = -1$ as originally studied by Heston and Nandi and demonstrates that an approximate formula for implied volatility derived in Chapter 3 works particularly well in this limit. As a result, we are able to find parameters of LV and SV models that generate almost identical European option prices.

Contents of the Book: Chapters' Short Descriptions

Chapter 5 'Adding Jumps': Explores the modelling of jumps showing first why jumps are required; introduces then characteristic function techniques and applies these to the computation of IV in models with jumps; concludes by showing that the SVJ (SV with jumps in the stock price) is capable of generating a volatility surface that has most of the features of the empirical surface.

Chapter 6 'Modelling Default Risk': Applies the work on jumps to Merton's jump-to-ruin model of default; explains the Credit-Grades model.

Contents of the Book: Chapters' Short Descriptions

Chapter 7 'Volatility Surface Asymptotics': examines the asymptotic properties of the VS showing that all models with SV and jumps generate VS that are roughly the same shape.

Chapter 8 'Dynamics of the Volatility Surface': shows how the dynamics of V can be deduced from the time series properties of VS; shows also why it is the dynamics of the VS generated by LV models are highly unrealistic.

Contents of the Book: Chapters' Short Descriptions

Chapter 9 'Barrier Options': presents various types of barrier options (BO) and shows how intuition may be developed for these by studying two simple limiting cases.

Chapter 10 'Exotic Cliquets': studies in detail three actual exotic cliquet transactions that happen to have matured so that we can explore both pricing and ex post performance, specifically, a locally capped and globally floored cliquet, a reverse cliquet, and a Napoleon.

Chapter 11 'Volatility Derivatives' (the longest of all): focuses on the pricing and hedging of claims whose underlying is quadratic variation, explains why the market in volatility derivatives is surprisingly active and liquid.

Chapter 1: 'Stochastic Volatility (SV) and Local Volatility (LV)'

In this chapter, author introduces SV-the notion that volatility varies in a random fashion.

LV then is then shown to be conditional expectation of the instantaneous variance so that various quantities of interest (such as option prices) may sometimes be computed as though future volatility were deterministic rather than stochastic.

Chapter 1: 'Stochastic Volatility (SV)

Sense to model SV as a RV: observations of equity market (the stock market crash of October 1987).

SV models are useful because they explain in a self-consistent way why options with different strikes and maturities have different Black-Scholes implied volatilities-that is, the 'volatility smile'.

Moreover, unlike alternative models that can fit the smile (such as LV models, e.g.), SV models assume realistic dynamics for the underlying.

Chapter 1: 'Stochastic Volatility (SV)

Although SV price processes are sometimes accused by being *ad hoc*, on the contrary, they can be viewed as arising from Brownian motion subordinated to a random clock.

This clock time, often referred to as *trading time*, may be identified with the volume of trades or the frequency of trading (Clark 1973); the idea is that as trading activity fluctuates, so does volatility.

Chapter 1: 'Stochastic Volatility (SV)

From a hedging perspective, traders who use the Black-Scholes model must continuously change the volatility assumption in order to match market prices.

Their hedge ratios change accordingly in an uncontrolled way; SV models bring some order into this chaos.

The prices of exotic options given by models based on Black-Scholes assumptions can be widely wrong and dealers in such options are motivated to find models that can take the volatility smile into account.

Chapter 1: 'Stochastic Volatility (SV)

In Figure 1.1, the log returns of SPX over a 15-year period is plotted; we see that large moves follow large moves and small moves follow small moves (so-called 'volatility clustering').

SPX Daily Log Returns (Dec 31, 1984-Dec 31, 2004)

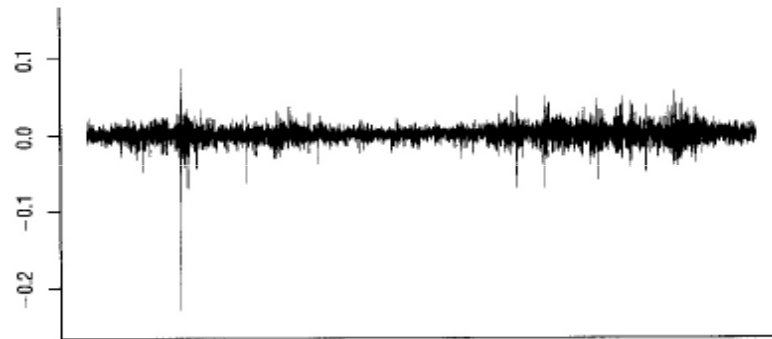


FIGURE 1.1 SPX daily log returns from December 31, 1984, to December 31, 2004. Note the -22.9% return on October 19, 1987!

Chapter 1: 'Stochastic Volatility (SV)

In Figure 1.2, the frequency distribution of SPX log returns over the 77-year period from 1928 to 2005 is plotted. The distribution is highly peaked and fat-tailed relative to the normal distribution.

Frequency Distribution of 77 Years of SPX Daily Log Returns

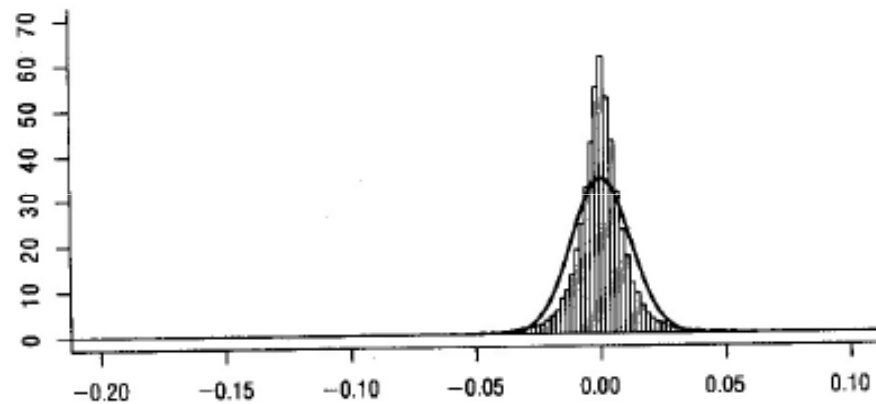


FIGURE 1.2 Frequency distribution of (77 years of) SPX daily log returns compared with the normal distribution. Although the -22.9% return on October 19, 1987, is not directly visible, the x-axis has been extended to the left to accommodate it!

Chapter 1: 'Stochastic Volatility (SV)

The Q-Q plot in Figure 1.3 shows just how extreme the tails of the empirical distribution of returns are relative to the normal distribution. (This plot would be a straight line if the empirical distribution were normal).

Q-Q Plot of SPX Daily Log Returns

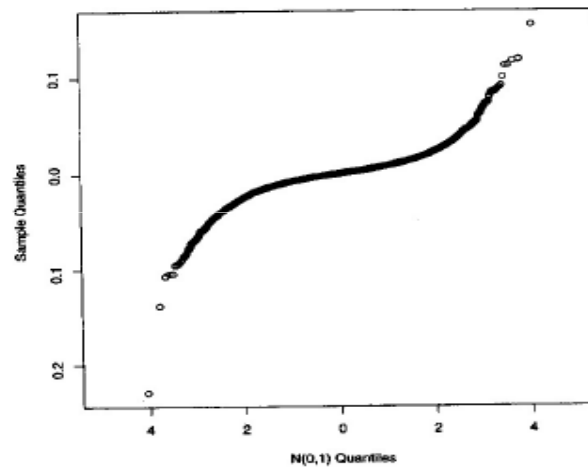


FIGURE 1.3 Q-Q plot of SPX daily log returns compared with the normal distribution. Note the extreme tails.

Chapter 1: 'Stochastic Volatility (SV)

Fat tails and the high central peak are characteristics of mixtures of distributions with different variances.

This motivates to model variance as a RV.

The volatility clustering feature implies that volatility (or variance) is auto-correlated. In the model, this is a consequence of the mean reversion of volatility.

Chapter 1: 'Stochastic Volatility (SV) (Derivation of the Valuation Equation)

In this section, author follows Wilmott (2000). Suppose that the stock price S and its variance v satisfy the following SDEs:

$$\begin{aligned}dS_t &= \mu_t S_t dt + \sqrt{v_t} S_t dZ_1 \\dv_t &= \alpha(S_t, v_t, t) dt + \eta \beta(S_t, v_t, t) \sqrt{v_t} dZ_2\end{aligned}\tag{1}$$

with

$$\langle dZ_1, dZ_2 \rangle = \rho dt.\tag{2}$$

Chapter 1: 'Stochastic Volatility (SV) (Derivation of the Valuation Equation)

The process S_t followed by the stock price is equivalent to the one assumed in the derivation of B-S (1973). This ensures that the standard time-dependent volatility version of the B-S formula (as derived in Sec. 8.6 of Wilmott (2000), e.g.) may be retrieved in the limit $\eta \rightarrow 0$.

In practical applications, this is a key requirement of a SV option pricing model as practitioners' intuition for the behaviour of option prices is invariably expressed within the framework of the B-S formula.

No assumption wrt the form of functions α and β .

Chapter 1: 'Stochastic Volatility (SV) (Derivation of the Valuation Equation)

In the B-S case, there is only one source of randomness, the stock price, which can be hedged with stock. In the present case, random changes in volatility also need to be hedged in order to form a riskless portfolio.

So we set up a portfolio Π containing the option being priced, whose value we denote by $V(S, v, t)$, a quantity $-\Delta$ of the stock and a quantity $-\Delta_1$ of another asset whose value V_1 depends on volatility.

Chapter 1: 'Stochastic Volatility (SV) (Derivation of the Valuation Equation)

We have:

$$\Pi = V - \Delta \times S - \Delta_1 V_1.$$

After applying the Itô formula and making the portfolio instantaneously risk-free by choosing

$$\frac{\partial V}{\partial S} - \Delta_1 \frac{\partial V_1}{\partial S} - \Delta = 0,$$

to eliminate dS terms, and

$$\frac{\partial V}{\partial v} - \Delta_1 \frac{\partial V_1}{\partial v} = 0,$$

to eliminate dv terms, we get

Chapter 1: 'Stochastic Volatility (SV) (Derivation of the Valuation Equation)

$$d\Pi = r\Pi dt = r(V - \Delta \times S - \Delta_1 V_1)dt,$$

where we have used the fact that the return on a risk-free portfolio must equal the risk-free rate r , which is deterministic.

Chapter 1: 'Stochastic Volatility (SV) (Derivation of the Valuation Equation)

Collecting all V terms on the left-hand side and all V_1 terms on the right-side, we get

$$\begin{aligned} \frac{\partial V}{\partial t} + \frac{1}{2}vS^2\frac{\partial^2 V}{\partial S^2} + \rho\eta v\beta S\frac{\partial^2 V}{\partial v\partial S} + \frac{1}{2}\eta^2v\beta^2\frac{\partial^2 V}{\partial v^2} + rS\frac{\partial V}{\partial S} - rV \\ = -(\alpha - \phi\beta\sqrt{v})\frac{\partial V}{\partial v}, \end{aligned}$$

where we have written the arbitrary function of S, v, t as $(\alpha - \phi\beta\sqrt{v})$, where α and β are the drift and volatility functions from the SDE (1).

Chapter 1: 'Stochastic Volatility (SV) (Market Price of Volatility Risk)

Function $\phi(S, v, t)$ is called *market price of volatility risk*. To see why, we follow Wilmott's argument.

Consider the portfolio Π_1 consisting of a delta-hedged option V . Then

$$\Pi_1 = V - \frac{\partial V}{\partial S} S$$

and applying Itô's lemma and taking into account of delta-hedged option (the coefficient of dS is zero), we have

$$d\Pi_1 - r\Pi_1 dt = \beta\sqrt{v}\frac{\partial V}{\partial v}[\phi dt + dZ_2].$$

Chapter 1: 'Stochastic Volatility (SV) (Market Price of Volatility Risk)

We see that the extra return per unit of volatility risk dZ_2 is given by ϕdt and so in analogy with the CAPM, ϕ is known as the *market price of volatility risk*.

Defining the *risk-neutral drift* as

$$\alpha' = \alpha - \beta\sqrt{v}\phi$$

we see that, as far as pricing of options is concerned, we could have started with the risk-neutral SDE for v ,

$$dv = \alpha' dt + \beta\sqrt{v}dZ_2$$

and got identical results with no explicit price of risk term because we are in the risk neutral world.

Chapter 1: 'Stochastic Volatility (SV) (Risk-Neutral World)

In what follows, we always assume that the SDEs for S and v are in risk-neutral terms because we are invariably interested in fitting models to option prices.

Effectively, we assume that we are imputing the risk-neutral measure directly by fitting the parameters of the process that we are imposing.

Chapter 1: Local Volatility (LV), History

Given the computational complexity of SV models and difficulty of fitting parameters to the current prices of vanilla options, practitioners sought a simpler way of pricing exotic options consistently with the volatility skew.

Since before Breeden and Litzenberger (1978), it was understood that the risk-neutral density could be derived from the market prices of European options.

Chapter 1: Local Volatility (LV), History

The breakthrough came when Dupire (1994) (continuous time version) and Derman and Kani (1994) (discrete time binomial tree version) noted that under risk neutrality, there was a unique diffusion process consistent with these distributions.

The corresponding unique state-dependent diffusion coefficient $\sigma_L(S, t)$, consistent with current European option prices, is known as the *local volatility function*.

Chapter 1: Local Volatility (LV), History

As if any proof were needed, Dumas, Fleming and Whaley (1998) performed an empirical analysis that confirmed that the dynamics of the implied volatility surface were not consistent with the assumption of constant local volatilities.

Later on, we show that LV is indeed an average over instantaneous volatilities, formalizing the intuition of those practitioners who first introduced the concept.

Chapter 1: Local Volatility (LV): A Brief Review of Dupire's Work

For a given expiration T and current stock price S_0 , the collection $C(S_0, K, T)$ of undiscounted option prices of different strikes yields the risk-neutral density function ϕ of the final spot S_T through the relationship

$$C(S_0, K, T) = \int_K^{\infty} \phi(S_T, T; S_0) dS_T (S_T - K).$$

Chapter 1: Local Volatility (LV): A Brief Review of Dupire's Work

Differentiate this twice wrt K to obtain

$$\phi(S_T, T; S_0) = \frac{\partial^2 C}{\partial K^2}$$

so the Arrow-Debreu prices for each expiration may be recovered by twice differentiating the undiscounted option price wrt K .

Chapter 1: Local Volatility (LV): A Brief Review of Dupire's Work

Given the distribution of final spot prices S_T for each time T conditional on some starting spot price S_0 , Dupire shows that there is a unique risk neutral diffusion process which generates these distributions.

That is, given the set of all European option prices, we may determine the functional form of the diffusion parameter (LV) of the unique risk neutral diffusion process which generates these prices.

Chapter 1: Local Volatility (LV): A Brief Review of Dupire's Work

Let us write the process as

$$\frac{dS}{S} = \mu_t dt + \sigma(S_t, t; S_0) dZ.$$

Application of Itô lemma together with risk neutrality, gives rise to a partial DE for functions of the stock price, which is a straightforward generalization of B-S. In particular, the pseudo-probability densities $\phi = \frac{\partial^2 C}{\partial K^2}$ must satisfy the Fokker-Planck equation.

Chapter 1: Local Volatility (LV): A Brief Review of Dupire's Work

This leads to the following equation for the undiscounted option price C in terms of the strike price:

$$\frac{\partial C}{\partial T} = \frac{\sigma^2 K^2}{2} \frac{\partial^2 C}{\partial K^2} + (r_t - D_t) \left(C - K \frac{\partial C}{\partial K} \right),$$

where r_t is the risk-free rate, D_t is the dividend yield and C is short for $C(S_0, K, T)$.

Chapter 1: Local Volatility (LV): Derivation of the Dupire Equation

Suppose the stock price diffuses with risk-neutral drift $\mu_t = r_t - D_t$ and LV $\sigma(S, t)$ according to the equation:

$$\frac{dS}{S} = \mu_t dt + \sigma(S_t, t) dZ.$$

The undiscounted risk-neutral value $C(S_0, K, T)$ of a European option with strike K and expiration T is given by

$$C(S_0, K, T) = \int_K^{\infty} \phi(S_T, T; S_0) dS_T (S_T - K).$$

Chapter 1: Local Volatility (LV): Derivation of the Dupire Equation

Here, ϕ is the pseudo-probability density of the final spot time T .

It evolves according to the Fokker-Planck equation:

$$\frac{1}{2} \frac{\partial^2}{\partial S_T^2} (\sigma^2 S_T^2) - S \frac{\partial(\mu S_T \phi)}{\partial S_T} = \frac{\partial \phi}{\partial T}.$$

Chapter 1: Local Volatility (LV): Derivation of the Dupire Equation

Differentiating wrt K gives

$$\frac{\partial C}{\partial K} = - \int_K^{+\infty} dS_T \phi(S_T, T; S_0),$$

and

$$\frac{\partial^2 C}{\partial K^2} = \phi.$$

Chapter 1: Local Volatility (LV): Derivation of the Dupire Equation

Differentiating C wrt T (and taking into account above $\frac{\partial \phi}{\partial T}$) gives

$$\frac{\partial C}{\partial T} = \int_K^{+\infty} dS_T \left\{ \frac{1}{2} \frac{\partial^2}{\partial S_T^2} (\sigma^2 S_T^2) - S \frac{\partial(\mu S_T \phi)}{\partial S_T} \right\} (S_T - K).$$

Integrating by parts twice gives:

$$\frac{\partial C}{\partial T} = \frac{\sigma^2 K^2}{2} \frac{\partial^2 C}{\partial K^2} + \mu(T) \left(-K \frac{\partial C}{\partial K} \right)$$

which is the *Dupire equation* when the underlying stock has risk-neutral drift μ .

Chapter 1: Local Volatility (LV): Derivation of the Dupire Equation

That is, the forward price of the stock at time T is given by

$$F_T = S_0 \exp\left(\int_0^T \mu_t dt\right).$$

Were we to express the option price as a function of the forward price F_T , we would get the same expression minus the drift term.

That is,

$$\frac{\partial C}{\partial T} = \frac{1}{2} \sigma^2 K^2 \frac{\partial^2 C}{\partial K^2},$$

where $C = C(F_T, K, T)$.

Chapter 1: Local Volatility (LV): Derivation of the Dupire Equation

Inverting this gives

$$\sigma^2(K, T, S_0) = \frac{\frac{\partial C}{\partial T}}{\frac{1}{2}K^2 \frac{\partial^2 C}{\partial K^2}}$$

The right-hand side of the last equation can be computed from known European option prices. So, given a complete set of European option prices for all strikes and expirations, LVs are given by the last equation.

We can view this equation as a definition of the LV function regardless of what kind of process (SV, e.g.) actually governs the evolution of volatility.

Chapter 1: Local Volatility (LV): LV in Terms of Implied Volatility (IV)

Market prices of options are quoted in terms of B-S IV $\sigma_{BS}(K, T; S_0)$. In other words, we may write

$$C(S_0, K, T) = C_{BS}(S_0, K, \sigma_{BS}(K, T; S_0), T).$$

More convenient to work in terms of two dimensionless variables: the B-S implied total variance w defined by

$$w(S_0, K, T) := \sigma_{BS}^2(K, T; S_0)T$$

and the log-strike y defined by

$$y = \log\left(\frac{K}{F_T}\right),$$

where $F_T = S_0 \exp(\int_0^T \mu_t dt)$ gives the forward price of the stock at time 0.

Chapter 1: Local Volatility (LV): LV in Terms of Implied Volatility (IV)

In terms of these variables, the B-S formula for the future value of the option price becomes

$$\begin{aligned} C_{BS}(F_T, y, w) &= F_T [N(d_1) - e^y N(d_2)] \\ &= F_T \left[N\left(-\frac{y}{\sqrt{w}} + \frac{\sqrt{w}}{2}\right) - e^y N\left(-\frac{y}{\sqrt{w}} - \frac{\sqrt{w}}{2}\right) \right] \end{aligned}$$

and the Dupire equation becomes

$$\frac{\partial C}{\partial T} = \frac{v_L}{2} \left[\frac{\partial^2 C}{\partial y^2} - \frac{\partial C}{\partial y} \right] + \mu(T)C$$

with $v_L = \sigma^2(S_0, K, T)$ representing the LV.

Chapter 1: Local Volatility (LV): LV in Terms of Implied Volatility (IV)

By taking derivatives of the B-S formula wrt w (twice), yw (mixed), y (twice), transforming the Dupire equation into an equation in terms of IV (by making substitutions for derivatives C wrt y , y^2 twice and T), we get the equation for $\frac{\partial C_{BS}}{\partial w} \frac{\partial w}{\partial T}$. Taking out a factor of $\frac{\partial C_{BS}}{\partial w}$ and simplifying, we get $\frac{\partial w}{\partial T}$ and inverting this gives the final result for LV:

$$v_L = \frac{\frac{\partial w}{\partial T}}{1 - \frac{y}{w} \frac{\partial w}{\partial y} + \frac{1}{4} \left(-\frac{1}{4} - \frac{1}{w} + \frac{y^2}{w^2} \right) \left(\frac{\partial w}{\partial y} \right)^2 + \frac{1}{2} \frac{\partial^2 w}{\partial y^2}}$$

Chapter 1: Local Volatility (LV): LV in Terms of Implied Volatility (IV): Special Case: No Skew

If the skew $\frac{\partial w}{\partial y}$ is zero, we must have (see previous slide)

$$v_L = \frac{\partial w}{\partial T}.$$

So the LV in this case reduces to the forward B-S IV. The solution to this is

$$w(T) = \int_0^T v_L(t) dt.$$

Chapter 1: Local Volatility (LV): LV as a Conditional Expectation of Instantaneous Variance

This result was originally independently derived by Dupire (1996) and Derman and Kani (1998). Following now the elegant derivation by Derman and Kani, assume the same stochastic process for the stock price $dS_t = \mu_t S_t dt + \sqrt{v_t} S_t dZ$ but write it in terms of the forward price $F_{t,T} = S_t \exp[\int_t^T \mu_s ds]$:

$$dF_{t,T} = \sqrt{v_t} F_{t,T} dZ.$$

Note that $dF_{T,T} = dS_T$.

Chapter 1: Local Volatility (LV): LV as a Conditional Expectation of Instantaneous Variance

The undiscounted value of a European option with strike K expiring at time T is given by

$$C(S_0, K, T) = E_Q[(S_T - K)^+].$$

Differentiating once wrt K gives $\frac{\partial C}{\partial K} = -E_Q[\theta(S_T - K)]$, where θ is the Heaviside function. Differentiating again wrt K gives $\frac{\partial^2 C}{\partial K^2} = E_Q[\delta(S_T - K)]$, where δ is the Dirac δ function. Now a formal application of Itô's lemma to the terminal payoff of the option (and using $dF_{T,T} = dS_T$) gives the identity

$$d(S_T - K)^+ = \theta(S_T - K)dS_T + \frac{1}{2}v_L S_T^2 \delta(S_T - K)dT.$$

Chapter 1: Local Volatility (LV): LV as a Conditional Expectation of Instantaneous Variance

Taking conditional expectations of each side, and using the fact that $F_{t,T}$ is a martingale, we get

$$dC = dE_Q[(S_T - K)^+] = \frac{1}{2}E_Q[v_L S_T^2 \delta(S_T - K)]dT.$$

Also, we can write

$$\begin{aligned} E_Q[v_L S_T^2 \delta(S_T - K)] &= E_Q[V_L | S_T = K] \frac{1}{2} K^2 E_Q[\delta(S_T - K)] \\ &= E_Q[V_L | S_T = K] \frac{1}{2} K^2 \frac{\partial^2 C}{\partial K^2}. \end{aligned}$$

Chapter 1: Local Volatility (LV): LV as a Conditional Expectation of Instantaneous Variance

Putting this together, we get

$$\frac{\partial C}{\partial T} = E_Q[v_L | S_T = K] \frac{1}{2} K^2 \frac{\partial^2 C}{\partial K^2}.$$

Comparing this with the definition of LV, we see that

$$\sigma^2(S_0, K, T) = E_Q[v_L | S_T = K].$$

That is, LV is the risk-neutral expectation of the instantaneous variance conditional on the final stock price S_T being equal to the strike price K .

Conclusion

1. Introduction to the Book
2. Contents of the Book
3. Chapter 1: Stochastic Volatility and Local Volatility

The End of Chapter 1

Thank You for Your Time and Attention!



UNIVERSITY OF
CALGARY