



THE BUSINESS SCHOOL
FOR FINANCIAL MARKETS



Stochastic Local Volatility

CAROL ALEXANDER

ICMA Centre, University of Reading

LEONARDO M. NOGUEIRA

ICMA Centre, University of Reading and Banco Central do Brasil

First version: September 2004

This version: March 2008

ICMA Centre Discussion Papers in Finance DP2008-02

Copyright 2008 Alexander and Nogueira All rights reserved.

ICMA Centre • The University of Reading
Whiteknights • PO Box 242 • Reading RG6 6BA • UK
Tel: +44 (0)1183 787402 • Fax: +44 (0)1189 314741
Web: www.icmacentre.rdg.ac.uk
Director: Professor John Board, Chair in Finance
The ICMA Centre is supported by the International Capital Market Association



ABSTRACT

There are two unique volatility surfaces associated with any arbitrage-free set of standard European option prices, the implied volatility surface and the local volatility surface. Several papers have discussed the stochastic differential equations for implied volatilities that are consistent with these option prices but the static and dynamic no-arbitrage conditions are complex, mainly due to the large (or even infinite) dimensions of the state probability space. These no-arbitrage conditions are also instrument-specific and have been specified for some simple classes of options. However, the problem is easier to resolve when we specify stochastic differential equations for local volatilities instead. And the option prices and hedge ratios that are obtained by making local volatility stochastic are identical to those obtained by making instantaneous volatility or implied volatility stochastic. After proving that there is a one-to-one correspondence between the stochastic implied volatility and stochastic local volatility approaches, we derive a simple dynamic no-arbitrage condition for the stochastic local volatility model that is model-specific. The condition is very easy to check in local volatility models having only a few stochastic parameters.

JEL Classification: G13, C16

Keywords: Local volatility, stochastic volatility, unified theory of volatility, local volatility dynamics.

Corresponding Address:

ICMA Centre, University of Reading, PO Box 242, Reading, United Kingdom, RG6 6BA
Tel: +44 (0)1183 786675 Fax: +44 (0)1189 314741
E-mails: c.alexander@icmacentre.rdg.ac.uk and l.nogueira@icmacentre.rdg.ac.uk

Acknowledgements: We are most grateful for very useful comments on an earlier draft from Bruno Dupire of Bloomberg, New York; Hyungsok Ahn of Nomura Bank, London; Emanuel Derman of Columbia University, New York; Jacques Pézier of the ICMA Centre; and Nicole Branger of the University of Münster.

I. INTRODUCTION

Option pricing models seek to price and hedge exotic and path-dependent options consistently with the market prices of simple European calls and puts. Two main strands of research have been developed in a prolific literature: stochastic volatility as in Hull and White (1987), Stein and Stein (1991), Heston (1993) and many others, where the variance or volatility of the price process is stochastic; and local volatility as in Dupire (1994), Derman and Kani (1994) and Rubinstein (1994) where arbitrage-free forward volatilities are ‘locked-in’ by trading options today. The term ‘local volatility’ has subsequently been extended to cover any deterministic volatility model where forward volatilities are a function of time and asset price. For a detailed review of these models, see Jackwerth (1999), Skiadopoulos (2001), Bates (2003), Psychoyios et al. (2003) and Nogueira (2006).

Stochastic and local volatility models have been regarded as alternative and competing approaches to the same unobservable quantity, the instantaneous volatility of the underlying asset. But it is only when one takes a restricted view of volatility dynamics that they appear to be different. Under more general assumptions the two approaches yield identical claim prices and hedge ratios. This paper unifies the two approaches into a single theory by proving that a local volatility model with stochastic parameters yields identical implied volatility dynamics to those of Schönbucher’s (1999) market model of stochastic implied volatilities. We thus provide a model-based proof of the unified theory of volatility of Dupire (1996).

By modeling parameter dynamics under no-arbitrage conditions we are also implicitly modeling the dynamics of the local volatility surface, the implied volatility surface and the implied probability distribution. Hence our approach can be regarded as an alternative to directly modeling the dynamics of local volatilities (Dupire, 1996, Derman and Kani, 1998), implied volatilities (Schönbucher, 1999, Brace et al., 2001, Ledoit et al., 2002) and implied probabilities (Panigirtzoglou and Skiadopoulos, 2004).

Several recent papers examine no-arbitrage conditions in the stochastic implied volatility framework. By defining a measure under which the underlying price and the implied volatilities are martingales, Zilber (2007) applies the fundamental theorem of arbitrage to provide a constructive proof of the existence of market models of implied volatilities that admit neither static nor dynamic arbitrage. The absence of dynamic arbitrage places constraints on the state probability space, and further complex conditions for absence of static arbitrage are required, as specified by Carr and Madan (2005) or Davis and Hobson (2007). However, Schweizer and Wissel transform the market

model for implied volatilities into a market model for forward implied volatilities (Schweizer and Wissel, 2008) or for local implied volatilities (Schweizer and Wissel, 2007) and provided these are non-negative there can be no static arbitrage. This transformation simplifies the problem, and they show how absence of dynamic arbitrage yields restrictions on the drift terms in the forward implied volatility model. These conditions are analogous to the drift restrictions in the HJM market model for interest rates (Heath et al, 1992). For a particular choice of volatility coefficients, an arbitrage-free model only exists if the drift coefficients, which depend on the volatility coefficients, satisfy the drift conditions. The drift condition is specified for some particular instruments such as standard call options, power options and the log contract. However, generalizing these results to derive drift restrictions when implied volatilities with different maturities *and* strikes are modeled in the same framework is a complex problem because the admissible state space has a complicated structure.

By contrast, the derivation of no-arbitrage conditions in a stochastic local volatility framework is relatively straightforward, because the static no-arbitrage condition is simply that local volatilities are non-negative and the admissible state space is relatively simple.¹ In this paper we derive a single drift condition for the stochastic local volatility model to be arbitrage-free, and specify the condition for some specific examples of stochastic local volatility models. Recently a similar approach has been considered by Wissel (2007), who defines ‘local implied volatilities’ as the discrete analogue of the local volatility function derived from arbitrage-free prices of standard European options given by Dupire’s equation (Dupire, 1994). The main difference between this and our approach is that we specify stochastic dynamics for the *parameters* of the local volatility function, whereas Wissel uses local implied volatilities to parameterize European option prices and then specifies stochastic dynamics for these volatilities directly.

The remainder of this paper is as follows: Section II reviews the inter-relationship between local volatility and stochastic volatility. Section III introduces the stochastic local volatility (SLV) model, derives the no-arbitrage drift condition and examines this condition for some specific stochastic local volatility models; Section IV derives the implied variance, implied probability and implied volatility dynamics that are consistent with the SLV claim price dynamics. In particular it proves the duality between the SLV model and the market model of implied volatilities introduced by Schönbucher (1999); Section V summarizes and concludes.

¹ Also, whereas implied volatilities are instrument-specific, being defined as the inverse price of a standard European option, a local volatility surface can be defined without reference to tradable instruments.

II. UNIFIED THEORY OF VOLATILITY

This section provides a non-technical summary of the theory of unified volatility that was first introduced by Dupire (1996), and a formal introduction to the concept of stochastic local volatility.

II.1 Gyöngy's theorem

The local volatility model can be understood from the early work of Gyöngy (1986). Suppose that X_t is a real-valued one-dimensional Itô process starting at $X_0 = 0$ with dynamics:

$$dX_t = \beta(t, \omega)dt + \zeta(t, \omega)d\mathbf{B}_t \quad (1)$$

where $(\mathbf{B}_t)_{k \times 1}$ is a k -dimensional Wiener process on the probability space $(\Omega, \mathfrak{F}, P)$, the possibly random coefficients $\beta(t, \omega)$ and $(\zeta(t, \omega))_{1 \times k}$ satisfy the regularity conditions of an Itô process and $\omega \in \Omega$ denotes dependence on some arbitrary variables. In particular, suppose that $\zeta\zeta^T(t, \omega)$ is positive. Gyöngy (1986) proved that there exists another stochastic process \tilde{X}_t which is a solution of the stochastic differential equation:

$$d\tilde{X}_t = b(t, \tilde{X}_t)dt + v(t, \tilde{X}_t)d\tilde{\mathbf{B}}_t, \quad \tilde{X}_0 = 0 \quad (2)$$

with *non-random* coefficients b and v defined by:

$$\begin{aligned} b(t, x) &\stackrel{\text{def}}{=} E_p \left[\beta(t, \omega) \mid X_t = x \right] \\ v(t, x) &\stackrel{\text{def}}{=} \left(E_p \left[\zeta\zeta^T(t, \omega) \mid X_t = x \right] \right)^{\frac{1}{2}} \end{aligned} \quad (3)$$

and which admits the *same* marginal probability distribution as that of X_t for every $t > 0$. That is, for every Itô process of the type (1) there is a deterministic process (2) that ‘mimics’ the marginal distribution of X_t for every t .

Now we establish the link between Gyöngy's result and local volatility. Suppose $k = 1$ and define $X_t = \ln(S_t/S_0)$, $\beta(t, \omega) = \mu - \frac{1}{2}\sigma^2(t, \omega)$ and $\zeta(t, \omega) = \sigma(t, \omega)$. Then (1) becomes:

$$d \ln S_t = \left(\mu - \frac{1}{2}\sigma^2(t, \omega) \right) dt + \sigma(t, \omega) d\mathbf{B}_t$$

and by Itô's lemma:

$$\frac{dS_t}{S_t} = \mu dt + \sigma(t, \omega) d\mathbf{B}_t \quad (4)$$

which is the stochastic differential equation for a financial asset S_t with possibly stochastic volatility.

Denote by S an arbitrary realization of S_t for some $t \geq 0$ and put $x = \ln(S/S_0)$. Using (3) we have:

$$v^2(t, x) = v^2(t, \ln(S/S_0)) = E_p \left[\sigma^2(t, \omega) \middle| S_t = S \right] \stackrel{\text{def}}{=} \sigma_L^2(t, S)$$

$$b(t, x) = b(t, \ln(S/S_0)) = E_p \left[\mu - \frac{1}{2} \sigma^2(t, \omega) \middle| S_t = S \right] = \mu - \frac{1}{2} \sigma_L^2(t, S)$$

Replacing these into (2) with $\tilde{X}_t = \ln(\tilde{S}_t/\tilde{S}_0)$, $\tilde{S}_0 = S_0$ and using Itô's lemma, we obtain:

$$\frac{d\tilde{S}_t}{\tilde{S}_t} = \mu dt + \sigma_L(t, \tilde{S}_t) d\tilde{B}_t \quad (5)$$

which is the stochastic differential equation of \tilde{S}_t with the *deterministic* local volatility $\sigma_L(t, S)$. Thus \tilde{S}_t in the local volatility model (5) has the same marginal distribution as S_t in (4) for every t . Besides, as there is a one-to-one relationship between risk-neutral marginal probabilities and the prices of standard European options (Breedon and Litzenberger, 1978), both models (4) and (5) produce the same prices for simple calls and puts after a measure change from P to the risk-neutral measure.

II.2 Dupire's equation

For every arbitrage-free model of the form (4) there is a corresponding local volatility model, such as (5). In particular, Dupire (1994) showed that the local volatility in (5) is unique and is also given by the square root of

$$\sigma_L^2(t, S) \Big|_{t=T, S=K} = 2 \left(\frac{\partial f_0}{\partial T} + (r - q) K \frac{\partial f_0}{\partial K} + q f_0 \right) \Big/ K^2 \frac{\partial^2 f_0}{\partial K^2} \quad (6)$$

where $f_0 = f_0(K, T)$ is the price of a standard European option with strike K and maturity T at time $t = 0$ when the underlying asset price is S_0 and when the local volatility is calibrated; r and q denote respectively the risk-free interest rate and the dividend yield for equities (or their counterparts for other markets) both assumed constant and continuously-compounded.

Local volatility models are widely used in practice because they enable fast and accurate pricing of exotic claims when only marginal distributions are required. However, it would be a mistake to interpret local volatility as a complete representation of the true stochastic process driving the underlying asset price. Local volatility is merely a simplification that is practically useful for describing a price process with non-constant volatility. More precisely, although the marginal distributions are the same at the time when the local volatility is calibrated, clearly S_t and \tilde{S}_t do not follow the same dynamics (cf. (4) and (5)) hence options prices will have different dynamics under each model, and hedge ratios can differ substantially. Therefore, for pricing and hedging some

exotic options, the perfect fit to the current vanilla options prices provided by the local volatility theory may not be enough and a deeper understanding of the dynamics of S_t is required.

For instance, if the ‘true’ risk-neutral price dynamics followed a stochastic volatility model such as:

$$\begin{aligned} \frac{dS_t}{S_t} &= (r - q)dt + \sigma(Y_t)dB_t \\ dY_t &= \alpha(t, Y_t)dt + \beta(t, Y_t)dW_t \quad d\langle B, W \rangle_t = \rho dt \end{aligned} \quad (7)$$

the same local volatility function would be obtained from both Dupire’s equation (6) and:

$$\sigma_L^2(t, S) = E\left[\sigma^2(Y_t) \mid S_t = S\right] \quad (8)$$

where the expectation is taken under the risk-neutral measure. See Dupire (1996) and Derman and Kani (1998) for the proof of this result. But given that $\sigma(Y_t)$ is stochastic, the local volatility will also be stochastic.

II.3 Towards a theory of stochastic local volatility

Stochastic volatility, which indeed appeared in the literature before local volatility, is based on empirical evidence that volatility displays characteristics of a stochastic process on its own.² There is vast evidence, mostly from S&P 500 index options, that option prices are driven by at least one random factor other than the asset price. For instance, Bakshi *et al.* (2000) find that intraday option prices do not always follow the movements of the underlying price, which is inconsistent with the perfect correlation from deterministic volatility models. Buraschi and Jackwerth (2001) find that one cannot span the pricing kernel using only two assets, and so reject deterministic volatility models in favour of models that incorporate additional risk factors. Coval and Shumway (2001) find abnormal returns for calls and puts and conclude that this should be caused by at least one additional risk factor.

A typical stochastic volatility model defines the price process as in (7), where the instantaneous volatility σ is a function of the stochastic process Y_t , possibly correlated with the underlying price. For instance, Hull and White (1987) define $\sigma(Y_t) = \sqrt{Y_t}$ where Y_t is lognormally distributed but their model does not allow for mean reversion, so volatility can grow indefinitely. In the models by Scott (1987) and Stein and Stein (1991) Y_t is mean reverting but it is also normal and can go negative if mean reversion is not strong. Hence $\sigma(Y_t)$ is conveniently chosen to map a real number into a positive number. In the strong GARCH diffusion of Nelson (1990) Y_t follows a geometric,

² See e.g. Schwert (1989), Ghysels et al. (1996), Fouque et al. (2000), Psychoyios et al. (2003) and Gatheral (2005).
Copyright 2008. Alexander and Nogueira. All Rights Reserved.

mean-reverting process and in the Heston (1993) and Ball and Roma (1994) models, Y_t has a non-central chi-square distribution and is positive and mean reverting. In all the models above with the exception of Heston's, the correlation between price and volatility is zero, although extensions to correlated volatility processes have been derived in many cases. For a review of typical stochastic volatility processes and their distributional assumptions, see Psychoyios *et al.* (2003).

Stochastic volatility models have the drawback that the volatility process itself is not observable. It can only be estimated from historical data or calibrated to current option prices. But, either way the resulting process is still dependent on the particular structure assumed for $\sigma(Y_t)$ and Y_t . Thus, supported by the widespread use of the Black-Scholes model, several authors have instead modelled the behaviour of the (observable) Black-Scholes implied volatility over time and verified empirically the existence of multiple risk factors driving option prices. See Skiadopoulos *et al.* (1999), Alexander (2001), Cont and da Fonseca (2002), Cont *et al.* (2002), Fengler *et al.* (2003), Hafner (2004) and Fengler (2005).

These findings motivated the class of 'stochastic implied volatility' models, in which a stochastic process is explicitly assumed for the Black-Scholes implied volatility of vanilla options and used to price and hedge exotic options, with the stochastic process for the instantaneous volatility following from no-arbitrage arguments. In fact, since for any vanilla option there is one and only one implied volatility that is consistent with the option price, it follows that implied volatilities are as observable as option prices, so that standard econometric techniques may be used to search for representations of volatility that are both tractable and realistic. Well-known members of this class include the models of Schönbucher (1999), Ledoit and Santa-Clara (1998) and Brace *et al.* (2001).

But the Black-Scholes implied volatility is not the only form of volatility that is observable. In particular, for any complete set of arbitrage-free prices of vanilla calls and puts, Dupire (1994) has shown that there is a unique local volatility surface and that this surface can be computed directly from option prices using equation (6) above. Can the dynamics of local volatilities also be modelled directly and used to derive arbitrage-free prices and hedge ratios for exotic options that are consistent with the evolution of local volatility over time? Derman and Kani (1998) have addressed this question and proposed a 'stochastic local volatility' model based on stochastic implied trees. They used trees because, when working in continuous time, one has to deal with a challenging no-arbitrage condition for the drift of the local volatility process.

In this paper we derive a continuous stochastic local volatility model that is tractable, intuitive, arbitrage-free and that helps explain the alleged ‘instability’ of local volatility surfaces over time. As we shall see, this is achieved by assuming a parametric functional form for local volatility and modelling the time evolution of the local volatility parameters under no-arbitrage conditions. Our approach is analogous to that of a stochastic implied volatility (SIV) model. In a typical SIV model, one uses the Black-Scholes model to compute the implied volatility of vanilla options for each date in the sample and applies standard econometric methods, such as principal component analysis and ARIMA models, to produce a realistic dynamics for implied volatilities. Likewise, in the stochastic local volatility (SLV) model described in the next section we assume that a parametric local volatility is calibrated to the observed prices of vanilla options at each date in the sample and, again, may apply simple econometric methods to produce realistic dynamics for the local volatility parameters.

In fact, these models are more than analogous: we show that they are equivalent. That is, we prove that, under the assumption that the SLV model holds, one can derive the same implied volatility dynamics as observed for SIV models. This is an important result but it is not really a surprise. Since implied volatilities and local volatilities are derived from the same option prices, their dynamics should be consistent if the market is arbitrage-free. Nevertheless, it is still worthwhile checking that our theory is indeed self-consistent.

The SLV model is not a mere special case of the class of SIV models. We prove a theorem which shows that verifying a single drift condition is enough to guarantee absence of arbitrage in the SLV model. By contrast, the SIV model imposes a drift condition whose coefficients depend on the particular option strike and maturity, and verifying that these coefficients are consistent across different strikes and maturities is not straightforward, as Schweizer and Wissel (2008) have demonstrated.

II.4 Stochastic local volatility

This subsection formalizes the concept of stochastic local volatility and provides further motivation for the model derived in Section III. Following Gyöngy (1986), Dupire (1994) and our discussion above, it is always possible to define local volatility dynamics that mimic the marginal probability distribution of the underlying asset price. These dynamics, under the risk-neutral measure, could be written as:

$$\frac{d\tilde{S}_t}{\tilde{S}_t} = (r - q)dt + \sigma_L(t, \tilde{S}_t)d\tilde{B}_t \quad (9)$$

with r and q as above and, in the absence of arbitrage, the local volatility function $\sigma_L(t, S)$ uniquely determined according to (6).

However, since finding $\sigma_L(t, S)$ requires a continuum of traded option prices, direct computation of the local volatility function is problematic.³ The local volatility surface can be very irregular and sensitive to the interpolation methods used between quoted option prices and their extrapolation to boundary values (see e.g. Bouchouev and Isakov (1997, 1999) and Avellaneda et al. (1997)). Consequently most recent work on local volatility has introduced a variety of parametric forms for local volatility functions in which the parameters

$$\boldsymbol{\lambda}_t = (\lambda^1, \lambda^2, \dots, \lambda^n)'$$

are calibrated to the market prices of vanilla options at every t . If $\boldsymbol{\lambda}_0$ denotes the calibrated values of the parameters at time 0 the underlying asset price process assumed at this time is:

$$\frac{d\tilde{S}_t}{\tilde{S}_t} = (r - q)dt + \sigma_L(t, \tilde{S}_t, \boldsymbol{\lambda}_0) d\tilde{B}_t \quad (10)$$

Then, when the model is re-calibrated at time $t_1 > t_0$, we have:

$$\frac{d\tilde{S}_t}{\tilde{S}_t} = (r - q)dt + \sigma_L(t, \tilde{S}_t, \boldsymbol{\lambda}_1) d\tilde{B}_t \quad (11)$$

and this will differ from (10) unless $\boldsymbol{\lambda}_1 = \boldsymbol{\lambda}_0$. Whilst local volatility models assume parameters are constant, in practice parameters are not constant and change each time the model is re-calibrated. This is because (10) is not necessarily the true model despite the perfect fit to options prices it could achieve at time 0.

To summarize, local volatility models assume a deterministic instantaneous volatility for the price process and this implies a static local volatility.⁴ That is, the forward volatilities are obtained by calibration to current market prices and in the local volatility framework they should be realized with certainty. Yet this assumption is not necessary. Both Dupire (1996) and Derman and Kani (1998) recognize this and thus define the local variance as the conditional expectation of a stochastic variance, such as:

$$\sigma_L^2(t, S) = E_Q \left[\sigma^2(t, S_t, \boldsymbol{\omega}) \mid S_t = S, \mathfrak{F}_0 \right] \quad \text{at time 0} \quad (12)$$

³ There is a version of Dupire's equation using implied volatilities that is more stable. See e.g. Fengler (2005, ch. 3).

⁴ We assume that the local volatility is not a function of the current asset price S_0 otherwise it cannot be static.

where E_Q denotes the expectation under the implied risk-neutral probability Q , i.e. the distribution that is consistent with current prices of European puts and calls (see Jackwerth, 1999). The filtration \mathfrak{F}_0 includes all information up to time 0 and ω denotes all sources of uncertainty that may influence the instantaneous volatility process other than the asset price. The instantaneous variance in (12) is very general because ω can be any arbitrage-free set of continuous stochastic processes. In particular, this definition of local variance is consistent with any univariate diffusion stochastic volatility model. For this reason, Dupire (1996) named (12) the ‘unified theory of volatility’.

Note that taking the expectation in (12) ignores the residual uncertainty from ω and its influence on the instantaneous volatility. This uncertainty is transferred to the local volatility surface itself. That is, although locally (i.e. at each calibration) the surface is indeed a deterministic function of t and S , that surface has stochastic dynamics. This explains why the local volatility surface can be very unstable on re-calibration: the uncertainty from ω does not just disappear from the model.

III. STOCHASTIC LOCAL VOLATILITY (SLV) MODEL

III.1 Model specification

Suppose the true asset price process follows a geometric Brownian motion with some arbitrary stochastic volatility under the risk-neutral measure:

$$\frac{dS_t}{S_t} = (r - q)dt + \sigma(t, S_t, \omega)dB_t \quad (13)$$

The risk-free rate r and dividend yield q are assumed constant and the volatility process is, for the moment, only assumed to be bounded and continuous so that S_t is a valid Itô process. Now assume that a local volatility model is calibrated at each time $u \geq 0$ and produces the risk-neutral dynamics for all $t \geq u$ given by:

$$\frac{d\tilde{S}_t}{\tilde{S}_t} = (r - q)dt + \sigma_L(t, \tilde{S}_t, \lambda_u)d\tilde{B}_t \quad (14)$$

Hence, all the uncertainty in ω is assumed to be captured by the dynamics of the parameters of the local volatility model.⁵ This way, the local volatility is regarded as a function of t , S_t and the stochastic process λ_u of parameters that are not fixed until they are calibrated to the market at some future time $u \geq 0$. With this definition the local volatility function has an implicit dependence on the future values of the parameters so that the local volatility becomes stochastic. The standard

⁵ This is not a strong assumption. It is analogous to assuming that implied volatility dynamics capture all imperfections of the Black-Scholes model. Bates (2003) argues that daily re-calibration is a way of hiding the imperfections of a model so that it always fits the data. These imperfections would then express themselves in the dynamics of parameters.

local volatility model can be viewed as a restricted form of this model in the special case that the parameters are constant.

Next, assume that the risk-neutral dynamics for each parameter in (14) are as follows:

$$d\lambda_t^i = \alpha_i(t, \lambda_t^i)dt + \beta_i(t, \lambda_t^i)dW_t^i \quad (15)$$

$$dW_t^i = \varrho_{i,S}(t, S_t, \lambda_t^i)dB_t + \sqrt{1 - \varrho_{i,S}^2(t, S_t, \lambda_t^i)}dZ_t^i \quad (16)$$

where

$$d\langle W^i, W^j \rangle_t = \varrho_{i,j}(t, \lambda_t^i, \lambda_t^j)dt \quad d\langle Z^i, B \rangle_t = 0 \quad (17)$$

for $i, j \in \{1, 2, \dots, n\}$. Here $\varrho_{i,j}$ is the correlation between variations in λ_t^i and λ_t^j , and $\varrho_{i,S}$ is the correlation between variations in λ_t^i and S_t respectively. The dynamics (15) are also assumed to satisfy the regularity conditions for an Itô process.⁶ We call the model defined by (13) – (17) the stochastic local volatility (SLV) model.

By taking the limit of (12) when $t \rightarrow 0$ in the context of dynamics (13) and (14), it is easy to show that $\sigma_L(0, S_0, \lambda_0) = \sigma(0, S_0, \omega)$, so that the instantaneous volatility at time 0 is uniquely given by the local volatility calibrated at time 0, when the asset price is S_0 and the parameters are λ_0 . This result can be generalized to any calibration time $u \geq 0$ as $\sigma_L(u, S_u, \lambda_u) = \sigma(u, S_u, \omega)$. Thus, since the local volatility model (14) is calibrated for all u , the instantaneous volatility in (13) at any time t is easily obtained from the (possibly analytical) local volatility function by setting:

$$\sigma(t, S_t, \omega) = \sigma_L(t, S_t, \lambda_t). \quad (18)$$

Since the value of λ_t is discovered only when the local volatility model is calibrated at time t , it follows that $\sigma(t, S_t, \omega)$ is stochastic, as it should be by definition.

III.2 Claim price dynamics and hedge ratios

We now derive the price dynamics for a general contingent claim on S under the SLV model. For any time $t \geq 0$ denote by $f_t = f(t, S_t, \lambda_t)$ the model price of the claim that is calibrated at time t . Since λ_t contains the future parameters of a local volatility model, the claim price must satisfy the following partial differential equation at each time $t > 0$:⁷

⁶ Jumps in the instantaneous volatility, if any, can be modelled by adding Poisson jumps to the dynamics of the parameters of the local volatility model. In this case, our theoretical findings need to be extended accordingly for jumps.

⁷ Within a local volatility model the parameters are assumed constant, hence the claim price can be expressed as a function of t and S_t only. Then (19) follows from an application of Itô's lemma and the risk-neutrality argument.

$$\frac{\partial f_t}{\partial t} + (r - q)S_t \frac{\partial f_t}{\partial S_t} + \frac{1}{2}\sigma^2 S_t^2 \frac{\partial^2 f_t}{\partial S_t^2} = rf_t \quad (19)$$

where $\sigma = \sigma(t, S_t, \lambda_t)$. Note that (19) only holds locally, assuming the model is re-calibrated at each time t .

Theorem 1

In the model (13) – (17), the risk-neutral dynamics of the contingent claim price are given by:

$$df_t = rf_t dt + \left(\sigma S_t \frac{\partial f_t}{\partial S_t} + \sum_i \beta_i \varrho_{i,S} \frac{\partial f_t}{\partial \lambda_t^i} \right) dB_t + \sum_i \beta_i \sqrt{1 - \varrho_{i,S}^2} \frac{\partial f_t}{\partial \lambda_t^i} dZ_t^i \quad (20)$$

and for the absence of arbitrage the following drift condition must be satisfied for every $t > 0$:

$$\sum_i \left(\alpha_i \frac{\partial f_t}{\partial \lambda_t^i} + \sigma S_t \beta_i \varrho_{i,S} \frac{\partial^2 f_t}{\partial \lambda_t^i \partial S_t} + \frac{1}{2} \sum_j \varrho_{i,j} \beta_i \beta_j \frac{\partial^2 f_t}{\partial \lambda_t^i \partial \lambda_t^j} \right) = 0 \quad (21)$$

We now consider the claim price dynamics for some special cases of the model.

Case 1: If $\beta_i(t, \lambda_t^i) = 0 \quad \forall i$ we have the standard risk-neutral dynamics of a deterministic volatility model. In other words, when the local volatility parameters are non-stochastic:

$$df_t = rf_t dt + \sigma(t, S_t, \lambda_t) S_t \frac{\partial f_t}{\partial S_t} dB_t$$

Case 2: If $\beta_i(t, \lambda_t^i) \neq 0$ and $\varrho_{i,S}(t, S_t, \lambda_t^i) = 1 \quad \forall i$ then no new source of uncertainty is introduced, so (20) becomes:

$$df_t = rf_t dt + \left(\sigma S_t \frac{\partial f_t}{\partial S_t} + \sum_i \beta_i \frac{\partial f_t}{\partial \lambda_t^i} \right) dB_t$$

where the diffusion coefficient is modified but the claim is redundant (i.e. it can be replicated).

Case 3: If $\beta_i(t, \lambda_t^i) \neq 0$ and $\varrho_{i,S}(t, S_t, \lambda_t^i) = 0 \quad \forall i$ then the claim price dynamics will be driven by a multi-factor model:

$$df_t = rf_t dt + \sigma S_t \frac{\partial f_t}{\partial S_t} dB_t + \sum_i \beta_i \frac{\partial f_t}{\partial \lambda_t^i} dZ_t^i$$

Theorem 2

The first and second order sensitivities of a contingent claim price $f(t, S, \lambda)$ with respect to S , and the first order sensitivity to time t , are given by:

$$\begin{aligned}
 \delta &= \frac{\partial f}{\partial S} + \sum_{i=1}^n \frac{\beta_i Q_{i,S}}{\sigma S} \frac{\partial f}{\partial \lambda^i} \\
 \gamma &= \frac{\partial f}{\partial S^2} + \sum_{i=1}^n \frac{\beta_i Q_{i,S}}{\sigma S} \left(2 \frac{\partial^2 f}{\partial S \partial \lambda^i} - \frac{1}{S} \frac{\partial f}{\partial \lambda^i} + \sum_{j=1}^n \frac{\beta_j Q_{j,S}}{\sigma S} \frac{\partial^2 f}{\partial \lambda^i \partial \lambda^j} \right) \\
 \tau &= \frac{\partial f}{\partial t} + \sum_{i=1}^n \frac{\partial f}{\partial \lambda^i} \left(\alpha_i - (r - q) \frac{\beta_i Q_{i,S}}{\sigma} + \frac{1}{2} \sigma \beta_i Q_{i,S} \right)
 \end{aligned} \tag{22}$$

The second term on the right hand side of (22) is an adjustment factor that depends on the correlation between movements of each parameter and movements of the asset price S . In effect, we split each hedge ratio into two parts: a sensitivity derived from the standard view (i.e. calibrated to the smile at a fixed point in time) and an adjustment factor that depends on the dynamics of the stochastic parameters.

III.3 Absence of Arbitrage

The absence of static arbitrage is guaranteed provided the local volatilities are never negative, as in Schweizer and Wissel (2008). So we now consider the conditions for absence of dynamic arbitrage in the SLV model. A local volatility model is arbitrage-free under the assumption of a static local volatility surface, i.e. when the volatility at each node of a tree for the asset price is known with certainty. But since this surface is not found to be static in practice, the absence of dynamic arbitrage requirement imposes restrictions on local volatility dynamics. But we should view this in the context of constant parameter stochastic volatility models. It is usually that case that such a model is ‘arbitrage-free’ in theory. However if the calibrated parameters change over time, which is usually the case, there is *nothing* to ensure that the model is arbitrage-free in practice.

[Insert section on existence of arbitrage-free prices].

In particular, the drift condition (21) must hold for *any* claim, hence it places a strong constraint on the dynamics of the model parameters. For instance, if $\beta_i = 0 \forall i$ then also $\alpha_i = 0 \forall i$. Thus if the volatility surface moves at all it does so stochastically. And if there is only one stochastic parameter in the model, for instance as is there is in the ‘SABR’ model introduced by Hagan et al (2002), then the condition (21) implies

$$\frac{\partial^2 f_t}{\partial \alpha_i \partial S_t} = 0$$

where α_t here denotes the stochastic alpha parameter of this well-known model. In others words, in SLV models with only one stochastic parameter, the price of the claim is arbitrage-free if its delta does not depend on this parameter.

When two parameters are stochastic there is much more flexibility
 [Discuss other conditions for specific SLV models]

III.3 Similar Approaches

Dupire (1996) and Derman and Kani (1998) propose modeling the dynamics of local variance for each future price K and time T directly. For instance:

$$\frac{dS}{S} = (r - q)dt + \sigma(t)dB \quad (23)$$

$$d\sigma_{K,T}^2(t, S) = a_{K,T}(t, S)dt + \sum b_{K,T}^i(t, S)dW_i \quad (24)$$

In (23) $\sigma(t)$ represents the stochastic instantaneous volatility consistent with the local variance dynamics (24) and it can be derived from the integral form of (24):

$$\sigma^2(t) = \sigma_{S,t}^2(t_0, S_0) + \int_{t_0}^t a_{S,t}(u, S(u))du + \sum \int_{t_0}^t b_{S,t}^i(u, S(u))dW_i(u)$$

where $\sigma(t) = \sigma_{S,t}(t, S)$. Therefore, if $a_{K,T}$ and $b_{K,T}^i$ are known for every K and T , then the stochastic volatility $\sigma(t)$ is fully specified in terms of the local volatility surface today $\sigma_{S,t}(t_0, S_0)$ and its dynamics over time. Note that (24) holds for each pair (K, T) in the local variance surface and that the same vector of Wiener processes \mathbf{W} drives the dynamics of the whole surface, according to the coefficients $a_{K,T}$ and $b_{K,T}^i$. For instance, if $b_{K,T}^i = b^i$ for every (K, T) then any shock causes a parallel shift in the local variance surface. Since there are an infinite number of equations a functional form for $a_{K,T}$ and each $b_{K,T}^i$ may be required in practice.

This approach is attractive because it models local volatility directly and it is general enough to capture any continuous dynamics for volatility. However, we still need to impose a restriction on the drift terms $a_{K,T}$ to avoid arbitrage opportunities and, although the dependence of the drifts on K and T allows enough freedom to define a risk-neutral measure, the resulting expression is rather complex. Solving the no-arbitrage condition is not trivial in continuous time, so Derman and Kani (1998) employ stochastic implied trees, an extension to their original implied trinomial trees. By contrast, modeling the dynamics of the parameters as in the SLV model is tractable and can be easily put in practice because the vector λ_t of parameters is observable.

[Insert review of new paper by Wissel, 2007]

IV. EQUIVALENT APPROACHES

In this section we first derive the implied probability and local variance SDEs that are equivalent to the stochastic local volatility model (13) – (17). Then we prove that our model is equivalent to the market model of implied volatilities.

IV.1 Implied probability and local variance dynamics

Now consider the model's implied distribution at time T . This is defined as the risk-neutral marginal probability distribution of the asset price $S(T)$ that is consistent with the market prices of liquid options expiring at time T . Breeden and Litzenberger (1978) show that this can be obtained from a simple differentiation of the vanilla option price with respect to K :

$$\pi_t^T(S)\Big|_{S=K} = e^{r(T-t)} \frac{\partial^2 f(K, T)}{\partial K^2} \quad (25)$$

where $\pi_t^T(S)$ is the implied risk-neutral density of the asset price S at time $T > t$. Dupire (1996) shows that (25) is the undiscounted price of an infinitesimal butterfly spread centered at K , and hence $\pi_t^T(K)$ is a martingale whose risk-neutral dynamics under the SLV model are given by:

$$d\pi = \left(\sigma S \frac{\partial \pi}{\partial S} + \sum_i \beta_i Q_{i,S} \frac{\partial \pi}{\partial \lambda^i} \right) dB + \sum_i \beta_i \sqrt{1 - Q_{i,S}^2} \frac{\partial \pi}{\partial \lambda^i} dZ_i \quad (26)$$

where we have assumed an arbitrage-free market and used the result (20).⁸

Likewise, applying Itô's lemma to the local variance (6) under the assumption (21) we obtain:

$$dV = \sigma S \frac{\partial V}{\partial S} \left(dB - \left(\sigma S \frac{\partial G}{\partial S} + \sum_i \beta_i Q_{i,S} \frac{\partial G}{\partial \lambda^i} \right) dt \right) + \sum_i \beta_i \frac{\partial V}{\partial \lambda^i} \left(dW_i - \left(\sigma S Q_{i,S} \frac{\partial G}{\partial S} + \sum_j \beta_j Q_{i,j} \frac{\partial G}{\partial \lambda^j} \right) dt \right) \quad (27)$$

where $G(K, T) = \ln \frac{\partial^2 f(K, T)}{\partial K^2}$ and $V(K, T) = \sigma_L^2(t, S)\Big|_{t=T, S=K}$ is a martingale under the K -strike

T -maturity 'forward risk-adjusted' measure, i.e. the measure defined by the price of an infinitesimal butterfly spread centred at K as numeraire. See Derman and Kani (1998) or Fengler (2005, section 3.8).

⁸ It is interesting to compare this with the dynamics of Panigirtzoglou and Skiadopoulos (2004), who propose a simple mean-reverting model for the entire implied risk-neutral distribution based principal component analysis. But the correspondence is not one-to-one since the SLV model has more parameters.

IV.2 Stochastic implied volatility dynamics

Denote the Black-Scholes (B-S) price at time t of a standard European option with strike K and maturity T when the asset price is S and the implied volatility is $\theta_{K,T}$ by:

$$f_{K,T}^{BS} = f_{K,T}^{BS}(t, S, \theta_{K,T})$$

The market implied volatility $\theta_{K,T}^M(t, S)$ is that θ such that the B-S model price equals the observed market price of the option. Likewise, when the local volatility model is calibrated to a market implied volatility surface at each time t the model implied volatility $\theta_{K,T}^L(t, S, \lambda)$ is defined by equating the local volatility price to the B-S price:

$$f_{K,T}^L(t, S, \lambda) = f_{K,T}^{BS}(t, S, \theta_{K,T}^L(t, S, \lambda)) \quad (28)$$

Note that $\theta_{K,T}^M(t, S)$ and $\theta_{K,T}^L(t, S, \lambda)$ will be the same for all K and T if the local volatility model is able to fit market prices exactly. Yet, unfortunately this is hardly the case for a parametric model.

We now derive an explicit relationship between the stochastic local volatility price dynamics and the evolution of the model implied volatility. This proves that the SLV model has implied volatility dynamics that are identical to those specified by Schönbucher (1999) for a market model of implied volatilities. For this we shall need the following notation for the B-S price and volatility sensitivities:

$$\begin{aligned} \tau_{K,T}^{BS}(t, S, \theta_{K,T}) &= \frac{\partial f_{K,T}^{BS}}{\partial t}; & \delta_{K,T}^{BS}(t, S, \theta_{K,T}) &= \frac{\partial f_{K,T}^{BS}}{\partial S}; & \gamma_{K,T}^{BS}(t, S, \theta_{K,T}) &= \frac{\partial^2 f_{K,T}^{BS}}{\partial S^2} \\ \nu_{K,T}^{BS}(t, S, \theta_{K,T}) &= \frac{\partial f_{K,T}^{BS}}{\partial \theta}; & \varkappa_{K,T}^{BS}(t, S, \theta_{K,T}) &= \frac{\partial^2 f_{K,T}^{BS}}{\partial \theta^2}; & \omega_{K,T}^{BS}(t, S, \theta_{K,T}) &= \frac{\partial^2 f_{K,T}^{BS}}{\partial S \partial \theta} \end{aligned} \quad (29)$$

and the following notation for the local volatility price sensitivities:

$$\tau_{K,T}^L(t, S, \lambda) = \frac{\partial f_{K,T}^L}{\partial t}; \quad \delta_{K,T}^L(t, S, \lambda) = \frac{\partial f_{K,T}^L}{\partial S}; \quad \gamma_{K,T}^L(t, S, \lambda) = \frac{\partial^2 f_{K,T}^L}{\partial S^2} \quad (30)$$

Lemma 1

The model implied volatility $\theta_{K,T}^L(t, S, \lambda)$ has the following sensitivities to t , S and λ :

$$\begin{aligned} \frac{\partial \theta_{K,T}^L}{\partial t} &= \frac{\tau_{K,T}^L - \tau_{K,T}^{BS}}{\nu_{K,T}^{BS}}; & \frac{\partial \theta_{K,T}^L}{\partial S} &= \frac{\delta_{K,T}^L - \delta_{K,T}^{BS}}{\nu_{K,T}^{BS}}; & \frac{\partial \theta_{K,T}^L}{\partial \lambda^i} &= \frac{1}{\nu_{K,T}^{BS}} \frac{\partial f_{K,T}^L}{\partial \lambda^i} \\ \frac{\partial^2 \theta_{K,T}^L}{\partial S^2} &= \frac{1}{\nu_{K,T}^{BS}} \left(\gamma_{K,T}^L - \gamma_{K,T}^{BS} - \left(\frac{\delta_{K,T}^L - \delta_{K,T}^{BS}}{\nu_{K,T}^{BS}} \right) \left(2\omega_{K,T}^{BS} + \chi_{K,T}^{BS} \left(\frac{\delta_{K,T}^L - \delta_{K,T}^{BS}}{\nu_{K,T}^{BS}} \right) \right) \right) \\ \frac{\partial^2 \theta_{K,T}^L}{\partial S \partial \lambda^i} &= \frac{1}{\nu_{K,T}^{BS}} \left(\frac{\partial^2 f_{K,T}^L}{\partial S \partial \lambda^i} - \frac{1}{\nu_{K,T}^{BS}} \frac{\partial f_{K,T}^L}{\partial \lambda^i} \left(\omega_{K,T}^{BS} + \chi_{K,T}^{BS} \left(\frac{\delta_{K,T}^L - \delta_{K,T}^{BS}}{\nu_{K,T}^{BS}} \right) \right) \right) \\ \frac{\partial^2 \theta_{K,T}^L}{\partial \lambda^i \partial \lambda^j} &= \frac{1}{\nu_{K,T}^{BS}} \left(\frac{\partial^2 f_{K,T}^L}{\partial \lambda^i \partial \lambda^j} - \chi_{K,T}^{BS} \left(\frac{1}{\nu_{K,T}^{BS}} \right)^2 \frac{\partial f_{K,T}^L}{\partial \lambda^i} \frac{\partial f_{K,T}^L}{\partial \lambda^j} \right) \end{aligned}$$

The next lemma describes the partial differential equation that implied volatility must satisfy to be consistent with any local volatility model. In the following we use the shorthand notation $\theta = \theta_{K,T}^L(t, S, \boldsymbol{\lambda})$ and $\sigma = \sigma(t, S, \boldsymbol{\lambda})$, and define d_1 and d_2 as in the B-S formula:

Lemma 2

The model implied volatility must satisfy the following partial differential equation:

$$\frac{\partial \theta}{\partial t} + \left(r - q - \sigma^2 \frac{d_2}{\theta \sqrt{T-t}} \right) S \frac{\partial \theta}{\partial S} + \frac{1}{2} \sigma^2 S^2 \left(\frac{\partial^2 \theta}{\partial S^2} + \frac{d_1 d_2}{\theta} \left(\frac{\partial \theta}{\partial S} \right)^2 \right) + \frac{1}{2} \frac{\sigma^2 - \theta^2}{\theta(T-t)} = 0$$

Note that the differential equation in Lemma 2 has no partial derivative on the elements of $\boldsymbol{\lambda}$, even though the implied volatility is not independent of $\boldsymbol{\lambda}$. Hence even when the local volatility surface is static the implied volatility surface will move over time. However, problems may arise because the permissible movements are too restricted.

Now we show how the dynamics of the entire implied volatility surface will be governed by the same stochastic factors as those driving the local volatility and the option price. The main result of this section, a Theorem, proves the equivalence of the general SLV model with the ‘market model’ of stochastic implied volatilities specified by Schönbucher (1999).

Lemma 3

The dynamics of the model implied volatility in the SLV model are given by

$$d\theta = \xi dt + \sigma S \frac{\partial \theta}{\partial S} dB + \sum_{i=1}^n \frac{\partial \theta}{\partial \lambda^i} \beta_i dW_i \quad (31)$$

where the drift term must satisfy $\int_t^T |\xi(u)| du < \infty$ with

$$\xi(t) = \frac{1}{2} \frac{\theta^2 - \sigma^2}{\theta(T-t)} + \frac{d_2 \sigma \psi}{\theta \sqrt{T-t}} - \frac{1}{2} \frac{d_1 d_2 \eta^2}{\theta} \quad (32)$$

Here ψ is related to the covariance between implied volatility and asset price movements:

$$\psi = \sigma S \frac{\partial \theta}{\partial S} + \sum_{i=1}^n \beta_i \varrho_{i,S} \frac{\partial \theta}{\partial \lambda^i}$$

and η^2 is the variance of the implied volatility process:

$$\eta^2 = \psi^2 + \sum_{i=1}^n \sum_{j=1}^n \beta_i \beta_j (\varrho_{i,j} - \varrho_{i,S} \varrho_{j,S}) \frac{\partial \theta}{\partial \lambda^i} \frac{\partial \theta}{\partial \lambda^j}$$

and all partial derivatives of θ are as in Lemma 1.

Schönbucher (1999) expressed his version of the market model of implied volatilities in terms of uncorrelated Brownian motions. Thus, to prove equivalence, we re-write (31) using only uncorrelated Brownian motions as follows:

Theorem 3

The dynamics of the SLV model implied volatility may be written:

$$d\theta = \xi dt + \psi dB + \sum_{j=1}^n \zeta_j dZ_j^* \quad (33)$$

with ξ and ψ defined as in Lemma 3, and where $dBdZ_j^* = dZ_i^* dZ_j^* = 0$ for $i \neq j$ almost surely, and:

$$\zeta_j = \sum_{i=j}^n \beta_i \sqrt{1 - \varrho_{i,S}^2} C_{i,j} \frac{\partial \theta}{\partial \lambda^i}$$

where $C_{i,j}$ are the elements of the Cholesky decomposition \mathbf{C} of the correlation matrix $\mathbf{\Lambda}$ with:

$$\mathbf{C}^T \mathbf{C} = \mathbf{\Lambda}, \quad \Lambda_{i,j} = \frac{\varrho_{i,j} - \varrho_{i,S} \varrho_{j,S}}{\sqrt{(1 - \varrho_{i,S}^2)(1 - \varrho_{j,S}^2)}} \quad \text{and} \quad \eta^2 = \psi^2 + \sum_{j=1}^n \zeta_j^2.$$

Apart from minor differences in notation, equation (33) is precisely the same as equation (2.7) of Schönbucher (1999) for the dynamics of a stochastic implied volatility with the drift term given by equation (3.7) of that paper. Since Schönbucher's is a stochastic volatility model, this corollary proves that local and stochastic volatility models can indeed be unified under a single framework.

This result also enables one to specify the instantaneous correlation between the implied volatility and the asset price changes as:

$$\rho_{\theta,S} = \frac{\text{Cov}(d\theta, dS)}{\sqrt{\text{Var}(d\theta)\text{Var}(dS)}} = \frac{\psi\sigma S dt}{\sqrt{\eta^2 dt \sigma^2 S^2 dt}} = \frac{\psi}{\eta} = \frac{\psi}{\sqrt{\psi^2 + \sum_{j=1}^n \zeta_j^2}} \quad (34)$$

Hence the implied volatility and asset price movements will have perfect correlation, of ± 1 depending on the sign of the covariance ψ , if and only if $\zeta_j = 0$ for all j , i.e. when the local volatility surface is fixed. In other words, the instantaneous volatility is deterministic if and only if variations in implied volatility and the asset price are perfectly correlated.

Note that Schönbucher models the implied volatility dynamics for each strike K and maturity T separately, whilst we provide the dynamics for all strikes and maturities simultaneously. If there are options for k strikes and m maturities in the market the ‘market model’ specifies at least mk diffusions, one for each traded option because the drift term is option-dependent, and assuring that these diffusions are consistent and arbitrage-free among themselves is an issue still under research. On the other hand, the SLV approach parameterizes the smile surface with $n \ll mk$ parameters. This reduces the complexity compared to the market model from mk time series of implied volatilities to only n time series of parameters plus the asset price S .

V. CONCLUSIONS

Potential links between stochastic volatility and local volatility models were identified many years ago, yet these models have been developed in two separate strands of literature. Most research on stochastic volatility has specified a single factor diffusion for the instantaneous variance or volatility of the underlying asset; but research on local volatility models has assumed a deterministic instantaneous volatility function for the underlying asset price diffusion, with no reference to the dynamic evolution of volatility. Both approaches were incomplete, the former capturing the dynamic properties of volatility but only in a one-dimensional space, the latter focusing on the multi-dimensional aspects of volatility but ignoring its time-evolution. However, recent developments on diffusions for implied volatility have extended the stochastic volatility approach to be consistent with the cross-section of implied volatilities as well as their dynamics. To concord with this view, the deterministic local volatility model, which implies only a deterministic evolution for implied volatility, requires generalization.

We have shown that the stochastic volatility and local volatility approaches can be unified within a general framework and it is only when these approaches take a restricted view of volatility dynamics that they appear to be different. Following Dupire (1996) and Derman and Kani (1998) we regard

the deterministic local volatility model as merely a special case of a more general stochastic local volatility model. That is, we define local volatility as the square root of the conditional expectation of a future instantaneous variance that depends on stochastic parameters of the local volatility function as well as the underlying price. Thus we have modeled the stochastic evolution of a locally deterministic volatility surface over time and we have proved an important general result: that a stochastic parametric local volatility model induces implied volatility dynamics that are equivalent to those of a market model for stochastic implied volatilities. Hence the two models have identical claim prices. Explicit expressions for the delta, gamma and theta hedge ratios are easy to derive in the SLV framework, and these are identical to those in the market model for implied volatilities.

REFERENCES

- Alexander, C. (2001) 'Principles of the skew' *Risk* 14: 1 S29- S32. Reprinted in Alexander Lipton (Ed.) *Exotic Options*. Risk Publications, 2004, 57-64.
- Avellaneda, M., C. Friedman, R. Holmes and D. Samperi (1997) Calibrating Volatility Surfaces via Relative Entropy Minimization. *Applied Mathematical Finance*, 4, 37–64
- Bates, D.S. (2003) Empirical Option Pricing: a Retrospection. *Journal of Econometrics*, 116, 387-404.
- Black, F. and M. Scholes (1973) The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81, 637-659.
- Bouchouev, I. and V. Isakov (1997) The Inverse Problem of Option Pricing. *Inverse Problems*, 13, 11–7.
- Bouchouev, I. and V. Isakov (1999) Uniqueness, Stability and Numerical Methods for the Inverse Problem that Arises in Financial Markets. *Inverse Problems*, 15, 95–116.
- Brace, A., Goldys B., Klebaner F., and R. Womersley (2001) Market model of stochastic volatility with applications to the BGM model. *Working paper, University of New South Wales*.
- Brace, A., Goldys B., van de Hoek, J. and R. Womersley (2002) Markovian models in the stochastic implied volatility framework. *Working paper, University of New South Wales*.
- Breeden, D. and R. Litzenberger (1978) Prices of State-Contingent Claims Implicit in Option prices. *Journal of Business*, 51, 621-651.
- Brigo, D. and F. Mercurio (2002) Lognormal-Mixture Dynamics and Calibration to Market Volatility Smiles. *International Journal of Theoretical & Applied Finance* 5:4, 427-446.
- Brunner, B. and R. Hafner (2003) Arbitrage-Free Estimation of the Risk-Neutral Density from the Implied Volatility Smile. *Journal of Computational Finance*, 7:1, 75-106.
- Carr, P. and D.B. Madan (2005) A Note on Sufficient Conditions for No Arbitrage. *Finance Research Letters*, 2:3, 125-130.
- Davis, M., and D. Hobson (2007) The Range of Traded Option Prices. *Mathematical Finance* 17, 1–14.
- Derman, E. and I. Kani (1994) Riding on a Smile. *Risk*, 7:2, 32-39.
- Derman, E. and I. Kani (1998) Stochastic Implied Trees: Arbitrage Pricing with Stochastic Term and Strike Structure of Volatility. *International Journal of Theoretical and Applied Finance*, 1:1, 61-110.
- Dumas, B., Fleming, F. and R. Whaley (1998) Implied Volatility Functions: Empirical Tests. *Journal of Finance*, 53:6, 2059-2106.
- Dupire, B. (1994) Pricing with a Smile, *Risk*, 7:1, 18-20
- Dupire, B. (1996) A Unified Theory of Volatility. Working paper, now published in *Derivatives Pricing: The Classic Collection*, (Ed. P. Carr) Risk Books, 2004.
- Fengler, M.R. (2005) *Semiparametric Modelling of Implied Volatility*. 1st edition. Springer Finance Lecture Notes.
- Gyöngy, I. (1986) Mimicking the One-Dimensional Marginal Distributions of Processes Having an Ito Differential. *Probability Theory and Related Fields*, 71, 501-516.

- Hafner, R. (2004) *Stochastic implied volatility*. 1st edition. Lecture Notes in Economics and Mathematical Systems. Springer.
- Hagan, P.S, D. Kumar, A. S. Lesniewski, and D. E. Woodward (2002) Managing smile risk. *Wilmott Magazine* September, 84-108.
- Heath, D., Jarrow, R. and A. Morton (1992). Bond Pricing and the Term Structure of Interest Rates: a New Methodology for Contingent Claims Valuation. *Econometrica* 60, 77-105
- Heston, S (1993). A Closed Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *Review of Financial Studies*, 6:2, 327-343.
- Hull, J. and A. White (1987) The Pricing of Options on Assets with Stochastic Volatilities. *Journal of Finance*, 42:2, 281-300.
- Jackwerth, J.C. (1999) Option-Implied Risk-Neutral Distributions and Implied Binomial Trees: A Literature Review. *Journal of Derivatives* 7, 66-82.
- Ledoit, O., P. Santa-Clara and S. Yan (2002) Relative Pricing of Options with Stochastic Volatility. *Working paper*. Available from <http://www.personal.anderson.ucla.edu/pedro.santa-clara> .
- Nelson, D.B. (1990) ARCH Models as Diffusion Approximations. *Journal of Econometrics* 45, 7-38
- Nogueira, L. M. (2006) Hedging Options with Local and Stochastic Volatility Models. PhD Thesis, ICMA Centre, University of Reading.
- Panigirtzoglou, N. and G. Skiadopoulos (2004) A New Approach to Modeling the Dynamics of Implied Distributions: Theory and Evidence from the S&P 500 Options. *Journal of Banking and Finance*, 28:7, 1499-1520.
- Psychoyios, D., Skiadopoulos, G. and P. Alexakis (2003) A Review of Stochastic Volatility Processes: Properties and Implications. *Journal of Risk Finance* 4:3, 43-60.
- Rubinstein, M. (1994) Implied Binomial Trees. *Journal of Finance*, 49:3, 771-818.
- Schönbucher, P.J. (1999) A Market Model of Stochastic Implied Volatility. *Philosophical Transactions of the Royal Society*. Series A, 357, 2071-2092.
- Schweizer, M. and J. Wissel (2007) Arbitrage-free Market Models for Option Prices: The Multi-Strike case”, *Preprint, ETH Zurich*. Available from <http://www.nccr-finrisk.unizh.ch/wps>.
- Schweizer, M. and J. Wissel (2008) Term structures of implied volatilities: absence of arbitrage and existence results. *Mathematical Finance*, 18:1, 77-114.
- Skiadopoulos, G. (2001) Volatility Smile-Consistent Option Models: A Survey. *International Journal of Theoretical and Applied Finance*, 4:3, 403-438.
- Stein, E. M. and J. C. Stein (1991) Stock Price Distributions with Stochastic Volatility: an Analytic Approach. *Review of Financial Studies* 4, 727-752.
- Zilber, A. (2007) A Market Model for Stochastic Smiles. *Working paper*. Available from <http://ssrn.com/abstract=963572>.

APPENDIX

Proof of Theorem 1: Applying Itô's lemma to a claim price $f_t = f(t, S_t, \lambda_t)$ gives dynamics:

$$df_t = \frac{\partial f_t}{\partial t} dt + \frac{\partial f_t}{\partial S_t} dS_t + \frac{1}{2} \frac{\partial^2 f_t}{\partial S_t^2} d\langle S \rangle_t + \sum_i \frac{\partial f_t}{\partial \lambda_t^i} d\lambda_t^i + \sum_i \frac{\partial^2 f_t}{\partial \lambda_t^i \partial S} d\langle \lambda^i, S \rangle_t + \frac{1}{2} \sum_i \sum_j \frac{\partial^2 f_t}{\partial \lambda_t^i \partial \lambda_t^j} d\langle \lambda^i, \lambda^j \rangle_t$$

But using (13) - (17):

$$df_t = \Xi_t dt + \left(\sigma S_t \frac{\partial f_t}{\partial S_t} + \sum_i \beta_i \varrho_{i,S} \frac{\partial f_t}{\partial \lambda_t^i} \right) dB_t + \sum_i \beta_i \sqrt{1 - \varrho_{i,S}^2} \frac{\partial f_t}{\partial \lambda_t^i} dZ_t^i$$

with

$$\Xi_t = \frac{\partial f_t}{\partial t} + (r - q) S_t \frac{\partial f_t}{\partial S_t} + \frac{1}{2} \sigma^2 S_t^2 \frac{\partial^2 f_t}{\partial S_t^2} + \sum_i \left(\alpha_i \frac{\partial f_t}{\partial \lambda_t^i} + \sigma S_t \beta_i \varrho_{i,S} \frac{\partial^2 f_t}{\partial \lambda_t^i \partial S_t} + \frac{1}{2} \sum_j \varrho_{i,j} \beta_i \beta_j \frac{\partial^2 f_t}{\partial \lambda_t^i \partial \lambda_t^j} \right)$$

The drift condition (21) follows from (19) and on noting that under the risk-neutral probability the drift of the claim price must be the risk-free rate. \square

Proof of Theorem 2: When movements in S and λ are correlated, we can express each λ_i as a function of t , S and Z_i so that from Itô's formula:

$$d\lambda^i = \left(\frac{\partial \lambda^i}{\partial t} + (r - q) S \frac{\partial \lambda^i}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 \lambda^i}{\partial S^2} + \frac{1}{2} \frac{\partial^2 \lambda^i}{\partial Z_i^2} \right) dt + \sigma S \frac{\partial \lambda^i}{\partial S} dB + \frac{\partial \lambda^i}{\partial Z_i} dZ_i$$

Equating coefficients gives:

$$\frac{\partial \lambda^i}{\partial S} = \frac{\beta_i \varrho_{i,S}}{\sigma S} \Rightarrow \frac{\partial^2 \lambda^i}{\partial S^2} = -\frac{\beta_i \varrho_{i,S}}{\sigma S^2}$$

$$\frac{\partial \lambda^i}{\partial Z_i} = \beta_i \sqrt{1 - \varrho_{i,S}^2} \Rightarrow \frac{\partial^2 \lambda^i}{\partial Z_i^2} = 0$$

$$\frac{\partial \lambda^i}{\partial t} = \alpha_i - (r - q) \frac{\beta_i \varrho_{i,S}}{\sigma} + \frac{1}{2} \sigma \beta_i \varrho_{i,S}$$

Now the chain rule gives the first order price sensitivity as:

$$\delta = \frac{d}{dS} (f(t, S, \lambda)) = \frac{\partial f}{\partial S} + \sum_i \frac{\partial f}{\partial \lambda^i} \frac{\partial \lambda^i}{\partial S} = \frac{\partial f}{\partial S} + \sum_i \frac{\beta_i \varrho_{i,S}}{\sigma S} \frac{\partial f}{\partial \lambda^i}$$

Similarly the gamma and theta follow using:

$$\gamma = \frac{d}{dS} (\delta(t, S, \lambda)) \quad \tau = \frac{d}{dt} (f(t, S, \lambda)) \quad \square$$

Proof of Lemma 1. Differentiate (28) with respect to t , S and each parameter and apply the chain rule in the right-hand side whenever necessary. For instance:

$$\frac{\partial f_{K,T}^L}{\partial S} = \frac{\partial f_{K,T}^{BS}}{\partial S} + \frac{\partial f_{K,T}^{BS}}{\partial \theta} \frac{\partial \theta_{K,T}^L}{\partial S} \Rightarrow \frac{\partial \theta_{K,T}^L}{\partial S} = \frac{\delta_{K,T}^L - \delta_{K,T}^{BS}}{\nu_{K,T}^{BS}}$$

and so forth. □

Proof of Lemma 2. Subtract the Black-Scholes PDE from (19), apply (28) and Lemma 1, and use the relationship between the Black-Scholes sensitivities:

$$\nu_{K,T}^{BS} = \frac{1}{\theta(T-t)S^2} \nu_{K,T}^{BS}; \quad \kappa_{K,T}^{BS} = \frac{d_1 d_2}{\theta} \nu_{K,T}^{BS}; \quad \omega_{K,T}^{BS} = -\frac{d_2}{\theta S \sqrt{T-t}} \nu_{K,T}^{BS} \quad \square$$

Proof of Lemma 3. From Itô's lemma and using (13)-(17), the dynamics of the model implied volatility are given by:

$$\begin{aligned} d\theta &= \xi(t)dt + \sigma S \frac{\partial \theta}{\partial S} dB + \sum_i \frac{\partial \theta}{\partial \lambda^i} \beta_i dW_i \\ \xi(t) &= \frac{\partial \theta}{\partial t} + (r-q)S \frac{\partial \theta}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 \theta}{\partial S^2} + \sum_i \alpha_i \frac{\partial \theta}{\partial \lambda^i} + \sum_i \frac{\partial^2 \theta}{\partial S \partial \lambda^i} \sigma S \beta_{i,Q_{i,S}} + \frac{1}{2} \sum_i \sum_j \frac{\partial^2 \theta}{\partial \lambda^i \partial \lambda^j} \beta_i \beta_j Q_{i,j} \end{aligned}$$

Using Lemma 1, the drift expands to:

$$\begin{aligned} &\frac{\partial \theta}{\partial t} + (r-q)S \frac{\partial \theta}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 \theta}{\partial S^2} + \frac{1}{\nu_{BS}} \sum_i \left(\alpha_i \frac{\partial f^L}{\partial \lambda^i} + \sigma S \beta_{i,Q_{i,S}} \frac{\partial^2 f^L}{\partial S \partial \lambda^i} + \frac{1}{2} \sum_j \frac{\partial^2 f^L}{\partial \lambda^i \partial \lambda^j} \beta_i \beta_j Q_{i,j} \right) - \\ &\sum_i \frac{1}{\nu_{BS}} \left[\frac{\partial \theta}{\partial \lambda^i} \left(\Omega_{BS} + \kappa_{BS} \frac{\partial \theta}{\partial S} \right) \right] \sigma S \beta_{i,Q_{i,S}} - \frac{1}{2} \sum_i \sum_j \frac{1}{\nu_{BS}} \left[\kappa_{BS} \frac{\partial \theta}{\partial \lambda^i} \frac{\partial \theta}{\partial \lambda^j} \right] \beta_i \beta_j Q_{i,j} \end{aligned}$$

Next, using the drift condition (21) and lemmas 1 and 2, this re-arranges to (32), with $T > t$, $\theta > 0$, and ψ and η^2 as above. Finally, if θ is a valid Itô's process, then $\int_t^T |\xi(u)| du < \infty$, among other regularity conditions.⁹ □

Proof of Theorem 3. The result follows from Lemma 3 and the fact that:

$$dW_i dW_j = \left(Q_{i,S} Q_{j,S} + \sqrt{1-Q_{i,S}^2} \sqrt{1-Q_{j,S}^2} \Lambda_{i,j} \right) dt = Q_{i,j} dt.$$

with $dZ_i dZ_j \xrightarrow{a.s.} \Lambda_{i,j} dt$ and the Cholesky decomposition (see e.g. Hafner, 2004, section 6.1.1). □

⁹ Note that the option prices that are consistent with the implied volatility dynamics (31) must satisfy the same no-arbitrage conditions of Lemma 1. Besides this, there is an interesting singularity on the drift ξ as $t \rightarrow T$. However, this is not a problem as long as $\int_t^T |\xi(u)| du < \infty$ for all $T > t$.