The University of Reading

THE BUSINESS SCHOOL
FOR FINANCIAL MARKETS

A Taxonomy of Option Pricing Models:
Scale Invariant Volatility and Minimum Variance Hedging

ISMA Centre Discussion Papers in Finance DP2005-10

17th July 2005

Carol Alexander
ISMA Centre, University of Reading

Leonardo Nogueira
ISMA Centre, University of Reading

ISMA Centre • The University of Reading
Whiteknights • PO Box 242 • Reading RG6 6BA • UK
Tel: +44 (0)1183 788239 • Fax: +44 (0)1189 314741
Web: www.ismacentre.rdg.ac.uk
Director: Professor John Board, Chair in Finance
The ISMA Centre is supported by the International Securities Market Association

ISMA

Abstract

Most option pricing models are 'scale invariant', i.e. the equivalent local volatility and the model implied volatility are invariant under scaling in the price dimension. We derive several model free properties for these models, including some that may help curtail the proliferation of pricing models: when calibrated exactly to the smile, all scale invariant models have the same delta and gamma and even the minimum variance delta and gamma can be identical under certain parameter constraints. The dynamic delta and delta-gamma hedging performance of several well-known local and stochastic volatility models shows that scale invariant models perform worse than the Black-Scholes model for SP500 options but with minimum variance hedge ratios the performance is dramatically improved.

JEL Classification:
Keywords: Scale invariant volatility models, local volatility, stochastic volatility, jump models, dynamic hedging, minimum variance hedge ratios.

Author Details:
Professor Carol Alexander
Chair of Risk Management and Director of Research,
ISMA Centre, Business School,
The University of Reading
PO Box 242, Reading RG6 6BA
United Kingdom
Email: c.alexander@ismacentre.rdg.ac.uk
Tel: +44 (0)1183 786431 (ISMA Centre) Fax: +44 (0)1189 314741

Leonardo M. Nogueira
PhD Student,
ISMA Centre, Business School,
The University of Reading
PO Box 242, Reading RG6 6BA
United Kingdom
Email: l.nogueira@ismacentre.rdg.ac.uk
Tel: +44 (0)1183 786675 (ISMA Centre) Fax: +44 (0)1189 314741

Acknowledgements

We are most grateful for very useful comments from Dr. Bruno Dupire, Bloomberg New York; Dr. Hyungsok Ahn, Nomura, London; Dr. Jacques Pézier, Naoufel El-Bachir, Emese Lazar and Dmitri Lvov, all of the ISMA Centre, UK; Prof. Emanuel Derman, Columbia University, New York; Prof. Rama Cont, ESSEC, France; and participants at the Risk Europe 2005 conference in Monte Carlo and the Newton Institute workshop for Financial Mathematics in Cambridge, UK. Any remaining errors are our responsibility alone.

Introduction

A standard stylized fact in volatility theory is that the empirically observed 'smile' and 'skew' shapes in Black-Scholes implied volatilities contradict the assumptions of the Black-Scholes (1973) model. This has motivated an explosion of 'smile consistent' models with different asset price dynamics: Markovian/non-Markovian, diffusion/jump/Lévy processes, local volatility/stochastic volatility, complete market/incomplete market, and so forth. Each model aims to provide an intuitive interpretation of market behavior and to price more exotic or illiquid instruments consistently with the observed prices of liquid instruments.[i] The majority of smile consistent models in common use can fulfil this aim, some better than others. Given sufficient parameters, any model can fit the smile. What distinguishes a good model from a less-good model is the pricing robustness and the accuracy of its out-of-sample prices and hedge ratios. Not all prices are transacted, often they serve only to mark-to-model. The most important quality for option pricing models is the accuracy of their hedge ratios.

As the number of alternative models has grown so also has the confusion about their volatility dynamics and hedging properties. The main focus of the paper is to unify this literature by introducing a very general class of 'scale invariant volatility' (SIV) models (so-called because the equivalent local volatility and the model implied volatility are invariant under scaling in the price dimension) and to derive their properties. We show that this class of models includes some local volatility models, jump and Lévy models, most stochastic volatility models and all Black-Scholes mixture models.[ii] What is important about this classification is that many 'model free' properties of volatility and volatility dynamics can be derived for SIV models. In particular we show that the delta and gamma are model free – that is, all SIV models have the *same* delta and gamma – and the only differences between SIV model's deltas (or gammas) is the quality of the models' fit to the smile.

The delta and gamma of an option are defined as the partial derivatives of the option price with respect to the underlying asset price. The fact that these price hedge ratios are model free for many broad classes of option pricing models may curtail the proliferation of new models and improve our understanding of those already in popular use. Yet what matters for effective hedging of the underlying price is the *minimum variance* (MV) hedge ratio, i.e. the amount of the underlying that minimizes the dependence of the hedged portfolio on the underlying price – see Bakshi, Cao and Chen (1997). The MV delta and gamma may be expressed as total derivatives of the option price with respect to the underlying. They are not, in general, equal to the 'standard' delta and gamma because the MV hedge ratios also capture the expected change in volatility when the underlying price moves. This expected change may be covered by a separate vega hedge as in Bakshi and Kapardia (2003), in which case the standard price hedge ratios apply. Otherwise, when price and volatility are correlated and in the absence of a separate vega hedge, the MV delta and gamma hedge ratios should provide superior price hedges.

We shall derive explicit expressions for MV hedge ratios in some common option pricing models and show that scale invariant models only have identical MV hedge ratios under strong parameter constraints. Empirically, differences in calibration quality lead to only very small differences between SIV model's 'standard' price hedge ratios; however the freedom to calibrate parameters so that the smile is fit as closely as possible means that larger differences between their MV hedge ratios may be recorded.

The structure of the paper is as follows: Section I introduces SIV models and derives the properties they have in common; Section II classifies several well-known models; Section III discusses a paradox concerning the SABR model and motivates the distinction between standard and minimum variance price hedge ratios; Section IV investigates the properties of SIV models for pricing and hedging vanilla options. Section V provides an empirical study of hedging properties of several well-known option pricing models; and Section VI summarizes and concludes.

I. Scale invariant volatility models and their properties

A scale invariant volatility (SIV) model is defined by a Markov process for the underlying asset price:

$$\frac{dS}{S} = \alpha(t,X)\,dt + \sigma(t,X)\,dW + \Pi(t,X)\,dJ \quad \text{at time } t > 0 \tag{1}$$

where $X = S/S_0$ is the relative asset price at time t, S_0 is the value of S at time 0, and dW and dJ define respectively a Wiener process and a Poisson process with random intensity Π, independent of each other. The parameters α and σ are well-behaved (of bounded variance) functions of t, X and possibly other random factors and possibly other variables, except for S or S_0. Typically α is equal to the risk-free rate minus the annualized expected value of the jump $\Pi(t,X)dJ$ under the risk-neutral measure, and σ is interpreted as a volatility coefficient in the absence of jumps. The only requirement is that the model price dynamics are expressed in terms of Wiener and Poisson processes and α, σ and Π are not functions of S or S_0 separately.[iii]

The distinguishing feature of SIV models is that the right-hand side of (1) is at most a function of X and not of S or S_0 separately. Depending on the specification of the spot volatility σ it is called a 'local' volatility or a 'stochastic' volatility model. In local volatility models the spot volatility is typically defined as a deterministic function of t and S, but note that it must be a function of t and X only if the model is to be scale invariant.

In this section we derive several properties for SIV models. First we prove that in any SIV model neither the probability density of the relative price X nor the characteristic function of log X can be functions of the current asset price. Consequently, in such models the price f of a standard European option is a homogeneous function of degree 1 with respect to the current asset price S_0 and the strike price K.[iv] This homogeneity property is well-known for the Black-Scholes model and some stochastic volatility models; indeed it was recognised in Merton (1973). But here it is generalized to the much wider class of scale invariant volatility models. Next, applying Euler's theorem for homogeneous functions, we prove that the partial derivatives of f with respect to S_0 are defined uniquely as functions of f and its derivatives to K only. As a result, any SIV model should produce the *same* price sensitivities when calibrated to the same set of options prices, and any difference in practice can be justified only by different fits to options prices. That is, the usual delta and gamma in any SIV model are 'model free'.

Property 1: Independence

In any SIV model the probability density $\psi_t(x)$ *of X at time t is not a function of* S_0 *, i.e.*

$$\frac{\partial \psi_t(x)}{\partial S_0} = 0 \qquad \forall x \in \mathbb{R}^+, t > 0$$

Equivalently, in any SIV model the characteristic function of ln X, i.e.

$$\varphi_t(v) = E\left[\exp\left(iv \ln X\right)\right] = \int_0^\infty \exp\left(iv \ln x\right)\psi_t(x)\,dx \tag{2}$$

at time t is not a function of S_0 .

Proof: From the definition of X we have:

$$X = \frac{S}{S_0} \Rightarrow dX = \frac{dS}{S_0} = \frac{S}{S_0}\frac{dS}{S} \Rightarrow \frac{dX}{X} = \frac{dS}{S}$$

and using dynamics (1) we write:

$$\frac{dX}{X} = \alpha(t, X)\,dt + \sigma(t, X)\,dW + \Pi(t, X)\,dJ \tag{3}$$

Since S or S_0 do not appear individually in the right-hand side of (3) and since $X(0) = 1$, we conclude that $X(t)$ is independent of the current asset price S_0 for every $t > 0$, i.e. its density function is not a function of S_0. The equivalent rule for the characteristic function follows on differentiating (2) with respect to S_0. □

Despite the simplicity of Property 1, this is all we need to derive the remaining properties of SIV models. In particular, we note that the density $\psi_t(x)$ does not need to be known in closed form. As long as it is not a function of S_0 the model is scale invariant; and vice-versa.

We now turn to the option pricing and hedging properties of SIV models. Define the price of a standard European option with strike *K* and time to maturity *T* as:

$$f\left(S_0,K,T\right)=E\left[B\left(T\right)\max\left(w\left(S\left(T\right)-K\right),0\right)\right]=E\left[B\left(T\right)\max\left(w\left(S_0X\left(T\right)-K\right),0\right)\right]$$

where w is 1 for calls or -1 for puts, and $B\left(T\right)=\exp\left(-\int_0^T r\left(t\right)dt\right)$ is the (possibly stochastic) discount factor. The expectation is under the risk-neutral probability measure and the risk-free rate $r\left(t\right)$ is not a function of S or S_0 by assumption.

Property 2: Homogeneity

In any SIV model, the price of a standard European option is a homogeneous function of degree 1 with respect to S_0 *and K:*

$$f\left(uS_0,uK,T\right)=uf\left(S_0,K,T\right)\quad\forall u\in\mathbb{R}^+,T>0 \tag{4}$$

Proof: Using Property 1, we have:

$$f\left(uS_0,uK,T\right)=E\left[B\left(T\right)\max\left(w\left(uS_0X\left(T\right)-uK\right),0\right)\right]=uE\left[B\left(T\right)\max\left(w\left(S_0X\left(T\right)-K\right),0\right)\right]=uf\left(S_0,K,T\right)$$

which is possible only because $B\left(T\right)$ and $X\left(T\right)$ are not functions of S_0 or K in SIV models. □

Property 2 implies that if we scale both the current asset price S_0 and the strike K of a standard European option by the same amount u, the European option price is also scaled by u in a SIV model. Figure 1 depicts the evolution of the asset price in a SIV model. The figure shows that when the vertical axis is scaled by a positive real number u, with $0 < u < 1$ in this example, the volatility and price-volatility correlation remain unchanged. Hence the option price will be scaled according to (4).

[Figure 1 here]

The homogeneity property in SIV models is not limited to standard European options. For instance, the present value of a forward contract expiring at $T > 0$ is:

$$F\left(uS_0,uK,T\right)=E\left[B\left(T\right)\left(uS_0X\left(T\right)-uK\right)\right]=uE\left[B\left(T\right)\left(S_0X\left(T\right)-K\right)\right]=uF\left(S_0,K,T\right)$$

where again we used the fact that $B\left(T\right)$ and $X\left(T\right)$ are independent of S_0 in SIV models. Likewise in (4) we could include barrier options and indeed any option with characteristics in the price domain beyond a simple strike K. What matters is that all features of the claim that relate to the price dimension are scaled as the asset price is scaled. Yet property 2 cannot be generalized to all claims. For instance, a digital option paying 1 if $S\left(T\right)>K$ at maturity T is worth:

$$D\left(S_0,K,T\right)=E\left[B\left(T\right)1_{\left\{S_0X(T)>K\right\}}\right]=E\left[B\left(T\right)1_{\left\{uS_0X(T)>uK\right\}}\right]=D\left(uS_0,uK,T\right)$$

Hence the price D of the digital option is a homogeneous function of degree 0.

Property 3: Separability and proportionality

In any SIV model, the first and second order sensitivities of the price of a standard European option with respect to S_0 and K obey:

$$f\left(S_0,K,T\right)=S_0\frac{\partial f\left(S_0,K,T\right)}{\partial S_0}+K\frac{\partial f\left(S_0,K,T\right)}{\partial K} \tag{5}$$

$$\frac{\partial^2 f\left(S_0,K,T\right)}{\partial {S_0}^2}=\left(\frac{K}{S_0}\right)^2\frac{\partial^2 f\left(S_0,K,T\right)}{\partial K^2} \tag{6}$$

Proof: An immediate application of Euler's theorem gives (5). To prove (6) define $m=\frac{S_0}{K}$ as the 'moneyness' of the option and set $u=\frac{1}{K}$ in (4). Then:

$$f\left(S_0,K,T\right)=Kf\left(\frac{S_0}{K},1,T\right)=Kh\left(m,T\right) \tag{7}$$

where $h\left(m,T\right)$ a function of the moneyness and time to maturity only. Differentiating with respect to S_0 and K, we have:

$$\frac{\partial f\left(S_0,K,T\right)}{\partial S_0}=\frac{\partial h\left(m,T\right)}{\partial m} \qquad \Rightarrow\frac{\partial^2 f\left(S_0,K,T\right)}{\partial {S_0}^2}=\frac{1}{K}\frac{\partial^2 h\left(m,T\right)}{\partial m^2}$$

$$\frac{\partial f\left(S_0,K,T\right)}{\partial K}=h\left(m,T\right)-m\frac{\partial h\left(m,T\right)}{\partial m}\Rightarrow\frac{\partial^2 f\left(S_0,K,T\right)}{\partial K^2}=\frac{m^2}{K}\frac{\partial^2 h\left(m,T\right)}{\partial m^2}$$

and (6) follows. □

Property 3, which is consistent with Euler's theorem for homogeneous functions of degree 1,[v] states that the option price is an additively separable function of S_0 and K and their first order derivatives. Moreover, the second derivatives w.r.t. S_0 and K are proportional. It implies that the first and second order sensitivities to S_0 – i.e. the usual delta and gamma – are functions of the option price f and its sensitivities to the strike K. This finding has important implications for hedging, as we shall see later.

The *equivalent local volatility* is the volatility function that is consistent with the forward equation as defined by Dupire (1996) and Derman and Kani (1998):

$$\sigma_L^2\left(t,s;S_0\right)\Big|_{t=T,s=K} = 2\left(\frac{\partial f\left(S_0,K,T\right)}{\partial T} + \left(r-q\right)K\frac{\partial f\left(S_0,K,T\right)}{\partial K} + qf\right)\Bigg/ K^2\frac{\partial^2 f\left(S_0,K,T\right)}{\partial K^2} \tag{8}$$

where q denotes the dividend yield of a stock or the foreign risk-free rate in FX markets. Note that the equivalent local volatility is not the same as the instantaneous volatility $\sigma\left(t,X\right)$ defined in (1), although there is a direct relationship between them given by:

$$\sigma_L^2\left(t,s;S_0\right) = \mathrm{E}\left[\sigma^2\left(t,X\right)\middle|S\left(t\right)=s\right] \tag{9}$$

That is, the square of the local volatility for the forward time t and asset price s is the conditional expectation of the square of the stochastic volatility given the asset price at time t is equal to s.

Property 4: Invariance of equivalent local volatility

In any SIV model the equivalent local volatility $\sigma_L\left(t,s;S_0\right)$ *derived from standard European options prices is a homogeneous function of degree zero in both s and* S_0*, i.e.*

$$\sigma_L\left(t,us;uS_0\right) = \sigma_L\left(t,s;S_0\right)\forall u \in \mathbb{R}^+$$

Proof: Write $f\left(S_0,K,T\right) = Kh\left(m,T\right)$ in (8) to obtain:

$$\sigma_L^2\left(t,s;S_0\right)\Big|_{t=T,s=K} = 2\left(\frac{\partial h\left(m,T\right)}{\partial T} - rm\frac{\partial h\left(m,T\right)}{\partial m} + rh\left(m,T\right)\right)\Bigg/ m^2\frac{\partial^2 h\left(m,T\right)}{\partial m^2}$$

The right-hand side is a function of $m = \frac{S_0}{K}$ and not of S_0 and K separately. Therefore:

$$\sigma_L^2\left(t,us;uS_0\right)\Big|_{t=T,us=uK} = \sigma_L^2\left(t,s;S_0\right)\Big|_{t=T,s=K} \qquad \square$$

Figure 2 shows how the homogeneity property of the model carries over to the scale invariance of equivalent local volatility. The spot volatility σ_0 is the square root of the ATM local variance at time 0

$$\sigma_0^2 = \lim_{t\downarrow 0} E\left[\sigma_L{}^2\left(t,s;S_0\right)\right] \tag{10}$$

Defining the spot volatility as a limit prevents one ignoring the possible dependence between price and volatility at time 0. We shall discuss this crucial issue in more depth in example 5 of Section II.

[Figure 2 here]

Define the *model implied volatility* $\theta\left(S_0,K,T\right)$ as the Black-Scholes volatility that is implicit in the SIV model price. The next property shows that model implied volatility also inherits the scale invariant property:

Property 5: Invariance of model implied volatility

In any SIV model, the implied volatility surface $\theta(S_0, K, T)$ *that is obtained from the model prices for standard European options is a homogeneous function of degree zero in both K and* S_0*, i.e.*

$$\theta(uS_0, uK, T) = \theta(S_0, K, T)\ \forall u \in \mathbb{R}^+ \tag{11}$$

Proof: Set the option price f from the SIV model equal to the Black-Scholes price f_{BS} with the model implied volatility, and use the fact that the Black-Scholes model is also scale invariant to write:

$$f(S_0, K, T) = f_{BS}(S_0, K, T, \theta(S_0, K, T)) \Leftrightarrow h(m, T) = h_{BS}(m, T; \theta(S_0, K, T))$$

whence $\theta(S_0, K, T)$ is implicitly defined in terms of m and T only. □

Properties 4 and 5 can be seen as the motivation for the name 'scale invariant volatility' models, in the sense that both the equivalent local volatility and the implied volatility surface are invariant when we scale in the price dimension.

Our last property is the most important property of this section. It shows that the sensitivity of the implied volatility to S_0 has the opposite sign of the slope of implied volatility surface in the K axis. This result holds for all SIV models, and in the next section we show that this class includes.the majority of option pricing models in current use.

Property 6: **Model free implied volatility**

In any SIV model, the partial sensitivities of the implied volatility surface $\theta(S_0, K, T)$ *w.r.t* S_0 *and K are related by:*

$$\frac{\partial\theta(S_0, K, T)}{\partial S_0} = -\frac{K}{S_0}\frac{\partial\theta(S_0, K, T)}{\partial K} \tag{12}$$

Proof: Apply Euler's theorem to $\theta(S_0, K, T)$ since it is a homogeneous function of degree 0. □

Property 6 has surprising implications. Since implied volatilities are observable in liquid markets, it implies that the sensitivities to S_0 given by different SIV models are the same if these models fit the slope of the implied volatility surface equally well. In other words these sensitivities are 'model-free'.

Each one of the above properties is a necessary and sufficient condition for a model to be scale invariant. That is, if any of these properties is verified for a given model then this model is scale invariant and all other properties are automatically satisfied. In fact, properties 1 to 6 are alternative definitions of SIV models. But note that properties 2 to 6 were originally defined in the context of standard European options and, although they can be extended to other claims, they may not hold for all claims. Finally, note that not all SIV models will have dynamics such as in (1). For instance, using Property 1 we have that any Lévy process for the relative price X is a scale invariant model as long as the characteristic function of ln X is not a function of S_0, yet not all Lévy processes can be expressed in terms of a single Wiener and a single Poisson process as denoted in (1).

II. Classification of option pricing models

In this section we show that many popular option-pricing models are scale invariant and discuss those that are not. Beginning with the Black-Scholes model we consider increasingly complex models, up to models where prices are driven by Lévy processes under stochastic interest rates. We shall show that some local volatility models are scale invariant, when the volatility is a non-stochastic process that depends only on the relative price X. Also any stochastic volatility model of the form

$$\begin{aligned} \frac{dS}{S} &= rdt + \sigma(Y)dW \\ dY &= a(t,Y)dt + b(t,Y)dZ \quad \langle dW, dZ \rangle = \varrho dt \end{aligned} \tag{13}$$

is scale invariant regardless of the actual specification of the Y process, given only that Y is an Itô process and a and b are not functions of S and S_0. In particular, when $\sigma(Y) = \sqrt{Y}$, this includes both the Hull and White (1987) and the Heston (1993) models. When Π is non-zero, the dynamics (1) can include some Lévy processes and, in particular, Merton's (1976) jump diffusion model. Thus the class of SIV models as defined above is broad enough to cover most pricing models in the financial literature, and even models not explicitly examined in this paper including uncertain volatility models (such as in Avellaneda et al (1995)) and volatility jump models (such as Naik (1993)).

Example 1: The Black-Scholes (BS) model

The first model of interest is the Black and Scholes (1973) model. It defines the asset price risk-neutral dynamics in the absence of dividends as:

$$\frac{dS}{S} = rdt + \sigma dW$$

where r and σ are constant. Applying Ito's lemma and integrating with respect to t, we have:

$$X(T) = \frac{S(T)}{S_0} = \exp\left[\left(r - \tfrac{1}{2}\sigma^2\right)T + \sigma W(T)\right]$$

which is independent of S_0 in the right-hand side, so that all properties derived in Section I must hold.

In fact, the price of a standard European option given by the BS model is:

$$f_{BS}\left(S_0,K,T\right)=w\left[S_0\Phi\left(wd_+\right)-Ke^{-rT}\Phi\left(wd_-\right)\right] \quad \text{with} \quad d_\pm=\frac{1}{\sigma\sqrt{T}}\left(\ln\frac{S_0}{K}+\left(r\pm\tfrac{1}{2}\sigma^2\right)T\right)$$

where Φ is the cumulative standard normal distribution, so that:

$$f_{BS}\left(uS_0,uK,T\right)=w\left[uS_0\Phi\left(wd_+\right)-uKe^{-rT}\Phi\left(wd_-\right)\right]=uf_{BS}\left(S_0,K,T\right) \tag{14}$$

since $d_\pm$ is a function of $\dfrac{S_0}{K}$ and not of S_0 and K separately, and $\dfrac{uS_0}{uK}=\dfrac{S_0}{K}$.

The homogeneity property has been verified and therefore the BS model is scale invariant.

Example 2: BS mixture models

In BS mixture models the price of a standard European option is defined as a linear combination of Black-Scholes prices, i.e.

$$f_{mix}\left(S_0,K,T\right)=\sum_i\lambda^i f^i_{BS}\left(S_0,K,T\right) \quad \text{with} \quad \sum_i\lambda^i=1$$

The lognormal mixture diffusion of Brigo and Mercurio (2002) is a good example of mixture model. Assuming that the weights λ are not functions of S_0 or K, we may write:

$$f_{mix}\left(uS_0,uK,T\right)=\sum_i\lambda^i f^i_{BS}\left(uS_0,uK,T\right)=u\sum_i\lambda^i f^i_{BS}\left(S_0,K,T\right)=uf_{mix}\left(S_0,K,T\right)$$

where we have used (14) to prove that mixture models are also scale invariant according to Property 2.

Example 3: The CEV model

The risk-neutral dynamics of the underlying price in the CEV model, introduced by Cox (1975), are written as:

$$\frac{dS}{S}=rdt+\alpha S^\beta dW \tag{15}$$

where $\alpha>0$ and $\beta<0$. As S appears in the right-hand side of (15), the CEV model is clearly not scale invariant. To verify this, we follow Schroder (1989) and write the price of a vanilla call option in the CEV model (15) as:

$$\begin{gathered} f_{CEV}\left(S_0,K,T\right)=S_0e^{-qT}Q\left(2y,2-1/\beta,2x\right)-Ke^{-rT}Q\left(2y,2+1/\beta,2x\right) \\ w=\frac{\left(r-q\right)}{e^{-2\beta\left(r-q\right)T}-1}, \quad x=w\frac{S_0^{-2\beta}}{\beta\alpha^2}e^{-2\beta\left(r-q\right)T}, \quad y=w\frac{K^{-2\beta}}{\beta\alpha^2} \end{gathered} \tag{16}$$

where Q is the non-central chi-square density function, carefully described in Schroder (1989). As K and S_0 appear separately in the expressions for x and y in (16) and they do not re-combine as a fraction in the computation of Q, we conclude that the CEV model price is not homogeneous of degree 1 and hence the CEV model is indeed <u>not</u> scale invariant.

Example 4: The modified CEV model

Define the modified CEV model as:

$$\frac{dS}{S} = rdt + \sigma_0 \left(\frac{S}{S_0} \right)^{\beta} dW \tag{17}$$

where we have introduced a new parameterization of (15) by setting $\alpha = \frac{\sigma_0}{S_0^{\beta}}$. As α, β, σ_0 and S_0 are known at calibration time, the modified CEV model produces exactly the same prices as the CEV model (15). But model (17) is scale invariant because the right-hand side is a function of S/S_0 and not of S or S_0 separately. In effect, the option price of the modified CEV model is the same as in (16) except that x and y become:

$$x = w \frac{1}{\beta \sigma_0^{\,2}} e^{-2\beta(r-q)T}, \quad y = w \frac{\left(K/S_0\right)^{-2\beta}}{\beta \sigma_0^{\,2}} \tag{18}$$

after replacing α so that the option price is homogeneous since x and y are invariant when we scale K and S_0 by the same amount. That is, Property 2 is satisfied and the modified CEV model <u>is</u> scale invariant. At this stage this distinction may seem pointless to the reader since (17) is simply a re-parameterization of (15), but its relevance will be clarified shortly.

Example 5: Local volatility models

In a typical 'local' volatility model, the spot volatility is defined as a deterministic function of time and the asset price (Dupire (1994), Derman and Kani (1994) and Rubinstein (1994)). Thus the asset price dynamics under the risk-neutral measure are written as:

$$\frac{dS}{S} = rdt + \sigma\left(t, S\right) dW$$

where $\sigma\left(t, S\right)$ is a deterministic function of t, S and possibly other non-random parameters.

Local volatility models can be implemented using simple lattices and there are analytical solutions in some cases, hence they are very popular among practitioners. There is only one source of randomness in the model and the market is complete because perfect (dynamic) replication with the underlying asset and a money market account is possible. Of course whether or not this is realistic has led to intense debate in the literature, as in Dumas, Fleming and Whaley (1998) and Hagan *et al* (2002). We return to this issue in Section IV.

Local volatility models are <u>not</u> scale invariant in general. Yet Property 4 implies that local volatility models in which the spot volatility is a function of time and S/S_0 and not of S or S_0 separately are scale invariant. For instance, both the CEV and the modified CEV models discussed above are local volatility models, but only the modified CEV model is scale invariant.

In local volatility models that are not scale invariant the local volatility moves along the tree when S_0 moves. In other words the local volatility surface is fixed at the time of the calibration and the forward instantaneous conditional volatilities are 'locked in' by the current prices of vanilla options. See Dupire (1996) for further explanation of this point. However in scale invariant local volatility models, Property 4 implies:

$$\sigma_L{}^2\left(t,s;S_0\right)=\sigma_L{}^2\left(t,s/S_0;1\right)$$

Hence the spot volatility is:

$$\sigma_0\equiv\sigma\left(t\right)\Big|_{t=0}=\lim_{t\to0,s\to S_0}\sigma_L\left(t,s;S_0\right)=\sigma_L\left(0,S_0;S_0\right)=\sigma_L\left(0,1;1\right) \tag{19}$$

but this does not imply that the spot volatility σ_0 is independent of S_0. For instance, in the modified CEV model above:

$$\sigma_L\left(t,s;S_0\right)=\sigma_0\left(\frac{s}{S_0}\right)^{\beta}\Rightarrow\frac{\partial\sigma}{\partial S}\bigg|_{t=0}=\lim_{t\to0,s\to S_0}\frac{\partial}{\partial s}\left(\sigma_0\left(\frac{s}{S_0}\right)^{\beta}\right)=\lim_{t\to0,s\to S_0}\frac{\sigma_0\beta}{S_0}\left(\frac{s}{S_0}\right)^{\beta-1}=\frac{\sigma_0\beta}{S_0}$$

and in general:

$$\frac{\partial\sigma\left(t\right)}{\partial S}\bigg|_{t=0}=\lim_{t\to0,s\to S_0}\frac{\partial}{\partial s}\left(\sigma_L\left(t,s;S_0\right)\right)=\lim_{t\to0,s\to S_0}\frac{\partial}{\partial s}\left(\sigma_L\left(t,s/S_0;1\right)\right)=\lim_{t\to0,x\to1}\frac{\partial}{\partial x}\left(\sigma_L\left(t,x;1\right)\right)\frac{1}{S_0}\neq0$$

Indeed, to deduce from (19) that σ_0 is independent of S_0 would be to ignore the dynamic relationship between volatility and price at the time of calibration. This is inconsistent with the very idea of local volatility, which aims to model the dependence between volatility and asset price at *every* point in time. In fact, ignoring their dependence at time 0 would lead one to conclude that the only consistent definition of a scale invariant local volatility model is the Black-Scholes model. To see why, suppose σ_0 were indeed independent of S_0. Then all spot volatilities in the future would have to be σ_0. Hence either volatility would be constant and we return to the Black-Scholes model, or when S_0 moves the whole local volatility surface must also move.[vi] Note that this argument would not apply to SIV models with jumps or stochastic volatility because there are other sources of uncertainty in the model. But in local volatility models the forward volatility is assumed to be known with certainty, hence if the spot volatility and the current price *were* independent, the only scale invariant local volatility that would not open to arbitrage would be the BS model.

Example 6: Stochastic volatility models

Consider the general stochastic volatility model (13). Conditioning on a particular volatility path, it can be shown that the option price may be written:[vii]

$$f_{SV}\left(S_0,K,T\right)=E\left[f_{BS}\left(S_0\xi_T,K,T,\sqrt{\bar{\sigma}_T^2}\right)\right] \tag{20}$$

where the expectation is with respect to the volatility path under the risk-neutral measure, and:

$$\xi_T = \exp\left(\int_0^T \varrho\sigma dZ - \tfrac{1}{2}\int_0^T \varrho^2\sigma^2 dt\right) \quad \text{and} \quad \bar{\sigma}_T^2 = \frac{1}{T}\int_0^T \left(1-\varrho^2\right)\sigma^2 dt$$

It follows that:

$$f_{SV}\left(uS_0, uK, T\right) = E\left[f_{BS}\left(uS_0\xi_T, uK, T, \sqrt{\bar{\sigma}_T^2}\right)\right] = uE\left[f_{BS}\left(S_0\xi_T, K, T, \sqrt{\bar{\sigma}_T^2}\right)\right] = uf_{SV}\left(S_0, K, T\right)$$

so that the homogeneity property has again been verified and we conclude that the general stochastic volatility model (13) is scale invariant.

Example 7: Jump-diffusion models

We have seen that it is possible to specify a SIV model that combines a diffusion process with jumps. In particular, consider the jump diffusion model:

$$\frac{dS}{S} = \left(\alpha - \lambda k\right)dt + \sigma dW + dq$$

where dq defines a Poisson jump with random intensity, independent of dW by assumption, and $\lambda k dt = E\left[dq\right]$ is the expected jump over an infinitesimal time-step dt. Assuming lognormal jump size, Merton (1976) shows that the option price given by:

$$f_{JD}\left(S_0, K, T\right) = \sum_{n=0}^{\infty} \frac{e^{-\lambda' T}\left(\lambda' T\right)^n}{n!} f_{BS}\left(S_0, K, T, r_n, \sqrt{v_n^2}\right) \tag{21}$$

where $\lambda' \equiv \lambda\left(1+k\right)$ and r_n and v_n^2 are the adjusted risk-free rate and volatility as defined in that paper. Since r_n and v_n^2 are not functions of asset price S_0 or strike K it follows that (21) can be regarded as an infinite BS mixture model and that Merton's jump diffusion model is scale invariant.

Example 8: Lévy models

Suppose ln X is a Lévy process whose characteristic function is given by the Lévy-Khintchine representation as:[viii]

$$\varphi_t\left(v\right) = \exp\left[iv\omega t - \tfrac{1}{2}v^2\sigma^2 t + t\int_{-\infty}^{\infty}\left(iv\varepsilon - 1\right)\mu\left(\varepsilon\right)d\varepsilon\right]$$

When the drift ω and the Lévy density $\mu\left(\varepsilon\right)$ are not functions of S_0, neither will be $\varphi_t\left(v\right)$, and Property 1 implies that such a Lévy process is scale invariant. In fact, following Carr and Madan (1999) and Lewis (2000), Gatheral (2004) provides a formula for the European option price by inverting the characteristic function:

$$f\left(S_0, K, T\right) = S_0 - \sqrt{S_0 K}\frac{1}{\pi}\int_0^{\infty}\frac{dv}{v^2 + \tfrac{1}{4}}\text{Re}\left[\exp\left(-iv\ln\frac{K}{S_0}\right)\varphi_T\left(v - \tfrac{i}{2}\right)\right]$$

where $\varphi_T\left(v\right)$ is the characteristic function of ln X at time T. Then, since $\varphi_T\left(v\right)$ is not a function of S_0, we have:

$$f\left(uS_0, uK, T\right) = uS_0 - \sqrt{uS_0 uK}\,\frac{1}{\pi}\int_0^\infty \frac{dv}{v^2 + \frac{1}{4}}\,\mathrm{Re}\left[\exp\left(-iv\ln\frac{uK}{uS_0}\right)\varphi_T\left(v - \tfrac{i}{2}\right)\right] = uf\left(S_0, K, T\right)$$

and the model defined by $\varphi_T(v)$ is scale invariant because it satisfies the homogeneity property. Gatheral (2004) assumes zero interest rates yet adding non-zero and even stochastic interest rates should not change this conclusion unless the rates are functions of S or S_0.

In summary, scale invariance is a very general property that is common to many volatility models. It includes all stochastic and local volatility models where the parameters are a deterministic function of t and X, as well as more complex models that mix jumps in price with jumps in volatility or with stochastic volatility. The pricing and hedging of a number of these more complex models have been compared in Bakshi, Cao and Chen (1997) and it is notable that these authors used minimum variance hedge ratios in their study: indeed the standard delta and gamma are model free, as we have shown here.

III. The SABR model paradox

The 'stochastic-αβϱ' or SABR model, described by Hagan *et al.* (2002), is essentially an extension to the CEV model (15) where the alpha parameter is assumed stochastic and follows a diffusion that is correlated with dW (hence the nomenclature). It appears that practitioners are increasingly adopting this model as the new market standard. It is not a SIV model but it may be re-parameterized into SIV form, as we demonstrate here.

The SABR model price dynamics at time $t > 0$ are written as:

$$\begin{aligned} dF &= \alpha F^{\beta} dW \qquad && F(0) = F_0 \\ d\alpha &= \upsilon \alpha dZ \qquad && \alpha(0) = \alpha_0 \qquad \langle dW, dZ \rangle = \varrho dt \end{aligned} \tag{22}$$

where $F = S\,e^{\mu(T-t)}$ is the forward price at time $T > t$ with constant carry cost μ. Applying Itô's lemma, we have:

$$\begin{aligned} \frac{dS}{S} &= \mu dt + \alpha F^{\beta-1} dW \qquad && S(0) = S_0 = F_0 e^{-\mu T} \\ d\alpha &= \upsilon \alpha dZ \qquad && \alpha(0) = \alpha_0 \qquad \langle dW, dZ \rangle = \varrho dt \end{aligned} \tag{23}$$

so that the SABR model is <u>not</u> scale invariant because S appears in the r.h.s. of the price process through F.

This can be readily verified from the approximation to the implied volatility for the SABR model given by:[ix]

$$\theta\left(F_0, K, T; \alpha_0\right) \approx \frac{\alpha_0}{\left(F_0 K\right)^{(1-\beta)/2}}\left(\frac{z}{\chi(z)}\right)\frac{1 + \left[\frac{(1-\beta)^2}{24}\frac{\alpha_0^{\,2}}{\left(F_0 K\right)^{1-\beta}} + \frac{1}{4}\frac{\varrho\beta\upsilon\alpha_0}{\left(F_0 K\right)^{(1-\beta)/2}} + \frac{2-3\varrho^2}{24}\upsilon^2\right]T}{1 + \frac{(1-\beta)^2}{24}\ln^2\frac{F_0}{K} + \frac{(1-\beta)^4}{1920}\ln^4\frac{F_0}{K}} \tag{24}$$

where z and $\chi(z)$ are defined as:

$$z = \frac{\upsilon}{\alpha_0}\left(F_0K\right)^{(1-\beta)/2}\ln\frac{F_0}{K} \qquad \chi(z) = \ln\left\{\frac{\sqrt{1-2\varrho z + z^2} + z - \varrho}{1-\varrho}\right\}$$

From (24) we conclude that it is only when β = 1 that the SABR model (22) or (23) is scale invariant because the implied volatility $\theta(F_0, K, T; \alpha_0)$ is a function of $(F_0K)^{1-\beta}$. This is exactly when F drops from the r.h.s. of (23).

Following an argument similar to that in examples 3 and 4 above we introduce a new parameterization by defining $\sigma = \alpha F^{\beta-1}$ and then apply Itô's lemma to write this 'modified SABR' model as a stochastic volatility model without the power law in the price process:

$$\begin{aligned} \frac{dS}{S} &= \mu dt + \sigma dW \qquad S_0 = F_0 e^{-\mu T} \qquad \sigma_0 = \alpha_0 F_0^{\beta-1} \\ \frac{d\sigma}{\sigma} &= (\beta-1)\sigma\left[\left(\upsilon\varrho + \tfrac{1}{2}(\beta-2)\sigma\right)dt + dW\right] + \upsilon dZ \end{aligned} \tag{25}$$

Because there is no F on the right-hand side of both s.d.e.s in (25) this model <u>is</u> scale invariant whilst the SABR model is <u>not</u> scale invariant. Formally, replace $\alpha_0 = \sigma_0 F_0^{1-\beta}$ in (24) and note that the approximation for the implied volatility becomes a function of $\frac{F_0}{K} = \frac{S_0}{K}e^{\mu T}$ so that model (25) is indeed scale invariant by Property 5.

Model (25) is just a re-parameterization of the SABR model (22) so both models always produce the same options prices and they should have identical hedge ratios. Similarly, the CEV model (15) and the modified CEV model (17) always produce the same options prices and again should have the same hedge ratios, since one is just a re-parameterization of the other. However, since one model is scale invariant and the other is not, the sensitivities with respect to S_0 can be different and indeed they are, as can be easily verified.[x] Since the standard definition of delta is the first derivative with respect to S_0, we conclude that the SABR and the modified SABR models have different deltas (and likewise the CEV and the modified CEV models have different deltas). Hence these models, which should have identical price hedge ratios, have different deltas. The paradox arises from the definition of delta as the first derivative with respect to S_0. In the next section we consider the minimum variance hedge ratio and thus show how this paradox can be resolved.

IV. Dynamic hedging

In this section we first investigate the properties of SIV models for dynamic delta and delta-gamma hedging of vanilla options, i.e. standard European calls and puts, showing that the first and second derivatives of the option price with respect to the current spot price give delta and gamma hedge ratios that are model free. But an optimal hedge ratio for dynamic hedging should capture the *dynamic* relationship between the option price and the underlying price. This may require an adjustment to the standard delta and gamma to account for extra dynamic features such as jumps or any dependence between the underlying price and the other model parameters, including volatility. If the price and volatility are correlated then partial derivatives do not provide optimal price hedge ratios for delta or delta-gamma hedge strategies. This motivates the use of 'minimum variance' (MV) delta and gamma, which may be different from standard delta and gamma in SIV models, for instance when the price and volatility are correlated.

Our hedging study focuses on vanilla calls and puts only. This can be justified by noting that vanilla options can be used to replicate statically the payoff of exotic options. For instance, the price of a digital call can be written in terms of the prices of vanilla call options as:

$$D\left(S_0,K,T\right)=-\frac{\partial C\left(S_0,K,T\right)}{\partial K}=\lim_{\varepsilon\downarrow 0}\frac{C\left(S_0,K-\varepsilon,T\right)-C\left(S_0,K+\varepsilon,T\right)}{2\varepsilon}$$

Therefore, if vanilla calls are correctly priced by a model (because of direct calibration) then differentiating their model prices with respect to K should produce the correct prices for digital calls. This rule extends (at least) to all European-style claims whose payoff can be statically replicated with vanilla options, such as asset-or-nothing, chooser options, etc (see Derman, Ergener and Kani (1995) and Carr, Ellis and Gupta (1998)).

Likewise, if a model provides an estimate for the hedge ratios of vanilla options, then static replication allows the estimation of the hedge ratios for exotic options. For instance, the delta of a digital call option can be defined as:

$$\frac{\partial D\left(S_0,K,T\right)}{\partial S_0}=-\frac{\partial}{\partial S_0}\left(\frac{\partial C}{\partial K}\right)=-\frac{\partial}{\partial K}\left(\frac{\partial C}{\partial S_0}\right)=\lim_{\varepsilon\downarrow 0}\frac{1}{2\varepsilon}\left(\frac{\partial C\left(S_0,K-\varepsilon,T\right)}{\partial S_0}-\frac{\partial C\left(S_0,K+\varepsilon,T\right)}{\partial S_0}\right)$$

The delta of a digital call is obtained by differentiating the model price of a vanilla call on the same strike with respect to S_0 and K at the same time, and this extends to other exotic options from static replication. In short, it is important to derive the correct hedge ratios for vanilla options because these are often used as building blocks for more sophisticated instruments. By implication, if a volatility model cannot produce realistic hedge ratios even for vanilla options it is unlikely that it would do any better for exotic options.

IV.1 Scale invariant hedge ratios

The fact that all SIV models have the same deltas and the same gammas follows immediately from Property 3 in Section I. This showed that the first derivative of the standard European option price given by any SIV model with respect to the underlying asset price is:

$$\frac{\partial f(S_0,K,T)}{\partial S_0}=\frac{1}{S_0}\left(f(S_0,K,T)-K\frac{\partial f(S_0,K,T)}{\partial K}\right) \tag{26}$$

All SIV models should give the same sensitivity when calibrated to the same options, because the option price $f(S_0,K,T)$ and its derivatives with respect to the strike K are observable quantities (assuming a continuum of options). Any difference between the models' sensitivities can only be justified by a different fit to option prices. The result (26) motivates the definition of the 'SIV delta', δ_{SIV} as the first derivative of the standard European option price with respect to the underlying asset price:

$$\delta_{SIV}(S_0,K,T)=\frac{\partial f(S_0,K,T)}{\partial S_0} \tag{27}$$

and we have shown that this is model free. SIV models also have identical higher derivatives. For instance, from (6) the 'model free' SIV gamma is given by:

$$\gamma_{SIV}(S_0,K,T)=\frac{\partial^2 f(S_0,K,T)}{\partial {S_0}^2}=\left(\frac{K}{S_0}\right)^2\frac{\partial^2 f(S_0,K,T)}{\partial K^2} \tag{28}$$

Empirically, of course, they may differ but this can only be due to models having different fits to option prices.

Note that if the SIV model is calibrated to market data one can derive a useful relationship between the SIV delta and the Black-Scholes (BS) delta. Write:

$$f(S_0,K,T)=f_{BS}\left(S_0,K,T,\theta(S_0,K,T)\right)$$

and after differentiation with respect to S_0 we have:

$$\frac{\partial f(S_0,K,T)}{\partial S_0}=\frac{\partial f_{BS}\left(S_0,K,T,\theta(S_0,K,T)\right)}{\partial S_0}+\frac{\partial f_{BS}\left(S_0,K,T,\theta(S_0,K,T)\right)}{\partial \theta}\frac{\partial \theta(S_0,K,T)}{\partial S_0}$$

that is

$$\delta_{SIV}(S_0,K,T)=\delta_{BS}\left(S_0,K,T,\theta(S_0,K,T)\right)+v_{BS}\left(S_0,K,T,\theta(S_0,K,T)\right)\frac{\partial \theta(S_0,K,T)}{\partial S_0} \tag{29}$$

where δ_{BS} and v_{BS} are the Black-Scholes (BS) model delta and vega respectively.

IV.2 Minimum variance hedge ratios

The above sensitivities to S_0 are not the optimal hedge ratios when other sources of price uncertainty are not hedged. For instance, when prices can jump or when the underlying price and volatility are correlated, the above delta and gamma do not capture the full extent of the option price change when the underlying price changes. In a delta hedge strategy, the 'minimum variance' (MV) delta is the optimal amount δ of the underlying asset that reduces the covariance of a delta-hedged portfolio $\Pi = f - \delta S$ with the underlying asset S to zero. That is,

$$0 = \langle d\Pi, dS \rangle = \langle df - \delta dS, dS \rangle = \langle df, dS \rangle - \delta \langle dS, dS \rangle$$
$$\Rightarrow \delta = \frac{\langle df, dS \rangle}{\langle dS, dS \rangle}$$

where we assume a self-financing portfolio and, without loss of generality, that S pays no dividend. In the Black-Scholes model, the MV delta is the same as the first derivative of the option price with respect to S_0 but this is not the case for all models. In particular if any model component such as the volatility or interest rates is correlated with the asset price, the MV delta differs from the usual delta and the latter is not the optimal hedge ratio.

Consider first the case that the spot volatility σ (or variance) is a continuous and stochastic process itself. The dynamics of the option price $f(t, S, \sigma)$ are given by Ito's lemma as:

$$df = \frac{\partial f}{\partial t}dt + \frac{\partial f}{\partial S}dS + \frac{\partial f}{\partial \sigma}d\sigma + \tfrac{1}{2}\frac{\partial^2 f}{\partial S^2}dS^2 + \tfrac{1}{2}\frac{\partial^2 f}{\partial \sigma^2}d\sigma^2 + \frac{\partial^2 f}{\partial S \partial \sigma}dS d\sigma$$

where the cross-terms are of order *dt*. Therefore, the minimum variance (MV) delta of the stochastic volatility model is:

$$\delta_{SV} = \frac{\langle df, dS \rangle}{\langle dS, dS \rangle} = \frac{\left\langle \frac{\partial f}{\partial S}dS + \frac{\partial f}{\partial \sigma}d\sigma, dS \right\rangle}{\langle dS, dS \rangle} = \frac{\partial f}{\partial S} + \frac{\partial f}{\partial \sigma}\frac{\langle d\sigma, dS \rangle}{\langle dS, dS \rangle} \tag{30}$$

where the first term on the right-hand side is the usual definition of delta and the second term is non-zero if and only if the price-volatility correlation is non-zero. Computing the second term is normally easy once the processes for S and σ are specified. When correlation is strong this term may be large and the usual delta can be very different from the MV delta.

Intuitively, the MV delta (30) resembles a total derivative of the option price $f(t, S, \sigma)$ with respect to S:

$$\delta_{SV} = \frac{df}{dS} = \frac{\partial f}{\partial S} + \frac{\partial f}{\partial \sigma}\frac{d\sigma}{dS} \tag{31}$$

where the total derivatives are defined in expectation as:

$$\frac{df}{dS} \equiv \frac{\langle df, dS \rangle}{\langle dS, dS \rangle} \quad \text{and} \quad \frac{d\sigma}{dS} \equiv \frac{\langle d\sigma, dS \rangle}{\langle dS, dS \rangle}$$

The motivation for writing the MV delta as in (31) is because, when S moves, we expect a simultaneous movement in σ if it is correlated with S. Therefore, the total derivative $d\sigma / dS$ here denotes the *expected* change in volatility when S changes.

In a (univariate) stochastic volatility model vega hedging with another option completes the market and we can use the standard delta to hedge the remaining delta risk from the underlying. That is, partial derivative delta hedges remain optimal in stochastic volatility models if separate vega hedging is applied, as in Bakshi and Kapadia (2003) but not as in Yung and Zhang (2003). In effect, the term $d\sigma / dS$ in (31) is already captured by the vega hedge. But if we do not vega hedge then, when the price is correlated with the volatility, we have delta risk coming from both the underlying and the volatility. Hence we need MV hedge ratios instead of the standard delta or gamma.

Similarly, we define the 'minimum variance' (MV) gamma as:

$$\gamma_{SV} = \frac{d^2 f}{dS^2} = \frac{d}{dS}\left(\frac{\partial f}{\partial S} + \frac{\partial f}{\partial \sigma}\frac{d\sigma}{dS} \right) = \frac{\partial^2 f}{\partial S^2} + 2\frac{\partial^2 f}{\partial S \partial \sigma}\frac{d\sigma}{dS} + \frac{\partial^2 f}{\partial \sigma^2}\left(\frac{d\sigma}{dS} \right)^2 + \frac{\partial f}{\partial \sigma}\frac{d^2\sigma}{dS^2} \tag{32}$$

where the first term in the right-hand side is the usual gamma and the remaining terms are adjustments that account for the dependence between σ and S. The total derivatives are defined in expectation as before. In particular, in the last term of (32) $d^2\sigma / dS^2$ denotes the expected change in $d\sigma / dS$ when S moves. Taking the limit when $t \downarrow 0$ in (31) and (32) we have the optimal MV hedge ratios to be used in delta and delta-gamma hedge strategies.

Example 9: MV hedge ratios for the Heston model

Consider the Heston (1993) model:

$$\begin{aligned} &\frac{dS}{S} = r dt + \sqrt{V}\, dW \\ &dV = a\left(m - V\right)dt + b\sqrt{V}\, dZ \quad \langle dW, dZ \rangle = \varrho dt \end{aligned} \tag{33}$$

Then (30) gives

$$\frac{df_{Heston}}{dS} = \frac{\partial f_{Heston}}{\partial S} + \frac{\partial f_{Heston}}{\partial V}\frac{\langle b\sqrt{V}dZ, S\sqrt{V}dW \rangle}{\langle S\sqrt{V}dW, S\sqrt{V}dW \rangle} = \frac{\partial f_{Heston}}{\partial S} + \frac{\partial f_{Heston}}{\partial V}\frac{\varrho b}{S} \tag{34}$$

and when $t \downarrow 0$ we have the MV delta:

$$\delta_{MV}\left(S_0,K,T;V_0\right)=\frac{df_{Heston}}{dS_0}=\delta_{SIV}+\frac{\partial f_{Heston}}{\partial V_0}\frac{\varrho b}{S_0} \tag{35}$$

where we used (27) to replace the partial price derivative by δ_{SIV} because the Heston model is scale invariant. This emphasises that the only model-dependent part of (35) is the second term in the right-hand side.

Likewise, the MV gamma in the Heston model is given by:

$$\gamma_{MV}\left(S_0,K,T;V_0\right)=\frac{d^2 f_{Heston}}{dS_0^{\;2}}=\frac{\partial^2 f_{Heston}}{\partial S_0^{\;2}}+\frac{\varrho b}{S_0}\left(2\frac{\partial^2 f_{Heston}}{\partial S_0 \partial V_0}+\frac{\varrho b}{S_0}\frac{\partial^2 f_{Heston}}{\partial V_0^{\;2}}-\frac{1}{S_0}\frac{\partial f_{Heston}}{\partial V_0}\right) \tag{36}$$

where $\frac{\partial^2 f_{Heston}}{\partial S_0^{\;2}}=\gamma_{SIV}$ is the model free gamma (28).

Next, setting:

$$f_{Heston}\left(S_0,K,T;V_0\right)=f_{BS}\left(S_0,K,T,\theta\left(S_0,K,T;V_0\right)\right)$$

and using the chain rule, we have:

$$\frac{df_{Heston}}{dS_0}=\frac{\partial f_{BS}}{\partial S_0}+\frac{\partial f_{BS}}{\partial \theta}\frac{d\theta}{dS_0}\Rightarrow\frac{d\theta}{dS_0}=\frac{\delta_{MV}-\delta_{BS}}{v_{BS}} \tag{37}$$

which describes the *expected* change in implied volatility when S_0 moves, in terms of the MV delta of the Heston model and the BS delta and vega.

Figure 3(a) generates an implied volatility skew from the Heston model and figure 3(b) depicts the model's total and partial derivatives of implied volatility with respect to S_0, given by (37) and (29) respectively.[xi] Here δ_{MV} denotes the minimum variance (MV) delta hedge in the Heston model. From (35), we have that $\delta_{MV}=\delta_{SIV}\Leftrightarrow\varrho=0$. Figure 3(c) compares the BS delta δ_{BS} to the total sensitivity δ_{MV} and the partial sensitivity δ_{SIV} of the Heston model using typical parameters for equity index options. Clearly δ_{SIV} is greater than the BS delta but because the chosen correlation ϱ is negative and reasonably large δ_{MV} is below the BS delta. Figure 3(d) compares the BS gamma, the SIV gamma and Heston MV gamma for the same data as above. Again, it is only when $\varrho = 0$ that the SIV gamma is correct because the second term on the right hand side of (36) is zero.

[Figure 3 here]

IV.3 Local volatility hedge ratios

Some of the literature on the effectiveness of delta and delta-gamma hedging with local volatility models investigates scale invariant local volatility models. Hagan *et al.* (2002) show that scale invariant local volatility models predict the wrong dynamics for implied volatility and this explains their poor hedging performance. But in these models the standard price hedge ratios are model free, as we have shown above, so the conclusions in Hagan *et al.* (2002) extend to other SIV models when standard price hedge ratios are used. What about the MV price hedge ratios? In scale invariant local volatility models, the standard price hedge ratios are not the same as the MV price hedge ratios because there is perfect correlation between price and volatility. But most local volatility models are not scale invariant – what hedging properties do these models have? This section examines the price hedge ratios of both scale invariant and non-scale invariant local volatility models.

There have been a number of empirical studies of hedging with local volatility models, all employing standard price hedge ratios rather than MV price hedge ratios. These include Dumas, Fleming and Whaley (1998), Coleman *et al.* (2001) and Crépey (2004) and their findings are controversial. Dumas, Fleming and Whaley (1998) test several different parametric and semi-parametric forms of local volatility function. They calibrate the models to S&P 500 index options prices on a particular date, repeating this on a weekly basis, and subsequently compare the hedging performance of the local volatility models with that of the BS model. Their conclusion is that the BS deltas appear to be more reliable than any of the deltas from the local volatility models that they tested. McIntyre (2001) reaches a similar conclusion, although Coleman et al. (2001) find that over long hedging periods the local volatility deltas do improve somewhat. Crépey (2004) examines the hedging performance of local volatility deltas in four equity market regimes – slow rallies, fast rallies, slow sell-offs and fast sell-offs – and concludes that, on average, local volatility deltas are more effective than BS deltas.

All the studies mentioned above used the 'standard' definition for price hedge ratios. We now derive the MV hedge ratios for local volatility models. In the stochastic volatility case we used the second source of randomness from the volatility process to motivate an adjustment to the hedge ratios, but in local volatility models there is just one source of randomness so this technique cannot be applied immediately. Instead we use a simple trick to solve this problem. Because the spot volatility $\sigma(t,S)$ in a local volatility model is a function of S, it is a process itself and it has dynamics given by Itô's lemma as:

$$d\sigma = \left(\frac{\partial \sigma}{\partial t} + \tfrac{1}{2}\sigma^2 S^2 \frac{\partial^2 \sigma}{\partial S^2}\right)dt + \frac{\partial \sigma}{\partial S}dS$$

which can be interpreted as a stochastic volatility model with perfect correlation between the volatility and the underlying asset price. Therefore, following (30) and (32) the MV local volatility hedge ratios are:[xii]

$$\delta_{LV}\left(S_0,K,T;\sigma_0\right) = \frac{df_{LV}\left(S_0,K,T;\sigma_0\right)}{dS_0} = \frac{\partial f_{LV}\left(S_0,K,T;\sigma_0\right)}{\partial S_0} + \frac{\partial f_{LV}\left(S_0,K,T;\sigma_0\right)}{\partial \sigma_0}\left.\frac{\partial \sigma}{\partial S}\right|_{t=0}$$

$$\gamma_{LV}\left(S_0,K,T;\sigma_0\right) = \frac{d^2 f_{LV}\left(S_0,K,T;\sigma_0\right)}{dS_0^{\ 2}} = \frac{\partial^2 f_{LV}}{\partial S_0^{\ 2}} + 2\frac{\partial^2 f_{LV}}{\partial S_0 \partial \sigma_0}\left.\frac{\partial \sigma}{\partial S}\right|_{t=0} + \frac{\partial f_{LV}}{\partial \sigma_0^{\ 2}}\left.\left(\frac{\partial \sigma}{\partial S}\right)^2\right|_{t=0} + \frac{\partial f_{LV}}{\partial \sigma_0}\left.\frac{\partial^2 \sigma}{\partial S^2}\right|_{t=0}$$

Example 10: Relationship between MV hedge ratios for Heston and modified CEV models

The modified CEV model is a scale invariant local volatility model defined in example 4 in Section II. In this model:

$$\sigma(t,S)=\sigma_0\left(\frac{S}{S_0}\right)^{\beta}\Rightarrow\frac{\partial\sigma(t,S)}{\partial S}=\frac{\sigma(t,S)\beta}{S}\Rightarrow\frac{\partial^2\sigma(t,S)}{\partial S^2}=\frac{\sigma(t,S)\beta(\beta-1)}{S^2}$$

so that:

$$\delta_{MCEV}(S_0,K,T;\sigma_0)=\frac{\partial f_{MCEV}(S_0,K,T;\sigma_0)}{\partial S_0}+\frac{\sigma_0\beta}{S_0}\frac{\partial f_{MCEV}(S_0,K,T;\sigma_0)}{\partial\sigma_0} \tag{38}$$

$$\gamma_{MCEV}(S_0,K,T;\sigma_0)=\frac{\partial^2 f_{MCEV}}{\partial S_0^{\,2}}+\frac{\sigma_0\beta}{S_0}\left(2\frac{\partial^2 f_{MCEV}}{\partial S_0\partial\sigma_0}+\frac{\sigma_0\beta}{S_0}\frac{\partial f_{MCEV}}{\partial\sigma_0^{\,2}}+\frac{\beta-1}{S_0}\frac{\partial f_{MCEV}}{\partial\sigma_0}\right) \tag{39}$$

We investigate the connection between the hedge ratios (38) and (39) and those derived for another SIV model, i.e. the Heston model in example 9 above. Since both are SIV models we know that their 'standard' delta and gamma are theoretically identical if both models fit the smile equally well. Indeed, in this case there is an intuitive mapping between the parameters of the two models. Assuming both models fit the smile equally well we have:

$$\frac{\partial f_{MCEV}(S_0,K,T;\sigma_0)}{\partial S_0}=\frac{\partial f_{Heston}(S_0,K,T;V_0)}{\partial S_0}.$$

because both models are SIV. Also $f_{MCEV}(S_0,K,T;\sigma_0)=f_{Heston}(S_0,K,T;V_0)\quad\forall K,T$ so

$$\frac{\partial}{\partial V_0}\left(f_{Heston}(S_0,K,T;V_0)\right)=\frac{\partial}{\partial V_0}\left(f_{MCEV}(S_0,K,T;\sigma_0)\right)=\frac{\partial f_{MCEV}}{\partial\sigma_0}\frac{\partial\sigma_0}{\partial V_0}=\frac{\partial f_{MCEV}}{\partial\sigma_0}\frac{1}{2\sigma_0}$$

where V_0 is the spot variance in the Heston model and σ_0 is the spot volatility at time 0. We conclude that the MV deltas will also be identical under the following parameter constraint:

$$\delta_{MCEV}(S_0,K,T;\sigma_0)=\delta_{Heston}(S_0,K,T;V_0)\Leftrightarrow\beta=\frac{\varrho b}{2V_0} \tag{40}$$

However the argument above assumes that both models fit the smile equally well and this is unlikely to be true. The Heston model (33) has five parameters whilst the MCEV model has only two. The parameter β, which is responsible for the fit to the skew of the implied volatility smile, mixes the roles of the correlation ϱ and of the 'vol-of-vol' *b* in the Heston model. In particular, a negative value for ϱ requires a *negative* value for β, and this is exactly what is observed when we calibrate the MCEV model to S&P 500 index options. This means that the MCEV cannot separately describe the mean-reversion (captured by *a* and *m* in the Heston model) and the higher moments (skewness and kurtosis, captured by ϱ and *b*). That is, the MCEV is under-parameterized and in general will not fit the smile as well as the Heston model.

Recall that whenever local volatility is a function of S/S_0 and not of S or S_0 separately the local volatility model is scale invariant, see example 5. Then the standard deltas and gammas, which are model free, require an adjustment to become minimum variance hedge ratios, i.e. to account for the dependence between price and volatility at the time of calibration. But in most local volatility models the local volatility is a function of S but not of S_0 explicitly, and it turns out that no such adjustment is needed. To illustrate this point, consider the CEV model where:

$$\sigma_L{}^2(t,s) = E\left[\sigma^2(t,S)\middle| S=s\right] = E\left[\left(\alpha S^\beta\right)^2\middle| S=s\right] = \left(\alpha s^\beta\right)^2 \Leftrightarrow \sigma_L(t,s) = \alpha s^\beta$$

Now it is simple to verify that the MV delta of the MCEV model is equal to the standard delta of the CEV model, that is:

$$\delta_{MCEV} = \frac{df_{MCEV}}{dS_0} = \frac{\partial f_{CEV}}{\partial S_0} \tag{41}$$

Yet these models always produce the same option prices, so they should have the same hedge ratios. We conclude that no adjustment is necessary to make the CEV deltas into MV hedge ratios. More generally, when the local volatility is not scale invariant the 'standard' delta and gamma should be the same as the MV delta and gamma.

V.4 Hedge ratios in the SABR model

From Section III, the asset price dynamics under the SABR model can be written in the form (23). Then the MV delta is given by:

$$\delta = \frac{df}{dS} = \frac{\langle df, dS\rangle}{\langle dS, dS\rangle} = \frac{\partial f}{\partial S} + \frac{\partial f}{\partial \alpha}\frac{\langle d\alpha, dS\rangle}{\langle dS, dS\rangle} = \frac{\partial f}{\partial S} + \frac{e^{(1-\beta)\mu(T-t)}\varrho\upsilon}{S^\beta}\frac{\partial f}{\partial \alpha} \tag{42}$$

and letting $t \downarrow 0$ the MV delta may be written:

$$\delta_{SABR} = \frac{df_{SABR}}{dS_0} = \frac{\partial f_{SABR}}{\partial S_0} + \frac{e^{(1-\beta)\mu T}\varrho\upsilon}{S_0{}^\beta}\frac{\partial f_{SABR}}{\partial \alpha_0} \tag{43}$$

where $f_{SABR} \equiv f_{BS}\left(S_0, K, T, \theta\left(F_0, K, T; \alpha_0\right)\right)$ is the BS option price based on the implied volatility (24). Note that when $\beta = 1$ (43) is consistent with the stochastic volatility delta (30) where α has the same role as σ.

Likewise, differentiating (42) with respect to S and letting $t \downarrow 0$ we derive the MV gamma in the SABR model is:

$$\gamma_{SABR} = \frac{d^2 f_{SABR}}{dS_0{}^2} = \frac{\partial^2 f_{SABR}}{\partial S_0{}^2} + \frac{e^{(1-\beta)\mu T}\varrho\upsilon}{S_0{}^\beta}\left(2\frac{\partial^2 f_{SABR}}{\partial S_0 \partial \alpha_0} + \left(\frac{e^{(1-\beta)\mu T}\varrho\upsilon}{S_0{}^\beta}\right)\frac{\partial^2 f_{SABR}}{\partial \alpha_0{}^2} - \frac{\beta}{S_0}\frac{\partial f_{SABR}}{\partial \alpha_0}\right) \tag{44}$$

which is consistent with (32) when $\beta = 1$.

Note that (43) is not the same delta as the one derived by Hagan *et al.* (2002) who employ the partial derivative of the option price with respect to F_0 and not the total derivative that gives the MV delta. Using (43):

$$\frac{df_{SABR}}{dF_0} = \frac{df_{SABR}}{dS_0}\frac{dS_0}{dF_0} = \left(\frac{\partial f_{SABR}}{\partial S_0} + \frac{e^{(1-\beta)\mu T}\varrho\upsilon}{S_0^{\beta}}\frac{\partial f_{SABR}}{\partial \alpha_0}\right)e^{-\mu T} = \frac{\partial f_{SABR}}{\partial F_0} + \frac{\varrho\upsilon}{F_0^{\beta}}\frac{\partial f_{SABR}}{\partial \alpha_0}$$

hence the MV delta captures the correlation between S and α, as in other stochastic volatility models.

Briefly, we discuss the definition of the *backbone* of a model which aims to capture the dynamic relationship between ATM implied volatility and asset prices. Hagan *et al.* (2002) introduce the 'backbone' as the at-the-money ($K = F_0$) implied volatility defined as a function of the forward price at time 0:

$$\theta_{ATM} = \theta\left(F_0, K, T; \alpha_0\right)\Big|_{K=F_0}$$

However under this definition, all SIV models would have a 'flat' backbone (i.e. ATM implied volatility would be insensitive to changes in F_0) because implied volatilities are homogeneous in these models (see Property 5 in Section I). This definition would lead to the conclusion that stochastic volatility models, and some local volatility and jump models, predict the wrong dynamics for ATM implied volatility, because they have a flat backbone. Of course this cannot be true.

Indeed, similar to defining the spot volatility as in (19) so that we do not ignore the dependence between price and volatility at the time of calibration, and similar to using the total price sensitivity rather than the partial price sensitivity for hedge ratios, the backbone should be defined as:

$$\theta_{ATM} = \theta\left(F(t), K, T; \alpha(t)\right)\Big|_{K=F(t)} \tag{45}$$

where $F(t)$ and $\alpha(t)$ are *random* at time $t > 0$.

Under definition (45), the *total* sensitivity of ATM implied volatility is:

$$\frac{d\theta_{ATM}}{dF} \equiv \frac{\langle d\theta_{ATM}, dF\rangle}{\langle dF, dF\rangle} = \frac{\partial\theta_{ATM}}{\partial F} + \frac{\partial\theta_{ATM}}{\partial\alpha}\frac{\langle d\alpha, dF\rangle}{\langle dF, dF\rangle} = \frac{\partial\theta_{ATM}}{\partial F} + \frac{\varrho\upsilon}{F^{\beta}}\frac{\partial\theta_{ATM}}{\partial\alpha} \tag{46}$$

where the ATM volatility at time 0 is given by:

$$\theta_{ATM}\big|_{t=0} = \theta\left(F_0, F_0, T; \alpha_0\right) \approx \frac{\alpha_0}{F_0^{1-\beta}}\left(1 + \left[\frac{(1-\beta)^2}{24}\frac{\alpha_0^{\,2}}{F_0^{2(1-\beta)}} + \frac{1}{4}\frac{\varrho\beta\upsilon\alpha_0}{F_0^{1-\beta}} + \frac{2-3\varrho^2}{24}\upsilon^2\right]T\right) \approx \frac{\alpha_0}{F_0^{1-\beta}} \tag{47}$$

and the last approximation is possible if we assume that the term in […] is small. Now letting $t \downarrow 0$ and replacing (47) into (46) we have:

$$\frac{d\theta_{ATM}}{dF}\bigg|_{t=0} \approx (\beta - 1)\frac{\alpha_0}{F_0^{2-\beta}} + \frac{\varrho\upsilon}{F_0} \tag{48}$$

This is the slope of the backbone (in the same way that (43) is the total sensitivity of option price to S) and when $\beta < 1$ and $\varrho < 0$, (48) is negative. Hence the ATM implied volatility *decreases* if F increases, and vice-versa. Finally, integrating (48) with respect to F and using the starting condition (47) we have:

$$\theta_{ATM}(F) \approx \frac{\alpha_0}{F^{1-\beta}} + \varrho\upsilon \ln\frac{F}{F_0} \tag{49}$$

Note that the backbone is determined by both β and ϱ, and not only by β. Therefore, even when $\beta = 1$, the backbone is not flat if ϱ and υ are non-zero.

V. Empirical results

As customary, we shall finish the paper with a short empirical section to compare the hedging performance of several of the models described above, using both standard and minimum variance hedge ratios. We shall not report all of our results since the empirical testing of option pricing models is not the main focus of this paper. Nor shall we focus on perfect hedging, such as in Bakshi, Cao and Chen (1997). Our purpose here is simply to investigate whether using (a) the 'standard' price hedge ratios and (b) the minimum variance (MV) price hedge ratios leads to a significant difference between the delta and delta-gamma hedging performance of different option pricing models. Our theoretical results have shown that, under the assumption that the models' fit to option prices is equally good, all SIV models have the same (standard) delta and gamma and that even their MV hedge ratios can be identical under certain parameter constraints. However, these results only hold when the models fit the smile equally well. In practice we generally find that models with fewer parameters do not fit the smile as closely those with more parameters; but it is not clear whether a close fit to the smile is an advantage for hedging purposes.

Hence we now ask: how similar are the empirical SIV and MV hedge ratios and, where there are discernable differences in the hedge ratios, what is the effect on the hedging performance? The models we consider are: the Black-Scholes (BS) model, the entire class of SIV models using the usual 'model-free' delta (29) and gamma (28), the Heston model with minimum variance delta and gamma given by (35) and (36), the CEV model with minimum variance delta and gamma given by (38) and (39) (note also (41)), the SABR delta and gamma employed by Hagan *et al.* (2002) and, finally, the minimum variance SABR delta and gamma given by (43) and (44). For the SABR hedges we set $\beta = 0$.

We have obtained data from Bloomberg on the June 2004 European options on the SP500 index: i.e. daily close prices from 02 Jan 2004 to 15 June 2004 (111 business days) for 34 different strikes (from 1005 to 1200). Only the strikes within ±10% of the current index level were used for the model's calibration each day but all strikes were used for the hedging strategies. The delta hedge strategy consists of one delta-hedged short call in each option, rebalanced daily. That is, one call on each of the 34 strikes from 1005 to 1200 is sold on 16th January (or when the option is issued, if later than this) and hedged by buying an amount delta of the underlying asset, where delta is determined by both the model and the option's characteristics. The portfolio is rebalanced daily, stopping on 2nd June because from then until the options expiry the fit to the smile worsened considerably for most of the models. The delta-gamma hedge strategy again consists of a short call in each option, but this time an amount of the 1125 option, which is closest to ATM in general over the period, is bought. This way the gamma on each option is set to zero and then we delta hedge the portfolio as above. This option-by-option strategy on a large and complete database of liquid options allows one to assess the effectiveness of hedging by strike or moneyness of the option, and day-by-day as well as over the whole period. A data set of P&L with 1324 observations is obtained which allows an in-depth investigation of the hedging effectiveness of each model.

Each model was calibrated daily by minimizing the root mean square error (RMSE) between the model implied volatilities and the market implied volatilities of the options used in the calibration set. For the BS model the deltas and gammas are obtained directly from the market data and there is no need for model calibrations. For the Heston (1993) model we used the closed form price based on Fourier transforms (see Lewis (2000)), chose a volatility risk premium of zero and set the long-term volatility at 12%. The calculation of the CEV option price is based on the non-central chi-square distribution result of Schroder (1989). Finally, we considered several SIV models but found that their hedge ratios were so similar that it was not necessary to include them all in the results. The SIV hedge ratios shown here are based on the Brigo and Mercurio (2002) lognormal mixture model with two constant volatility components.[xiii]

The deltas and gammas of each model, whilst changing daily, exhibit some strong patterns. When they are plotted, by strike or by moneyness, on any particular day the same shapes emerge day after day. In figure 4 we compare the deltas and gammas from the different models on 21st May 2004, a day exhibiting typical patterns for the models' delta and gamma. It is well-known that the BS model *over* hedges the equity skew. The SIV 'model free' delta, i.e. the partial price sensitivity that is common to all SIV models, has an even greater delta than the BS model for all but the very high strikes, and the SABR model delta lies between the BS and SIV deltas. However, a different picture emerges when minimum variance hedge ratios are used. In the CEV and SABR models (which are not scale invariant) and in the Heston model (which is scale invariant) the deltas are generally much lower than the BS deltas and this may be the reason why their delta-hedging performance in equity markets is superior to that of the BS model, as our results below will show.

[Figure 4]

Tables I and II report the sample statistics of the aggregate daily P&L for each model, over all options and over all days in the hedging period. As mentioned above, the standard Heston delta and gamma hedges are indistinguishable from the SIV hedges and hence are not included. In both tables the models are ordered by the standard deviation of the daily P&L, since to minimize this is a prime objective of hedging. Small skewness and excess kurtosis in the P&L distribution is also desirable – high values for these sample statistics indicate that the model was spectacularly wrong on a few days in the sample. Another important performance criterion is that the P&L be uncorrelated with the underlying asset. In our case over hedging would result in a significant positive correlation between the hedge portfolio and the SP500 index return. We have therefore performed a regression, based on all 1324 P&L data points, where the P&L for each option is explained by a quadratic function of the SP500 returns. The lower the R^2 from this regression, reported in the last column of the tables, the more effective the hedge.

[Tables I and II]

According to these criteria the best delta hedging models are the CEV, SABR and Heston models, each using the minimum variance hedge ratios given by the total derivative of the option price with respect to the underlying asset price. These models outperform the BS model for delta hedging. These also have P&L that is closest to being normally distributed according to the skewness and excess kurtosis. Note that all of the usual price sensitivities that are commonly used for delta hedging perform worse than the Black-Scholes model. Apart from this the positive mean P&L is a result of the short volatility exposure and gamma effects, since we have only rebalanced daily. The delta-gamma hedge strategy results in table II show a mean P&L that is close to zero. For delta-gamma hedging it is remarkable that the BS model performs really well according to all criteria, whilst the other models ranked more or less as before. Also notable is that SABR model minimum variance hedge has the smallest R^2 in both tables.

We have more detailed results on delta-gamma hedge portfolio P&L sample statistics by strike and by moneyness, averaged over all days in the sample period, and by date averaged over all strikes. For brevity these are not reported, but are available by request. They show that the BS model performs best for low strike call options and that mid to high strike options are better hedged using the CEV and the SABR and Heston models with minimum variance delta and gamma hedges. It appears that the over hedging of the BS delta is compensated by the over-hedging of gamma in the two-instrument hedge. It is possible that the apparent superiority of the BS model for delta-gamma hedging of low strike options is a result of the gamma hedging strategy that we have chosen, yet Bakshi, Cao and Chen (1997) also find that BS performs well for low strike call options.[xiv] These authors show that once stochastic volatility is modelled, the inclusion of jumps leads to no discernable improvement in hedging performance, at least when the hedge is rebalanced frequently, because the likelihood of a jump during the hedging period is too small. They also find that the inclusion of stochastic interest rates can improve the hedging of long-dated OTM options, but for other options stochastic volatility is the most important factor to model.

Although minimum variance hedge ratios are not model free in general the similarity of the models' hedging performance using minimum variance hedge ratios is remarkable. Figure 5 plots the standard deviation of the hedging P&L averaged over all options in the data set, for the CEV model and for the SABR and Heston minimum variance hedges. The differences between the models' performance are very small except that the SABR model appears slightly worse than the other two for options with high strikes when delta-gamma hedging. This may be due to the fact that we have set $\beta = 0$. Nevertheless, the differences are not statistically significant.

[Figure 5 here]

VI. Summary and conclusions

This paper has introduced a class of scale invariant volatility (SIV) models that includes most of the option pricing models in the literature. We derived a number of properties common to all SIV models and thus proved that the partial price sensitivities, i.e. the 'standard' delta and gamma for price hedging are theoretically identical in all SIV models. The only difference between these models in terms of their delta and delta-gamma hedging performance is empirical, because when calibrated to the skew they have different fits to option prices, but the empirical differences between the hedge ratios (and consequently the hedging performance) of the SIV models that we studied were found to be very small.

However, in general the standard delta and gamma are not the same as the minimum variance (MV) hedge ratios and for each of the models we studied there were substantial differences between the standard and the MV hedge ratios (except in the BS and CEV models, where the two are identical). Our empirical results provide convincing evidence that the standard delta and gamma should be dropped in favour of MV delta and gamma, in the absence of a separate vega hedge. Applying the standard price hedge ratios (when they are different from the MV hedge ratios) led to a hedging performance that was uniformly worse than the Black-Scholes (BS) model. On the other hand the MV hedges performed very well and the Heston, SABR and CEV models all out-performed the BS model in MV delta hedging. Hence we arrive at the startling conclusion that much of the literature on delta and delta-gamma hedging with smile-consistent models should be re-examined in the light of the sub-optimal hedge ratios that were used.

It is notable that, whilst we considered models having quite different characteristics – the Heston model is a scale invariant stochastic volatility model, the SABR model is a non-scale invariant stochastic volatility model and the CEV is a non-scale invariant local volatility model – and their MV hedge ratios are not model free, there is a remarkable similarity in their MV hedging performance. In fact we found no significant differences between these models for MV delta and delta-gamma hedging purposes.

References

Andersen, L. and J. Andreasen (2000) Jump-diffusion processes: volatility smile fitting and numerical methods for option pricing. *Review of Derivatives Research* 4:3, 231-262.

Avellaneda, M., Levy, A. and A. Parás (1995) Pricing and hedging derivative securities in markets with uncertain volatilities. *Applied Mathematical Finance*, 2, 73-88.

Bakshi, G., C. Cao and Z. Chen (1997) Empirical performance of alternative option pricing models. *Journal of Finance* 52, 2003-2049.

Bakshi, G. and N Kapadia (2003) Delta hedged gains and the negative volatility risk premium. *Review of Financial Studies* 16, 527-566.

Black, F. and M. Scholes (1973) The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637-659.

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.

Carr, P., K. Ellis and V. Gupta (1998) Static hedging of exotic options. *Journal of Finance*, 53, 1165-1190.

Carr, P. and D. Madan (1999) Option pricing using the fast Fourier transform. *Journal of Computational Finance* 2, 61-73.

Coleman, T, Y. Kim, Y. Li and A. Verma (2001) Dynamic hedging with a deterministic local volatility function model *Journal of Risk*, 4:1, 63-89

Cox, J. (1975) Notes on option pricing I: constant elasticity of variance diffusions. Unpublished work. Graduate School of Business, Standard University.

Crépey, S. (2004) Delta-hedging vega risk? *Quantitative Finance* 4, 559-579.

Derman, E., D. Ergener and I. Kani (1995) Static options replication. *Journal of Derivatives*, 2, 78-95.

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.

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.

Dumas, B., Fleming, F., Whaley, R., (1998). Implied volatility functions: empirical tests. *Journal of Finance* 53:6, 2059-2106.

Fouque, J-P., G. Papanicolaou and K.R. Sircar (2000) *Derivatives in financial markets with stochastic volatility.* 1st edition. Cambridge University Press.

Gatheral, J. (2004) *Case studies in financial modeling course notes.* Courant Institute of Mathematical Sciences. Available from http://math.nyu.edu/fellows_fin_math/gatheral.

Hagan, P.S, D. Kumar, A. S. Lesniewski, and D. E. Woodward (2002) Managing smile risk. *Wilmott Magazine September*, 84-108.

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, 281-300.

Lewis, A. (2000) *Option valuation under stochastic volatility with Mathematica code.* 1st edition. Finance Press.

McIntyre, M., (2001). Performance of Dupire's implied diffusion approach under sparse and incomplete data. *Journal of Computational Finance* 4:4, 33-84.

Merton, R. (1973) Theory of rational option pricing. *The Bell Journal of Economics and Management Science* 4:1, 141-183.

Merton, R. (1976) Option pricing when the underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125-144.

Naik, V. (1993) 'Option valuation and hedging strategies with jumps in the volatility of asset returns', *The Journal of Finance* 48:5, 1969-1984

Rubinstein, M. (1994) Implied binomial trees, *Journal of Finance* 69, 771-818.

Schönbucher, P.J., (1999) A market model of stochastic implied volatility. *Philosophical Transactions of the Royal Society.* Series A, 357, 2071-2092.

Schoutens, W. (2003) *Lévy Processes in Finance.* 1st edition. John Wiley & Sons.

Schroder, M. (1989) Computing the constant elasticity of variance option pricing formula. *Journal of Finance* 44:1, p. 221-219.

Skiadopoulos, G. (2001) Volatility smile consistent option models: a survey. *International Journal of Theoretical and Applied Finance*, 4:3, 2001, 403-438.

Yung, H. H. M and H. Zhang (2003) An empirical investigation of the GARCH option pricing model: hedging performance. *Journal of Futures Markets*, 23:12, 1191-1207.

Figure 1: The Homogeneity Property

The figure depicts the evolution of the asset price in a scale invariant volatility model. The continuous and dotted grey curves indicate confidence limits for the price – these are defined by the price volatility. The price and volatility processes can be correlated and here we assume that volatility increases after a fall in asset price, as shown by the dotted grey line. The figure shows that when the vertical axis is scaled by a positive real number u, with $0 < u < 1$ in this example, the volatility and price-volatility correlation remain unchanged. So if f denotes the price of the option with strike K when the current asset price is S_0 then the price of a vanilla option with strike uK when the current asset price is uS_0 will be uf.

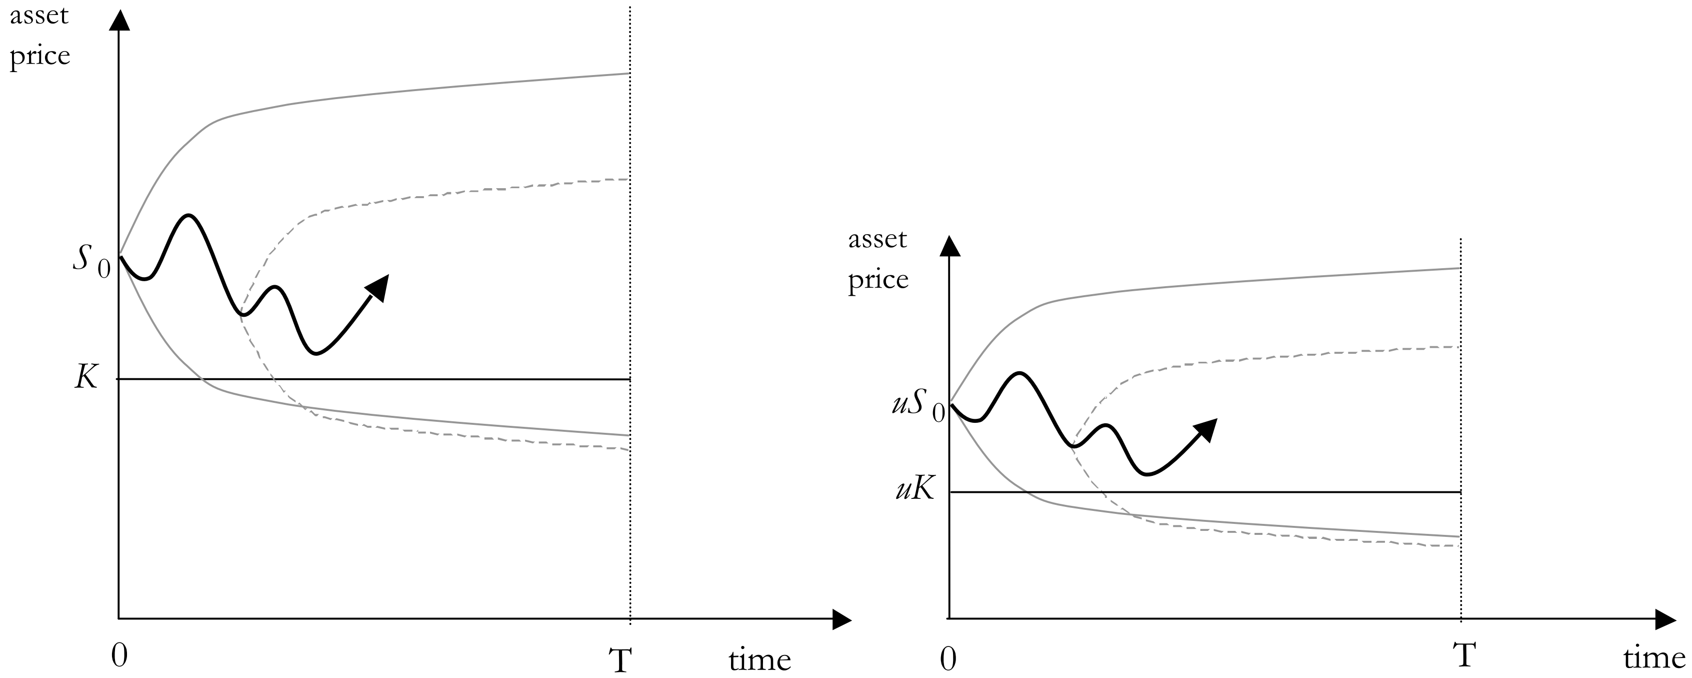

Figure 2: Invariance of Equivalent Local Volatility

Given a stochastic volatility process $\sigma(t)$ the equivalent local volatility is the square root of the conditional expectation of the variance given the asset price at time $t > 0$ is equal to s. That is:

$$\sigma_L^{\ 2}(t, s; S_0) = E_0\left[\sigma^2(t)\middle|S(t) = s\right].$$

The solid black line in figure (a) represents this volatility, along a path leading to s at time t with some individual stochastic volatility paths shown in the dotted vertical plane. The convex shape of the equivalent local volatility at time t, shown by the solid grey line, arises because when the asset price is the ATM forward price the conditional variance is at a minimum. Figure (b) is the translation of figure (a) under a scaling in the price dimension, with $0 < u < 1$ as in figure 1. Figure (c) compares the equivalent local volatility at time $t > 0$ before and after scaling in the price dimension. The equivalent local volatility, shown on the vertical axis, takes identical values before and after the scaling and so the unconditional expected variance at time 0, whose square root is marked by the horizontal line, is constant.

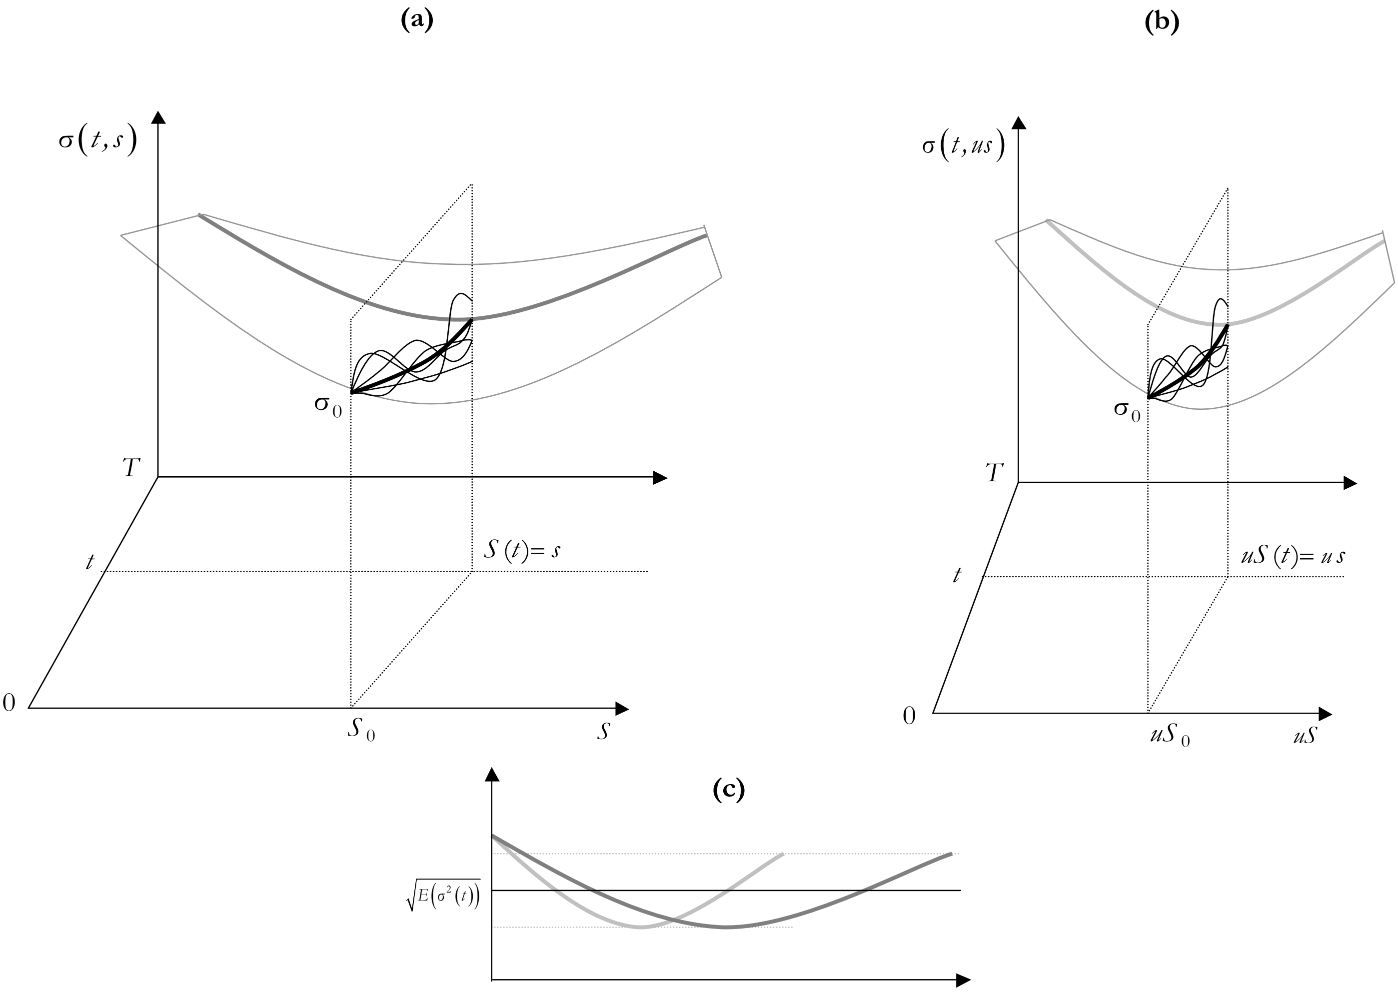

Figure 3: SIV versus MV Hedge Ratios in the Heston Model

The Heston (1993) model is scale invariant and so the (standard) delta is the SIV model-free delta. However, this is not the minimum variance (MV) hedge ratio. Here we use Lewis' (2000) closed-form solution to generate a typical skew, with the Heston model parameters $a = 1$, $b = 0.3$, $\varrho = -0.5$, $V_0 = 0.02$, $m = 0.02$ and assuming risk neutrality. Figure (a) shows the skew for options expiring in 3 months. The strong negative skew is explained by the correlation of -0.5. Figure (b) compares the implied volatility partial and total sensitivities. Note the minus sign in front of the sensitivity to K and that the partial sensitivities are related by Property 6. The partial and total implied volatility sensitivities to S_0 are derived using (29) and (37). These are remarkably different from each other. In fact, the difference, which is fairly constant across strikes and peaks near at-the-money, is zero only if correlation is zero. Figures (c) and (d) compare the standard Heston delta (denoted δ_SIV) and the MV delta (denoted δ_MV) using the Black-Scholes hedge ratios (denoted δ_BS) as the benchmark. We conclude that both the BS and the standard Heston deltas over-hedge in the presence of the skew, their gammas over-hedge for out-of-the-money options and under-hedge for in-the-money options. In addition, standard hedge ratios appear worse than BS hedge ratios in general. The horizontal axis in all charts is moneyness K/S_0 in percent units, i.e. the at-the-money option is at 100.

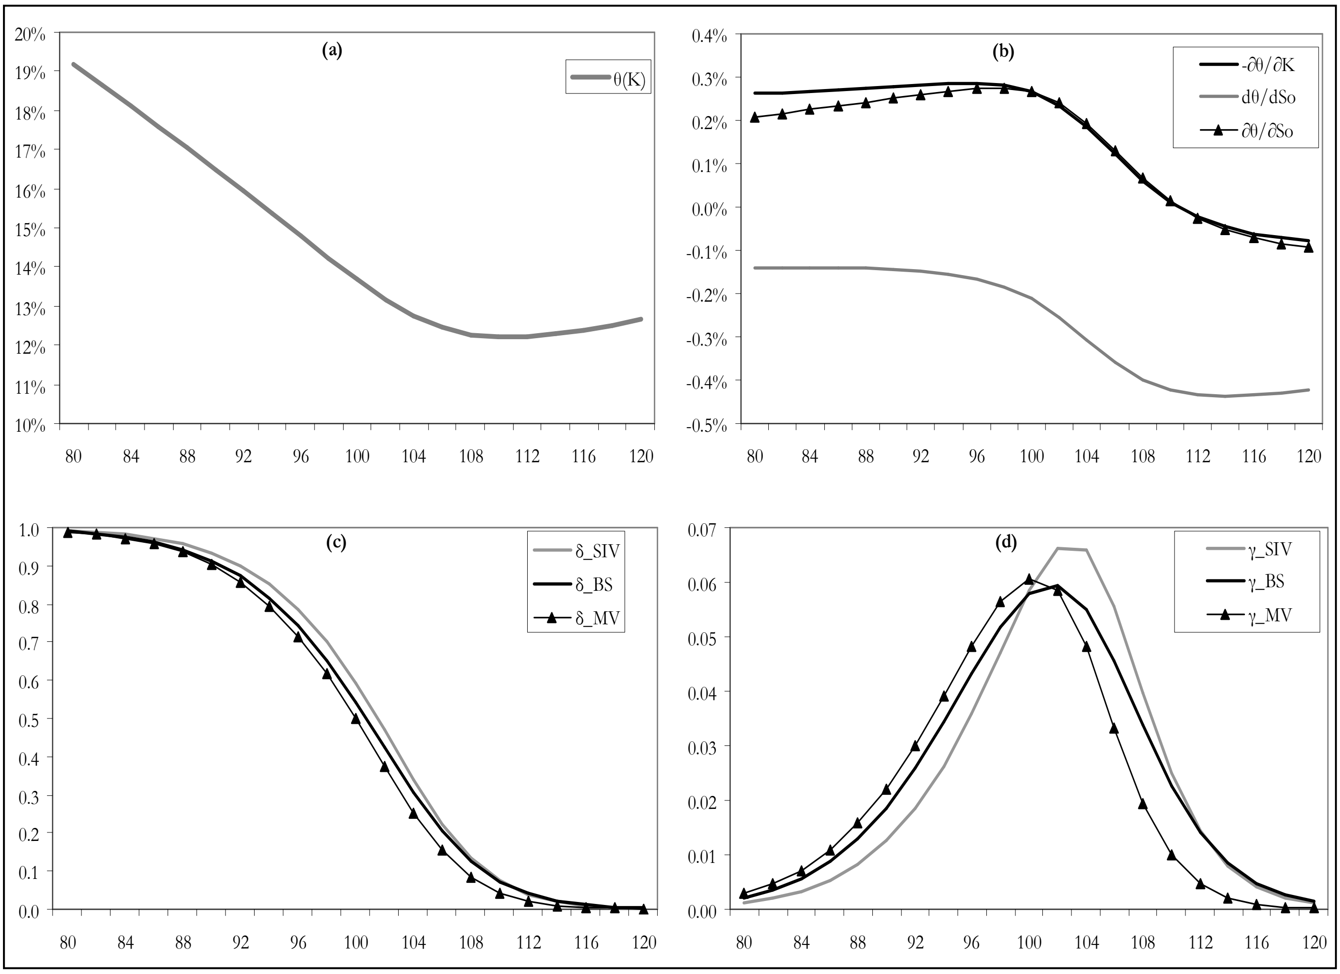

Figure 4: The model's delta and gamma by moneyness on May 21st 2004.

Figure (a) shows the standard and minimum variance (MV) delta of the Heston and SABR model, the SIV 'model free' delta, and the deltas of the CEV and BS models (for which the standard deltas are also MV). Figure (b) shows the corresponding gammas and in each figure they are drawn as a function of K/S_0. May 21st was chosen as a day when all the hedge ratios exhibited their typical pattern.

(a)

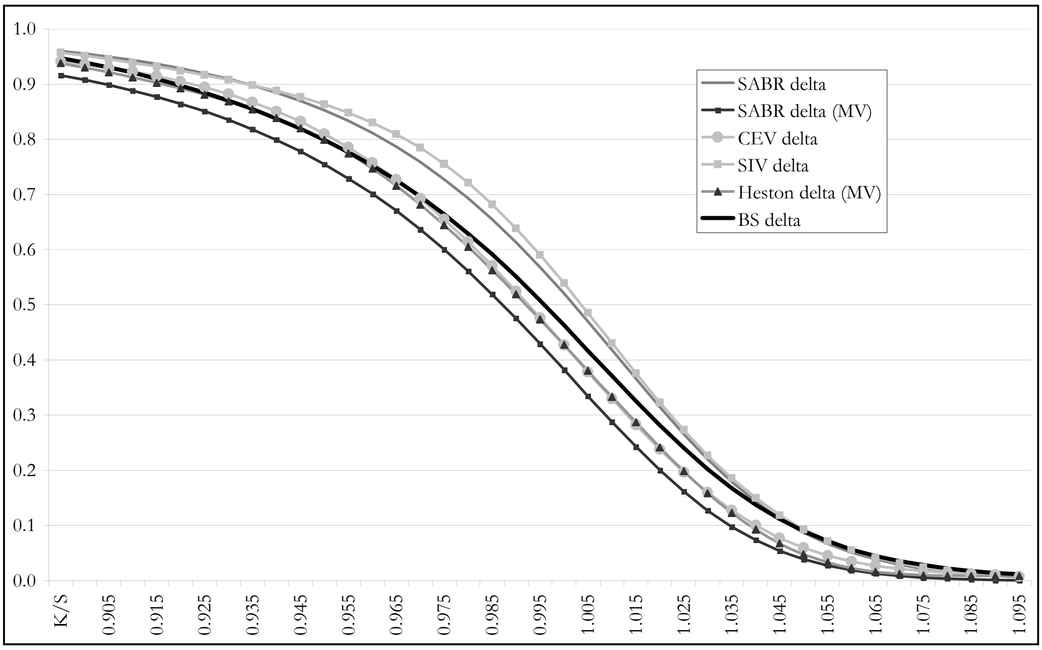

(b)

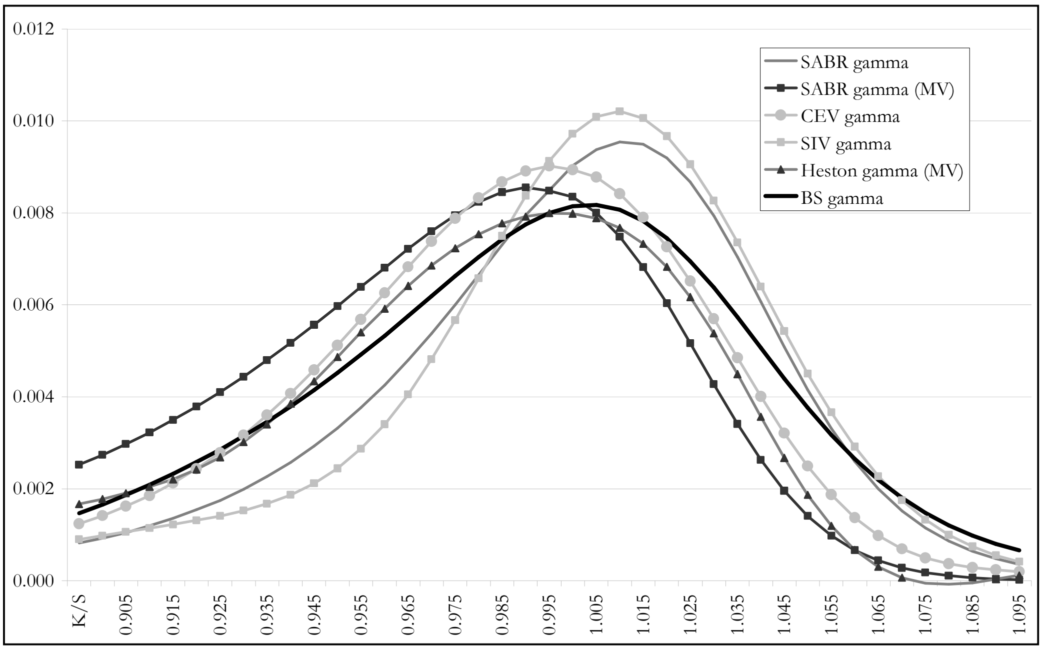

Figure 5: Standard Deviation of Hedging P&L

We consider here only the models that perform well according to our results in Tables 1 and 2, i.e. the CEV model, SABR and Heston models using minimum variance hedge ratios rather than the standard hedge ratios. For each option in the dataset the delta hedged and delta-gamma hedged portfolios are re-balanced daily and the out-of-sample P&L is computed as a daily time series and the standard deviation of each series is calculated. Here we plot these standard deviations as a function of K/S_0 for (a) the delta hedging strategy and (b) the delta-gamma hedging strategy. Whilst the SABR model has noticeably higher standard deviation for OTM call options, statistical tests show no significant difference between the models.

(a) Delta Hedging

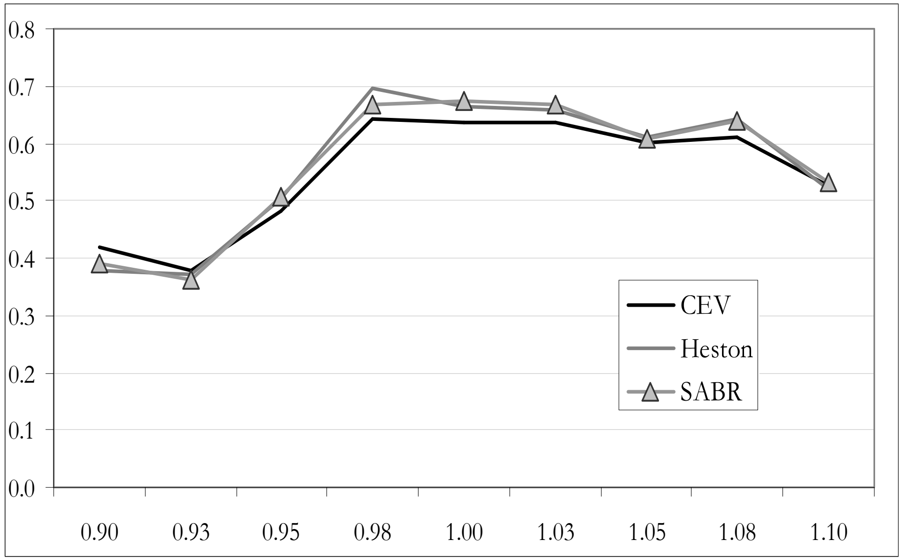

(b) Delta-Gamma hedging

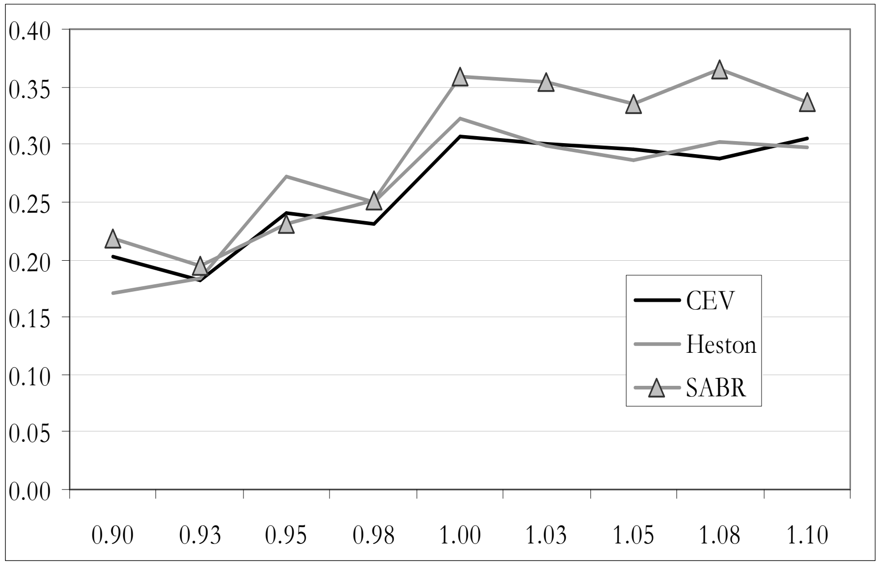

Table I: Sample Statistics of the Aggregate Daily P&L for Delta Hedging

This table reports the sample statistics of the aggregate daily P&L for each model, over all options and over all days in the hedging period, for the delta hedging strategy with daily rebalancing. The models are ordered by the standard deviation of the daily P&L, since to minimize this is the prime objective of dynamic delta hedging. Small skewness and excess kurtosis in the P&L distribution is also desirable – high values for these sample statistics indicate that the model was spectacularly wrong on a few days in the sample. Another important performance criterion is that the P&L be uncorrelated with the underlying asset. In our case over hedging would result in a significant positive correlation between the hedge portfolio and the SP500 index return. We have therefore performed a regression, based on all 1324 P&L data points, where the P&L for each option is explained by a quadratic function of the SP500 returns. The R^2 from this regression, reported in the last column of the table, should zero if the hedge is perfectly effective but will be positive of the model over-hedges.

Model	*Mean*	*Std Dev*	*Skewness*	*Excess Kurtosis*	R^2
CEV	0.1462	0.5847	-0.3424	0.7820	0.113
SABR_MV	0.1218	0.6080	-0.4040	0.8243	0.109
Heston_MV	0.1370	0.6103	-0.5704	1.6737	0.152
BS	0.1401	0.7451	-0.7029	2.0370	0.412
SABR	0.1427	0.9948	-0.6485	1.7099	0.629
SIV	0.1373	1.1788	-0.5928	1.4834	0.693

Table II: Sample Statistics of the Aggregate Daily P&L for Delta-Gamma Hedging

This table reports the sample statistics of the aggregate daily P&L for each model, over all options and over all days in the hedging period, for the delta-gamma hedging strategy with daily rebalancing. The explanation of statistics is given in the legend for Table I. The positive mean P&L in Table I was a result of the short volatility exposure and gamma effects, since we have only rebalanced daily. The delta-gamma hedge strategy results in this table show a mean P&L that is close to zero.

Model	*Mean*	*Std Dev*	*Skewness*	*Excess Kurtosis*	R^2
BS	-0.0014	0.2612	-0.4353	2.5297	0.020
CEV	0.0098	0.2691	-0.0291	3.0850	0.051
Heston_MV	0.0111	0.2789	0.1929	3.6019	0.029
SABR_MV	0.0044	0.3045	-0.3003	3.0032	0.016
SABR	0.0289	0.3821	-0.4845	5.0482	0.057
SIV	0.0428	0.4548	0.0208	4.0123	0.060

Notes:

---

[i] For a review of smile consistent models see e.g. Skiadopoulos (2001) and Gatheral (2004).

[ii] We say that a model is a 'Black-Scholes mixture model' when its price for vanilla options can be written as a linear combination of Black-Scholes prices. This includes, for instance, the class of lognormal mixture models introduced by Brigo and Mercurio (2002).

[iii] Extensions to models with several Wiener or Poisson processes, to non-Markovian models or to processes other than Wiener and Poisson are possible but not discussed here.

[iv] A real-valued function of several real variables $f(\mathbf{x})$ is homogeneous of degree z if and only if

$$f(u\mathbf{x}) = u^{z} f(\mathbf{x}) \quad \forall u \in \mathbb{R}$$

[v] Euler's theorem states that $f(\mathbf{x})$ is a homogeneous function of degree z if and only if

$$\sum_{i=1}^{n} x_i \frac{\partial f(\mathbf{x})}{\partial x_i} = z f(\mathbf{x}).$$

[vi] For this reason some authors have termed scale invariant local models 'floating smile' models. Another terminology that has been used for scale invariant local volatility models is 'sticky-delta', as opposed to the 'sticky-tree' models given by Dupire's original definition of local volatility.

[vii] See e.g. Fouque et al (2000, section 2.8.3) and Lewis (2000, chapter 4).

[viii] Refer to e.g. Schoutens (2003) and Gatheral (2004) on Lévy processes and the Lévy-Khintchine representation.

[ix] See Hagan *et al* (2002) for the proof of this approximation.

[x] For instance, set the CEV and the modified CEV model prices equal:

$$f_{MCEV}(S_0, K, T; \sigma_0, \beta) = f_{CEV}(S_0, K, T; \alpha, \beta)$$

and use the chain rule to derive:

$$\frac{\partial f_{MCEV}}{\partial S_0} = \frac{\partial f_{CEV}}{\partial S_0} + \frac{\partial f_{CEV}}{\partial \alpha}\frac{\partial \alpha}{\partial S_0} = \frac{\partial f_{CEV}}{\partial S_0} - \frac{\partial f_{CEV}}{\partial \alpha}\frac{\alpha\beta}{S_0}$$

[xi] For a discussion on the properties of implied volatility in the Heston model, see Gatheral (2004).

[xii] Note that these hedge ratios require the current spot volatility σ_0 to be a parameter of the local volatility model. If σ_0 is not an explicit parameter, it may be possible to re-parameterize the model in terms of σ_0. For instance, the CEV model defines the spot volatility as $\sigma(t, S) = \alpha S^{\beta}$ so the option price is a function of α and β, but not of σ_0 directly. However, setting $\alpha = \sigma_0 / S_0^{\beta}$ we obtain the modified CEV model where the option price is a function of σ_0. See examples 3 and 4 of Section II.

[xiii] Results for other SIV models are available from the authors by request.

[xiv] Also, no hedging costs have been included in the analysis and these costs should be greater for the over hedging strategies such as BS.

