The University of Reading

Optimal Hedging and Scale Invariance:
A Taxonomy of Option Pricing Models

ICMA Centre Discussion Papers in Finance DP2005-10

First Version: 17th July 2005
This Version: 15th November 2005

Carol Alexander
ICMA Centre, University of Reading

Leonardo Nogueira
ICMA Centre, University of Reading

ICMA 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.icmacentre.rdg.ac.uk
Director: Professor John Board, Chair in Finance
The ICMA Centre is supported by the International Capital Market Association

Abstract

The assumption that the probability distribution of returns is independent of the current level of the asset price is an intuitive property for option pricing models on financial assets. This 'scale invariance' feature is common to the Black-Scholes (1973) model, most stochastic volatility models and most jump-diffusion models. In this paper we extend the scale-invariant property to other models, including some local volatility, Lévy and mixture models, and derive a set of equivalence properties that are useful for classifying their hedging performance. Bates (2005) shows that, if calibrated exactly to the implied volatility smile, scale-invariant models have the same 'model-free' partial price sensitivities for vanilla options. We show that these model-free price hedge ratios are not optimal hedge ratios for many scale-invariant models. We derive optimal hedge ratios for stochastic and local volatility models that have not always been used in the literature. An empirical comparison of well-known models applied to SP 500 index options shows that optimal hedges are similar in all the smile-consistent models considered and they perform better than the Black-Scholes model on average. The partial price sensitivities of scale-invariant models provide the poorest hedges.

Keywords: scale-invariant volatility models, optimal hedging, pricing and hedging of options, minimum variance hedge ratios.

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

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

Acknowledgements: We are most grateful for the guidance of Dr. Bruno Dupire, Bloomberg New York. Very useful comments are also gratefully acknowledged from Dr. Hyungsok Ahn, Nomura, London; Dr. Jacques Pézier, ICMA Centre, UK; Prof. Emanuel Derman, Columbia University, New York; Prof. Rama Cont, ESSEC; Naoufel El-Bachir, ICMA Centre and participants at the Risk Quant 2005 conferences in London and New York and the 2nd Advances in Financial Forecasting symposium in Loutraki, Greece. Any errors are of course our own responsibility.

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.[1] The majority of smile-consistent models in common use can fulfil this aim, some better than others. 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.

As the number of alternative models grows it becomes increasingly important to distinguish between their properties. Generally speaking, volatility dynamics and hedging properties have been studied within a given model class, e.g. stochastic volatility, local volatility or Markovian jump models. In this paper we extend the work of Merton (1973) and Bates (2005) on 'scale-invariant' option pricing models for financial assets to a broader range of models, deriving volatility properties that are equivalent to scale-invariance and focussing on the models' hedging properties. A model is scale-invariant if and only if the density of returns is independent of the current asset price. That is, there is no 'level illusion' in the price – the price dynamics are the same whether the underlying is currently at 100 or 1000. This intuitive property has strong economic justification (see Hoogland and Neumann, 2001), and is an intuitive and desirable property for pricing options on financial assets or indices, but not necessarily for pricing options on economic fundamentals such as interest rates and inflation.

Most asset or index option pricing models are scale-invariant, including the Black-Scholes model, mixture models, stochastic volatility models, as well as some local volatility models and Lévy processes. The necessary and sufficient conditions for a model to be scale-invariant derived below provide a classification of complex models according to scale-invariance, or otherwise. In particular, we prove that the equivalent local volatility surface and the model implied volatility surface are invariant under scaling in the price dimension. For this reason, we use the term 'scale-invariant volatility' (SIV) for all the models that have these properties.

When there is any dependence between the asset price and volatility (or any other parameter of the price process) at time zero in a scale-invariant volatility model then:

$$\left.\frac{\partial f(S)}{\partial S}\right|_{t=0} \neq \frac{\partial f(S_0)}{\partial S_0} \tag{1}$$

---

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

The right-hand side of (1) is the partial derivative of the option price with respect to the current underlying price and it is 'model-free' in SIV models, as shown by Bates (2005). The left-hand side is the sensitivity of the option price to the underlying, evaluated at time 0. It is exact if the underlying is the only source of uncertainty (e.g. in local volatility models) and otherwise (e.g. in stochastic volatility and jump models) it is defined in expectation as:

$$\frac{\partial f(S)}{\partial S} \equiv \frac{\langle df, dS \rangle}{\langle dS, dS \rangle}$$

where $\langle ... \rangle$ denotes the quadratic variation between two random processes. Since the objective of dynamic delta hedging is to minimize the co-variation between changes in the asset price and changes in the value of the hedged portfolio, this 'minimum variance' hedge ratio is the optimal ratio for a pure delta hedge strategy. A similar rule applies to the gamma and to all higher partial price sensitivities.

The minimum variance hedge ratio captures the expected change in volatility, or any other stochastic variable in the price process, when the underlying price moves. If a separate hedge is not used to cover this change (as in Bakshi and Kapadia, 2003) the partial derivative of the option price with respect to the current underlying price will only be optimal when price and volatility are uncorrelated. This distinction is known in the context of jump and stochastic volatility models (see Bakshi, Cao and Chen, 1997) but it is not obvious in local volatility models or in non-standard price processes. Indeed the 'model free' delta on the right-hand side of (1) can lead to serious over-hedging of equity options, as our empirical application shows. We develop a general methodology for obtaining minimum variance hedge ratios, including those for scale-invariant local volatility models, and derive parameter constraints under which scale-invariant models may have model-free minimum variance hedge ratios.

The structure of the paper is as follows: Section I derives the equivalent properties of scale-invariant models and uses these to classify a large number of models according as they are or are not scale-invariant; Section II derives the 'model-free' hedge ratios for scale-invariant models and the optimal hedge ratios for some models; Section III provides an empirical study of the hedging properties of some popular option pricing models; and Section IV summarizes and concludes.

I. Properties of scale-invariant volatility models

We define a scale-invariant volatility (SIV) model by the following process for the underlying asset price:

$$\frac{dS}{S} = G(t, X, ...) \text{ at time } t > 0 \tag{2}$$

where G is a general function, which may typically combine Wiener and Poisson processes, $X = S / S_0$ is the relative asset price at time t, and S_0 is the value of S at time 0. G may depend on a variety of random factors (e.g.

stochastic volatility, stochastic interest rates or jumps) and possibly other variables, except for S or S_0. The only requirement is that S is a proper Itô or Lévy process and that G is at most a function of the relative price X, but not of S or S_0 separately. Note that S does not need to be Markovian.

Property 1: Independence

A model is SIV if and only if 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.

$$\varphi_t(v) = E\left[\exp(iv \ln X)\right] = \int_0^\infty \exp(iv \ln x)\psi_t(x)\,dx \tag{3}$$

at time t *is not a function of* S_0 .

Proof:

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

hence (2) holds iff

$$\frac{dX}{X} = G(t, X, ...) \tag{4}$$

Since $X(0) = 1$, $X(t)$ is independent of S_0 for every $t > 0$. The equivalent rule for the characteristic function follows on differentiating (3) with respect to S_0. A closed-form density $\psi_t(x)$ does not need to be known. If it is not a function of S_0 the model is scale-invariant; and vice-versa. □

Define the price of a standard European option with strike K and time to maturity T as:

$$f(S_0, K, T) = E\Big[B(T)\max\big(w(S(T) - K), 0\big)\Big] = E\Big[B(T)\max\big(w(S_0 X(T) - K), 0\big)\Big]$$

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

Property 2: Homogeneity

A model is SIV if and only if the price of a standard European option is a homogeneous function of degree one with respect to S_0 *and* K*:*

$$f\left(uS_0,uK,T\right)=uf\left(S_0,K,T\right)\quad \forall u>0,T>0 \tag{5}$$

Proof: Using Property 1, we have:

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

which is possible iff $B\left(T\right)$ and $X\left(T\right)$ are not functions of S_0 or K. □

Hence if we scale both the current asset price S_0 and the strike K of a standard European option by the same positive 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 if 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. It is for this reason that the option price must also be scaled by the factor u.

[Figure 1 here]

As shown by Merton (1973), properties 1 and 2 are consistent with no economies of scale with respect to transaction costs and no problem with indivisibilities, such as in penny stocks. They also imply the absence of 'level illusion' as defined in Bates (2005), so that the distribution of asset returns is independent of the level of the asset itself. Additional economic justification is given by Hoogland and Neumann (2001), who also explore the link between scale invariance and the martingale pricing theory.

Property 2 is not limited to standard European options. For instance, the present value of a forward contract expiring at $T > 0$ is:

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

where again we used the fact that $B\left(T\right)$ and $X\left(T\right)$ are independent of S_0. Likewise in (5) we could include American and barrier options and indeed any path-dependent option with characteristics in the price domain beyond a simple strike K, provided that all features of the claim that relate to the price dimension are scaled with the asset price. Yet property 2 cannot be generalized to all claims. For instance, a binary option paying 1 if $S\left(T\right)>K$ at maturity T is worth:

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

Hence the price D of the binary option is a homogeneous function of degree zero.

Property 3: Separability and proportionality

A model is SIV if and only if the first and second order sensitivities of the price of a standard European option with respect to S_0 and K obey:

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

$$\frac{\partial^2 f\left(S_0,K,T\right)}{\partial {S_0}^2}=\left(\frac{K}{S_0}\right)^2\frac{\partial^2 f\left(S_0,K,T\right)}{\partial K^2} \tag{7}$$

Proof: This is proved in Bates (2005): it follows from a straightforward application of Euler's theorem.[2] □

Property 3 states that the option price is an additively separable function of S_0 and K and their first order derivatives and that the second derivatives w.r.t. S_0 and K are proportional. It implies that the first and second order sensitivities to S_0 – often called delta and gamma – are functions of the option price f and its sensitivities to the strike K.

Define the 'equivalent' local volatility as the deterministic function that is consistent with the forward equation 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. In local volatility models the 'equivalent' local volatility is just the local volatility itself. In stochastic volatility models without jumps there is a direct relationship between the equivalent local volatility and the instantaneous volatility $\sigma\left(t\right)$, given by:

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

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

A model is SIV if and only if 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.

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

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

$$\sigma_L\left(t, us; uS_0\right) = \sigma_L\left(t, s; S_0\right) \forall u > 0$$

Proof: The equivalent local volatility defines an equilibrium model for the asset price given by:

$$\frac{dS}{S} = \left(\mu - q\right)dt + \sigma_L\left(t, S; S_0\right)dW$$

This is the minimal model that can fit the observed prices for vanilla options exactly (see Derman and Kani, 1998). Clearly $\sigma_L\left(t, S; S_0\right)$ is a homogenous function of degree zero in S and S_0 iff it is a function of S/S_0 only (and not of these variables separately), and hence also iff the model is scale-invariant, by definition. An alternative verification of property 4 is to use the homogeneity property, writing $f\left(S_0, K, T\right) = Kh\left(m, T\right)$ with $m = \frac{S_0}{K}$ in Dupire's equation to obtain:

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

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 a stochastic volatility model carries over to the scale invariance of equivalent local volatility. The spot volatility σ_0 is the limit of the at-the-money (ATM) local volatility at time 0

$$\sigma_0 = \lim_{t \to 0, s \to S_0} \sigma_L\left(t, s; S_0\right) \tag{10}$$

Defining the spot volatility as a limit prevents one ignoring the possible dependence between price and volatility at time 0. Property 4 implies that the equivalent local volatility surface for any SIV model moves as the underlying asset price changes.[3] On the other hand, in non-scale-invariant models the local volatility $\sigma_L\left(t, s\right)$ is static if and only if it is deterministic and it is not a function of S_0.[4] Non-static local volatility is not a problem for pricing options that can be statically replicated with vanilla options (because calibration guarantees that vanilla options are correctly priced) but it is a problem for hedging as will be shown in Section III.

---

[3] At calibration time t_0 with asset price at S_0 the ATM local volatility is $\sigma_L\left(t_0, S_0; S_0\right) = \sigma_L\left(t_0, 1; 1\right)$ by property 4. At any future time $t > t_0$ and price $S_t = uS_0$, the local volatility is $\sigma_L\left(t, uS_0; S_0\right) = \sigma_L\left(t, u; 1\right)$ again by property 4. Consider calibration at time t, we update time to $t_0^* = t$ and $S_0^* = uS_0$, and require $\sigma_L^*\left(t_0^*, S_0^*; S_0^*\right) = \sigma_L\left(t, uS_0; S_0\right)$ so that the new local volatility is consistent with the previous one. This implies $\sigma_L^*\left(t, 1; 1\right) = \sigma_L\left(t, u; 1\right)$ (again by property 4) so the local volatility surface cannot be static. This proof requires local volatility to be deterministic, but similar result holds when it is stochastic.

[4] The same rationale produces $\sigma_L^*\left(t_0^*, S_0^*\right) = \sigma_L\left(t, uS_0\right)$ which is true for every t and u if and only if $\sigma_L^* \equiv \sigma_L$ (assuming a deterministic local volatility).

[Figure 2 here]

Property 5: Invariance of model implied volatility

A model is SIV if and only if the model implied volatility $\theta(S_0, K, T)$ *is a homogeneous function of degree zero in both K and* S_0*, i.e.*

$$\theta(uS_0, uK, T) = \theta(S_0, K, T) \; \forall u > 0 \tag{11}$$

Proof: The model implied volatility is the Black-Scholes volatility that is implicit in the model price of standard European options. Setting the model option price f equal to the Black-Scholes price f_{BS} with the model implied volatility, and using the fact that the Black-Scholes model is scale-invariant gives:

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

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.

Property 6: **Model-free implied volatility sensitivities**

A model is SIV if and only if 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 zero. □

This property shows (in a non-parametric way) that the sensitivity of the implied volatility to S_0 and the slope of the implied volatility surface in the K axis have the opposite sign. 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 implied volatility surface exactly. In other words these sensitivities are 'model-free' and any observed difference can only be due to a difference in the empirical fit to the smile.

Although some of the properties above are already known for specific models, here they are derived under fewer assumptions and are proven to be very general. Each one of the above properties is a necessary and sufficient condition for a model to be scale-invariant. If any one of the above properties holds, then so do all the other properties and the model is SIV.

The class of SIV models is broad enough to cover most pricing models in the financial literature, and even models not explicitly examined below including uncertain volatility models (such as in Avellaneda et al (1995)) and volatility jump models (Naik (1993)).

Example 1: The Black-Scholes (BS) model

The Black and Scholes (1973) model 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 Itô'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 the BS model is scale-invariant.

Example 2: Mixture models

The term 'mixture model' here refers to any model in which the price of a standard European option is a linear combination of scale-invariant model option prices with weights that are not functions of S_0 or K. For instance, the lognormal mixture diffusion of Brigo and Mercurio (2002) is an example of a Black-Scholes mixture model. Such models are scale-invariant because property 2 holds, as can be seen on writing:

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

where f_{mix} denote the price based on the mixture model and f_{SI}^i denotes the price based on the i^{th} scale-invariant model.

Example 3: The CEV model

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

$$\frac{dS}{S} = rdt + \alpha S^{\beta} dW \tag{13}$$

where $\alpha > 0$ and $\beta < 0$. As S appears in the right-hand side of (13), the CEV model is clearly not scale-invariant.

Example 4: The modified CEV model

Define the modified CEV model as:

$$\frac{dS}{S} = rdt + \sigma_0\left(\frac{S}{S_0}\right)^{\beta} dW \tag{14}$$

where we have introduced a new parameterization of (13) 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 (13), but model (14) is scale-invariant.

Example 5: Local volatility models

Here the asset price dynamics under the risk-neutral measure are:

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

where $\sigma(t,S)$ is a deterministic function of *t*, *S* and possibly other non-random parameters (see Dupire (1994), Derman and Kani (1994) and Rubinstein (1994)). 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.

Overall local volatility models are <u>not</u> scale-invariant. In particular any 'implied tree' model with a static local volatility surface is not scale-invariant. This is because, by property 4, scale invariance only holds when the local volatility is a function of time and S/S_0 and not of S or S_0 separately, and this requires the local volatility surface to move with the underlying.

Scale-invariant local volatilities satisfy:

$$\sigma_L(t,s;S_0) = \sigma_L(t,s/S_0;1)$$

Hence the spot volatility is:

$$\sigma_0 \equiv \sigma(t)\big|_{t=0} = \lim_{t\to 0, s\to S_0} \sigma_L(t,s;S_0) = \sigma_L(0,S_0;S_0) = \sigma_L(0,1;1) \tag{15}$$

but this does not imply that the spot volatility is independent of S at time 0. For instance, in the modified CEV model above:

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

and in general:

$$\frac{\partial\sigma}{\partial S}\bigg|_{t=0} = \lim_{t\to 0, s\to S_0} \frac{\partial\sigma_L(t,s;S_0)}{\partial s} = \lim_{t\to 0, s\to S_0} \frac{\partial\sigma_L(t,s/S_0;1)}{\partial s} = \lim_{t\to 0, x\to 1} \frac{\partial\sigma_L(t,x;1)}{\partial x}\frac{1}{S_0} \neq 0 \tag{16}$$

This price-volatility dependence has important implications for dynamic hedging with scale-invariant local volatility models as we shall see later.

Example 6: Stochastic volatility models

Consider the general stochastic volatility model

$$\begin{aligned} \frac{dS}{S} &= r\,dt + \sigma(Y)\,dW \\ dY &= a(t,Y)\,dt + b(t,Y)\,dZ \quad \langle dW, dZ \rangle = \varrho\, dt \end{aligned} \tag{17}$$

Conditioning on a particular volatility path, it can be shown that the option price may be written:[5]

$$f_{SV}(S_0,K,T) = E\left[f_{BS}\left(S_0 \xi_T, K, T, \sqrt{\bar{\sigma}_T^2} \right) \right] \tag{18}$$

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:

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

so that the homogeneity property holds and the general stochastic volatility model above is scale-invariant.

Example 7: Jump-diffusion models

Consider the jump diffusion model:

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

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

$$f_{JD}(S_0,K,T) = \sum_{n=0}^{\infty} \frac{e^{-\lambda' T}(\lambda' T)^n}{n!} f_{BS}\left(S_0, K, T, r_n, \sqrt{v_n^2} \right) \tag{19}$$

where $\lambda' \equiv \lambda(1+k)$ 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 (19) 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:[6]

---

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

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

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

When the drift ω and the Lévy density $\mu(\varepsilon)$ are not functions of S_0, neither will be $\varphi_t(v)$, and Property 1 implies that such a Lévy process is scale-invariant.

In summary, scale invariance is a very general property that is common to many option pricing 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, and models with stochastic discount rates (unless the rates are a function of the asset price level).

II. Dynamic hedging

Our hedging study focuses on standard 'vanilla' European options only.[7] Property 3 implies that the first partial derivative of a standard European option price $f(S_0, K, T)$ in any SIV model with respect to the underlying asset price is:

$$\frac{\partial f(S_0,K,T)}{\partial S_0} = \frac{1}{S_0}\left(f(S_0,K,T) - K\frac{\partial f(S_0,K,T)}{\partial K}\right) \tag{20}$$

Since the option price and its derivatives with respect to the strike K are observable quantities (assuming a continuum of options) all smile-consistent SIV models should have the same price sensitivity when calibrated to the same options, i.e. the price sensitivities (20) are 'model-free'. This result, which was proved by Bates (2005), motivates the definition of the 'SIV delta', δ_{SIV} as the first partial derivative of the standard European option price with respect to the current underlying asset price:

$$\delta_{SIV}(S_0,K,T) = \frac{\partial f(S_0,K,T)}{\partial S_0} \tag{21}$$

SIV models also have identical higher derivatives.[8] For instance, from (7) the 'model-free' SIV gamma is given by:

$$\gamma_{SIV}(S_0,K,T) = \frac{\partial^2 f(S_0,K,T)}{\partial {S_0}^2} = \left(\frac{K}{S_0}\right)^2 \frac{\partial^2 f(S_0,K,T)}{\partial K^2} \tag{22}$$

---

[7] This can be justified by noting that vanilla calls and puts can be used to replicate statically the payoff of some exotic options (see Derman, Ergener and Kani (1995) and Carr, Ellis and Gupta (1998)). 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.

[8] Using mathematical induction, we show that all derivatives with respect to S_0 are linear functions of f and its derivatives with respect to K. Taking (20) as the base rule and defining $\frac{\partial^0 f}{\partial K^0} \equiv f$ for simplicity, we assume $\frac{\partial^n f}{\partial S_0^n} = \sum_{i=0}^{n} w_i(S_0,K)\frac{\partial^i f}{\partial K^i}$, $n > 0$ so that:

$$\frac{\partial^{n+1} f}{\partial S_0^{n+1}} = \sum_{i=0}^{n}\left(w_i \frac{\partial^i}{\partial K^i}\left(\frac{\partial f}{\partial S_0}\right) + \frac{\partial w_i}{\partial S_0}\frac{\partial^i f}{\partial K^i}\right) = \sum_{i=0}^{n}\left(\left(\frac{w_i}{S_0} + \frac{\partial w_i}{\partial S_0}\right)\frac{\partial^i f}{\partial K^i} - w_i\frac{K}{S_0}\frac{\partial^{i+1} f}{\partial K^{i+1}}\right) = \sum_{i=0}^{n+1} w_i^*(S_0,K)\frac{\partial^i f}{\partial K^i}$$

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

However, (21) and (22) are not the optimal hedge ratios for SIV models, although they are sometimes mistakenly used in practice. For example, Meindl and Primbs (2004) consider the analytical solution for the Heston model using Fourier transforms. Differentiating it with respect to S_0 they produce an estimate for (21), but this is not the correct minimum variance delta for the Heston model, as we show later. Primbs and Yamada (2005) make similar mistake for Merton's (1976) jump diffusion model and conclude that the Black-Scholes model outperforms Merton's model for delta hedging. A more subtle example is the derivation of hedge ratios for the EGARCH option pricing model considered by Yung and Zhang (2003), who also report a poor delta hedging performance for this model. There is no known analytical solution for the option price, so the authors employ Monte Carlo simulation and obtain deltas probably by bumping the current asset price, such as described in James (2003, chapter 10). This is equivalent to using (21). Yet models of the GARCH family are scale-invariant, hence the authors should have used minimum variance deltas (or something similar to the discrete 'realized delta' of Hayashi and Mikland (2005)).

In equity markets the model-free delta (21) will provide a very poor hedge. To see why, note that calibration to market data requires

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

and differentiating with respect to S_0 gives:

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

that is

$$\delta_{SIV}\left(S_0,K,T\right)=\delta_{BS}\left(S_0,K,T,\theta\left(S_0,K,T\right)\right)-v_{BS}\left(S_0,K,T,\theta\left(S_0,K,T\right)\right)\left(\frac{K}{S_0}\right)\frac{\partial \theta\left(S_0,K,T\right)}{\partial K} \tag{23}$$

where δ_{BS} and v_{BS} are the Black-Scholes (BS) model delta and vega respectively and we have used Property 6. Since the global crash of 1987 market implied volatilities have exhibited a pronounced 'skew' in equity markets, that is

$$\frac{\partial \theta\left(S_0,K,T\right)}{\partial K}<0$$

except sometimes for options with very high strikes. Hence the scale-invariant delta will be greater than the BS delta for most options. Since it is well-known that the BS model over-hedges in equity markets (and our empirical results confirm that) the delta hedge from scale-invariant models will perform worse than the BS delta. Nevertheless (23) is popular among FX traders because of its implicit 'sticky-delta' property, where volatility is uncorrelated with the underlying asset, and (23) does give the optimal delta in this case.

The 'minimum variance' (MV) delta, δ is the amount 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,

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

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 smile-consistent scale-invariant models when any model component such as the volatility or interest rates is correlated with the asset price.

Suppose 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 Itô'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 MV delta of a stochastic volatility model is:

$$
\delta_{SV} = \frac{\langle df, dS \rangle}{\langle dS, dS \rangle} = \frac{\left\langle \frac{\partial f}{\partial S} dS + \frac{\partial f}{\partial \sigma} d\sigma, dS \right\rangle}{\langle dS, dS \rangle} = \frac{\partial f}{\partial S} + \frac{\partial f}{\partial \sigma} \frac{\langle d\sigma, dS \rangle}{\langle dS, dS \rangle} \tag{24}
$$

where the first term on the right-hand side is the SIV delta (21) at time 0 and the second term is non-zero if and only if the price-volatility correlation is non-zero. Intuitively, (24) resembles a total derivative of the option price $f(t, S, \sigma)$ with respect to S:

$$
\delta_{SV} = \frac{df(t, S, \sigma)}{dS} = \frac{\partial f}{\partial S} + \frac{\partial f}{\partial \sigma} \frac{d\sigma}{dS} \tag{25}
$$

where the total derivatives are defined in expectation as:

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

In a univariate stochastic volatility model, vega hedging with another option completes the market and we can use the SIV 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), because the term $d\sigma / dS$ in (25) is already captured by the vega hedge. But in the absence of a vega hedge the delta risk arises from the underlying and, if there is price-volatility correlation, from the volatility as well.

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

$$\gamma_{SV} = \frac{d^2 f(t,S,\sigma)}{dS^2} = \frac{d}{dS}\left(\frac{\partial f}{\partial S} + \frac{\partial f}{\partial \sigma}\frac{d\sigma}{dS}\right) = \frac{\partial^2 f}{\partial S^2} + 2\frac{\partial^2 f}{\partial S \partial \sigma}\frac{d\sigma}{dS} + \frac{\partial^2 f}{\partial \sigma^2}\left(\frac{d\sigma}{dS}\right)^2 + \frac{\partial f}{\partial \sigma}\frac{d^2\sigma}{dS^2} \tag{26}$$

where the first term in the right-hand side is the SIV gamma (22) 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 (26) $d^2\sigma / dS^2$ denotes the expected change in $d\sigma / dS$ when S moves. Taking the limit when $t \downarrow 0$ in (25) and (26) 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

In the Heston (1993) model:

$$\begin{aligned} &\frac{dS}{S} = rdt + \sqrt{V}\,dW \\ &dV = a(m-V)\,dt + b\sqrt{V}\,dZ \quad \langle dW, dZ\rangle = \varrho dt \end{aligned} \tag{27}$$

(24) gives

$$\frac{df_{Heston}}{dS} = \frac{\partial f_{Heston}}{\partial S} + \frac{\partial f_{Heston}}{\partial V}\frac{\left\langle b\sqrt{V}dZ, S\sqrt{V}dW\right\rangle}{\left\langle S\sqrt{V}dW, S\sqrt{V}dW\right\rangle} = \frac{\partial f_{Heston}}{\partial S} + \frac{\partial f_{Heston}}{\partial V}\frac{\varrho b}{S} \tag{28}$$

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

$$\delta_{Heston}(S_0, K, T; V_0) = \left.\frac{df_{Heston}}{dS}\right|_{t=0} = \delta_{SIV} + \frac{\partial f_{Heston}}{\partial V_0}\frac{\varrho b}{S_0} \tag{29}$$

where we have used (21) 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 (29) is the second term in the right-hand side.

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

$$\gamma_{Heston}(S_0, K, T; V_0) = \left.\frac{d^2 f_{Heston}}{dS^2}\right|_{t=0} = \frac{\partial^2 f_{Heston}}{\partial S_0^2} + \frac{\varrho b}{S_0}\left(2\frac{\partial^2 f_{Heston}}{\partial S_0 \partial V_0} + \frac{\varrho b}{S_0}\frac{\partial^2 f_{Heston}}{\partial {V_0}^2} - \frac{1}{S_0}\frac{\partial f_{Heston}}{\partial V_0}\right) \tag{30}$$

where $\frac{\partial^2 f_{Heston}}{\partial {S_0}^2} = \gamma_{SIV}$ is the model-free gamma (22).

Figure 3(a) generates an implied volatility skew from the Heston model using typical parameters for equity index options and figure 3(b) depicts the model's total and partial derivatives of implied volatility with respect to S_0 as a function of strike. The partial derivative is very close to the slope of the smile in absolute value, which follows from (12), but the total derivative is much lower at every strike. It even has the opposite sign except for very high strikes, as the smile begins to slope upwards. Figure 3(c) compares the BS delta to the total and partial price sensitivities. The 'model-free' delta is greater than the BS delta but because the chosen correlation ϱ is negative and reasonably large the MV delta is below the BS delta. Figure 3(d) compares the BS gamma, the SIV gamma and MV gamma for the same data as above. Clearly the 'model-free' hedge ratios for equity options can be even more biased than the Black-Scholes deltas and gammas in equity markets.

[Figure 3 here]

Turning now to local volatility models, we find a number of empirical hedging studies on equity indices, including Dumas, Fleming and Whaley (1998), McIntyre (2001), Coleman *et al.* (2001) and Crépey (2004) but none of these use scale-invariant local volatility functions. Their findings are controversial possibly because of the difficulty in calibrating the local volatility function. Dumas, Fleming and Whaley (1998) consider different parametric and semi-parametric forms of local volatility function. They calibrate the models to SP 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 in the FTSE 100 index options, but Coleman *et al.* (2001) find that local volatility deltas for the SP 500 index do improve over long hedging periods. Crépey (2004) examines the local volatility hedging performance 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 these authors use the partial price derivative to obtain delta. Whilst this is not necessarily a problem, because none of the local volatility models considered is scale-invariant, we question the validity of applying these models to equity markets. Since many authors (e.g. Merton (1973), Hoogland and Neumann (2001) and Bates (2005)) argue that scale-invariance is an intuitive property for pricing options on financial assets, the hedging of non-scale-invariant local volatility models might be better tested when an interest rate or another economic fundamental is the option's underlying.

We now derive the MV hedge ratios for local volatility models, scale-invariant and otherwise. 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. Nevertheless, because the spot volatility $\sigma(t,S)$ in a local volatility model is a function of S, it is also a continuous process 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 (24) and (26) the MV local volatility hedge ratios are:

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

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

assuming the spot volatility at time 0 is an explicit parameter of the model.[9]

In non-scale-invariant local volatility models the standard delta and gamma can be equal to the MV delta and gamma. To illustrate this point, consider the CEV model (example 3 above) where:

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

It is simple to verify that the MV delta of the modified CEV model (example 4) is equal to the standard delta of the CEV model, that is:

$$\delta_{MCEV} = \left.\frac{df_{MCEV}}{dS}\right|_{t=0} = \frac{\partial f_{CEV}}{\partial S_0} \tag{31}$$

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.

We now turn to conditions under which two scale-invariant models have identical MV deltas. First suppose that two SIV models fit the smile equally well,[10] so (in the absence of jumps) they have the same equivalent local volatility.[11] By property 4, this may be written

$$\sigma_L\left(t, s; S_0\right) = \sigma_L\left(t, x; 1\right) \text{ with } x = s / S_0 .$$

Now using

---

[9] 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 $\sigma\left(t,S\right) = \alpha S^\beta$ so the option price is a function of α and β, but not of σ_0 directly. However, setting $\alpha = \sigma_0 \big/ {S_0}^\beta$ we obtain the modified CEV model where the option price is a function of σ_0. See examples 3 and 4 of Section II.

[10] Note that it is unlikely to be true that two models fit the smile equally well. For instance, the Heston model (27) has five parameters whilst the modified CEV 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. A negative value for ϱ requires a negative value for β, and this is exactly what is observed when we calibrate these models to S&P 500 index options. But the MCEV cannot separately describe the mean-reversion (captured by *a* and *m* in the Heston model above) and the higher moments (skewness and kurtosis, captured by ϱ and *b*). That is, the modified CEV is under-parameterized and in general will not fit the smile as well as the Heston model.

[11] The result is derived using the equivalent local volatility, which for jump models is a complex issue on its own, so this statement on the equality of minimum variance hedge ratios requires refinements in the presence of jumps.

$$\left.\frac{\partial\sigma}{\partial S}\right|_{t=0} = \lim_{t\to 0, x\to 1}\frac{\partial\sigma_L(t,x;1)}{\partial x}\frac{1}{S_0}$$

as in (16) we have that two SIV models produce the same MV hedge ratios if and only if $\frac{\partial f}{\partial\sigma}$ is the same for both models.

To illustrate this result we derive conditions under which the Heston (1993) stochastic volatility model and the modified CEV local volatility model have identical MV hedge ratios. In the modified CEV:

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

so that:

$$\delta_{MCEV}(S_0,K,T;\sigma_0) = \frac{\partial f_{MCEV}(S_0,K,T;\sigma_0)}{\partial S_0} + \frac{\sigma_0\beta}{S_0}\frac{\partial f_{MCEV}(S_0,K,T;\sigma_0)}{\partial\sigma_0} \tag{32}$$

$$\gamma_{MCEV}(S_0,K,T;\sigma_0) = \frac{\partial^2 f_{MCEV}}{\partial {S_0}^2} + \frac{\sigma_0\beta}{S_0}\left(2\frac{\partial^2 f_{MCEV}}{\partial S_0\partial\sigma_0} + \frac{\sigma_0\beta}{S_0}\frac{\partial f_{MCEV}}{\partial{\sigma_0}^2} + \frac{\beta-1}{S_0}\frac{\partial f_{MCEV}}{\partial\sigma_0}\right) \tag{33}$$

Both models are scale-invariant so:

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

If they fit the smile equally well, i.e. $f_{MCEV}(S_0,K,T;\sigma_0) = f_{Heston}(S_0,K,T;V_0) \quad \forall K,T$ then

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

where V_0 is the spot variance in the Heston model and σ_0 is the spot volatility at time 0. Hence:

$$\delta_{MCEV}(S_0,K,T;\sigma_0) = \delta_{Heston}(S_0,K,T;V_0) \Leftrightarrow \beta = \frac{\varrho b}{2V_0} \tag{34}$$

We conclude this section by deriving the optimal hedge ratios in the SABR model of Hagan *et al.* (2002), a model that has recently become very popular amongst practitioners. The asset price dynamics under the SABR model can be written in the form

$$\begin{aligned} \frac{dS}{S} &= \mu dt + \alpha F^{\beta-1}dW \qquad S(0) = S_0 = F_0 e^{-\mu T} \\ d\alpha &= \upsilon\alpha dZ \qquad \alpha(0) = \alpha_0 \qquad \langle dW, dZ\rangle = \varrho dt \end{aligned} \tag{35}$$

so that the SABR model is not scale-invariant because S appears in the price process through F. The MV delta is given by:

$$\delta_{SABR} = \frac{\partial f}{\partial S} + \frac{\partial f}{\partial \alpha}\frac{\langle d\alpha, dS\rangle}{\langle dS, dS\rangle} = \frac{\partial f}{\partial S} + \frac{e^{(1-\beta)\mu(T-t)}\varrho\upsilon}{S^{\beta}}\frac{\partial f}{\partial \alpha} \tag{36}$$

and letting $t \downarrow 0$ yields:

$$\delta_{SABR} = \frac{\partial f_{SABR}}{\partial S_0} + \frac{e^{(1-\beta)\mu T}\varrho\upsilon}{S_0^{\ \beta}}\frac{\partial f_{SABR}}{\partial \alpha_0} \tag{37}$$

When $\beta = 1$ (37) is consistent with the stochastic volatility delta (24) where α has the same role as σ. Likewise, differentiating (36) with respect to S and letting $t \downarrow 0$ we derive the MV gamma in the SABR model:

$$\gamma_{SABR} = \frac{d^2 f_{SABR}}{dS_0^{\ 2}} = \frac{\partial^2 f_{SABR}}{\partial S_0^{\ 2}} + \frac{e^{(1-\beta)\mu T}\varrho\upsilon}{S_0^{\ \beta}}\left(2\frac{\partial^2 f_{SABR}}{\partial S_0 \partial \alpha_0} + \left(\frac{e^{(1-\beta)\mu T}\varrho\upsilon}{S_0^{\ \beta}}\right)\frac{\partial^2 f_{SABR}}{\partial \alpha_0^{\ 2}} - \frac{\beta}{S_0}\frac{\partial f_{SABR}}{\partial \alpha_0}\right) \tag{38}$$

which is consistent with (26) when $\beta = 1$.

Note that (37) 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. The total derivative:

$$\left.\frac{df_{SABR}}{dF}\right|_{t=0} = \left[\frac{df_{SABR}}{dS}\frac{dS}{dF}\right]_{t=0} = \frac{\partial f_{SABR}}{\partial F_0} + \frac{\varrho\upsilon}{F_0^{\ \beta}}\frac{\partial f_{SABR}}{\partial \alpha_0}$$

captures the correlation between S and α, as in other stochastic volatility models. This example is illustrative since the SABR model is not scale-invariant but it also requires an adjustment to the partial price derivatives to obtain minimum variance hedge ratios, as opposed to the CEV model, in which the MV hedge ratios are the same as the partial price derivatives.

III. Empirical results

Our theoretical results have shown that, under the assumption that the models' fit to option prices is equally good, all scale-invariant models have the same partial price sensitivities (the SIV delta and gamma) and that even their MV hedge ratios can be identical under certain parameter constraints. In practice we generally find that models with fewer parameters do not fit the smile as closely as those with more parameters; but it is not clear whether a very close fit is an advantage for hedging purposes, since we do not want to model noise from the data.

We now compare the hedging performance of some of the models described above: the Black-Scholes (BS) model, the entire class of SIV models using the model-free (and often incorrect) delta (21) and gamma (22), the Heston model with minimum variance delta and gamma given by (29) and (30), the CEV model with minimum variance delta and gamma given by (32) and (33) (note also (31)), the SABR delta and gamma employed by Hagan *et al.* (2002) and, finally, the minimum variance SABR delta and gamma given by (37) and (38). For the SABR hedges we set $\beta = 0$.

We have obtained data from Bloomberg on the June 2004 European options on the SP 500 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.

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 and different means.[12]

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 of SP 500 call options. The SIV 'model-free' delta, i.e. the partial price sensitivity that is common to all SIV models, produces a delta that is greater than the BS delta for all but the very high

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

strikes. This is the same as the delta proposed by Bates (2005). The SABR model delta used by Hagan *et al.* (2002), which is also based on the partial price sensitivity to S_0, lies between the BS and SIV deltas. So if the BS model over-hedges in presence of the skew, then both SIV and SABR deltas should perform worse than the BS model. 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 MV deltas are generally 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 will show. A similar pattern is observed for gammas. SIV and SABR gammas are lower than the BS gamma for ITM calls and greater than the BS gamma for OTM calls (except for very far OTM options) while the opposite is observed when minimum variance gammas are considered. So partial price sensitivities will under-hedge/over-hedge the gamma risk for ITM/OTM calls respectively, relative to the BS hedges.

[Figure 4]

To show that the patterns in figure 4 are not specific for that particular date, we perform a test for the differences between the model hedge ratios over all days in the hedging period. Table I reports the differences averaged by moneyness. They are significant at the 0.1% level except for those marked with an asterisk. The null hypothesis is that the difference between each pair of hedge ratios is zero. For example, the significant and negative value for 'CEV – BS' for every moneyness implies that CEV deltas are typically smaller than BS deltas for SP 500 call options. By contrast, the significant and positive value for 'SIV – BS' in the same table implies that SIV deltas, which are calculated using (21), are greater than BS deltas in general. Overall the differences between deltas are consistent with our theoretical results in Section II: from (23) we expect the SIV delta to be greater than the BS delta in the presence of the skew, and it is indeed substantially greater. Likewise, if the BS delta over-hedges then MV deltas should be smaller than the BS delta to account for this over-hedging, and this is verified for all three MV deltas, which report negative differences when compared to the BS delta. Finally, it is notable that there is almost no significant difference between CEV and Heston deltas. This is remarkable in light of the discussion in Section II, since these models cannot fit the equity skew equally well. The result indicates that the goodness-of-fit to option prices may not be important for delta hedging.

[Table I]

Table I also illustrates the substantial difference between optimal and sub-optimal gammas. Whenever the optimal (minimum variance) gamma is greater than the BS gamma, its sub-optimal counterpart (i.e. the partial price sensitivity) is typically smaller than the BS gamma, and vice-versa. The differences between the MV gammas and the BS gamma also switch from positive to negative close to the ATM option. This implies that the BS gamma may be fairly accurate for near-the-money options.

Tables II and III report the sample statistics of the aggregate daily P&L for each model, over all options and over all days in the hedging period. In both tables the models are ordered by the standard deviation of the daily P&L, since to minimize this is the 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 II and III]

According to these criteria the best delta hedging models are the CEV, SABR(MV) and Heston(MV) deltas, each using the minimum variance hedge ratios. These models outperform the BS model for delta hedging. These also have P&L that is closest to being normally distributed according to the observed skewness and excess kurtosis. Note that both SABR and SIV deltas, which are based on the partial derivative of the option price with respect to S_0, perform worse than the Black-Scholes delta and this confirms that these are indeed the wrong hedge ratios. 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 III show a mean P&L that is close to zero. For delta-gamma hedging it is remarkable that the BS model performs very 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.

Results on hedged portfolio P&L standard deviation by moneyness, averaged over all days in the sample period, are given in Table IV. This table shows that the apparent superiority of the BS model for delta-gamma hedging is only due to its success at hedging the strikes slightly higher than ATM. This may be linked to our finding in Table I that the BS gamma is fairly accurate for near-the-money options. For OTM options the optimal hedge ratios from the Heston model give the lowest standard deviation of hedged portfolio P&L.

Figures 5(a) and (b) illustrate how the delta and delta-gamma hedge performances, as measured by the standard deviation of the hedged portfolio P&L, are distributed across moneyness. Figure 5(a) shows that the three minimum variance delta hedges (i.e., CEV, Heston and the minimum variance SABR hedge) give remarkably similar results for all options, providing the strongest evidence yet that the parameter constraints for equality of MV hedge ratios do hold in practice. In Figure 5(b) we see that the BS model has the smallest standard deviation for the delta-gamma hedge P&L of near-the-money options and that all models deteriorate for out-of-the-money call options, with increasing standard deviation.[13] It is possible that the apparent superiority of the BS model for delta-gamma hedging mid 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 except for low strike ITM call options.[14] In both figures we see that hedging performance is particularly bad when the 'standard' hedge ratios for SABR and SIV models are used.

[Figure 5 here]

VI. Summary and conclusions

Merton (1973) was probably to first to identify that level-independent asset returns lead to the homogeneity of vanilla option prices. More recently Bates (2005) proved that this 'scale invariance' also implies that vanilla option price sensitivities are model-free. Both authors argue that scale invariance is a natural and intuitive property to require for models that price options on financial assets. Yet these authors examined only a limited set of models, applied to only standard European options and they did not explore the full set of properties that are shared by scale-invariant models. Moreover Bates did not consider the optimality of option price sensitivities as hedge ratios.

This paper has (i) widened the class of 'scale-invariant' models beyond standard Markovian stochastic volatility or jump models, to include local volatility, mixture and Levy process models that need not be Markovian; (ii) derived properties, for European and path-dependent option prices and for a model's volatility surfaces, that are equivalent to scale invariance; (iii) shown that scale-invariant 'model-free' hedge ratios can be seriously biased when the asset price and volatility are dependent (e.g. in equity markets); (iv) derived the optimal hedge ratios to be used in delta or delta-gamma hedge strategies, in general and specifically for various well-known models; (v) derived conditions under which minimum variance hedge ratios are also equivalent, for some scale-invariant models; and (vi) performed an empirical study of the hedging performance of popular option pricing models.

The first and second derivatives of the option price with respect to the current spot price give model-free delta and gamma hedge ratios for scale-invariant models, but these are not the optimal hedge ratios for scale-invariant models except when the price and volatility (or any other random factor in the model) are uncorrelated. In dynamic hedging an optimal hedge ratio should minimize the covariance between the changes in the underlying asset and changes in the hedge portfolio value, and such a hedge ratio requires an adjustment to the model-free delta and gamma to account for extra dynamic features such as jumps or any dependence between the underlying

[13] Also, no hedging costs have been included in the analysis and these costs should be greater for over-hedging/under-hedging strategies.

[14] These authors also show that once stochastic volatility is modeled, 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.

price and the other model parameters, including volatility. In the Black-Scholes model the minimum variance hedge is equal to the standard, partial price derivative with respect to the underlying, but this only happens because the Black-Scholes model is not smile-consistent.

The academic literature on hedging with stochastic volatility has, in some cases, used inappropriate price hedge ratios. And to our knowledge all the studies on hedging with local volatility have applied non-scale-invariant models to equity markets. Hence there remains much scope for further empirical research in the light of our theoretical results. For this purpose we have considered models with 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. Their minimum variance hedge ratios are not model-free and they will only be equal under certain parameter constraints, which have been derived above.

Our empirical results on the SP 500 index show that, whilst the model-free hedge ratios of scale-invariant models perform worse than the Black-Scholes model, the optimal hedge ratios provide better hedges on average. Our results also reveal a remarkable similarity in the optimal hedging performance of all the models considered, indicating that the theoretical parameter restrictions for equality of minimum variance hedge ratios do hold in practice, at least for deltas. In the absence of jumps, the optimal delta and gamma hedge ratios differ from the 'model-free' partial price sensitivities because they also account for volatility risk, for instance. Hence there should not be much difference between the volatility hedging of the models considered here, at least according to our empirical results. Clearly, further empirical work should be done in this area. It may be that the most significant difference between the hedging performances of option pricing models lies in their ability to hedge jump risk.

References

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.

Bates, D. (2005) Hedging the smirk. *Financial Research Letters*, *forthcoming.*

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.

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.

Hayashi, T. and P.A. Mykland (2005) Evaluating hedging errors: an asymptotic approach. *Mathematical Finance* 15:2, 309-343.

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.

Hoogland, J.K. and C.D.D. Neumann (2001) Local scale invariance and contingent claim pricing, *International Journal of Theoretical and Applied Finance* 4:1, 1-21.

James, P. (2003) *Option Theory.* 1st edition. Wiley Finance.

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.

Meindl, P.J. and J.A. Primbs (2004) Dynamic hedging with transaction costs using receding horizon control. *Proceedings of the 2nd IASTED international conference*, FEA 2004, Cambridge, USA.

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

Primbs, J.A. and Y. Yamada (2005) A moment computation algorithm for the error in discrete dynamic hedging. *Journal of Banking and Finance, forthcoming*.

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

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, 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.

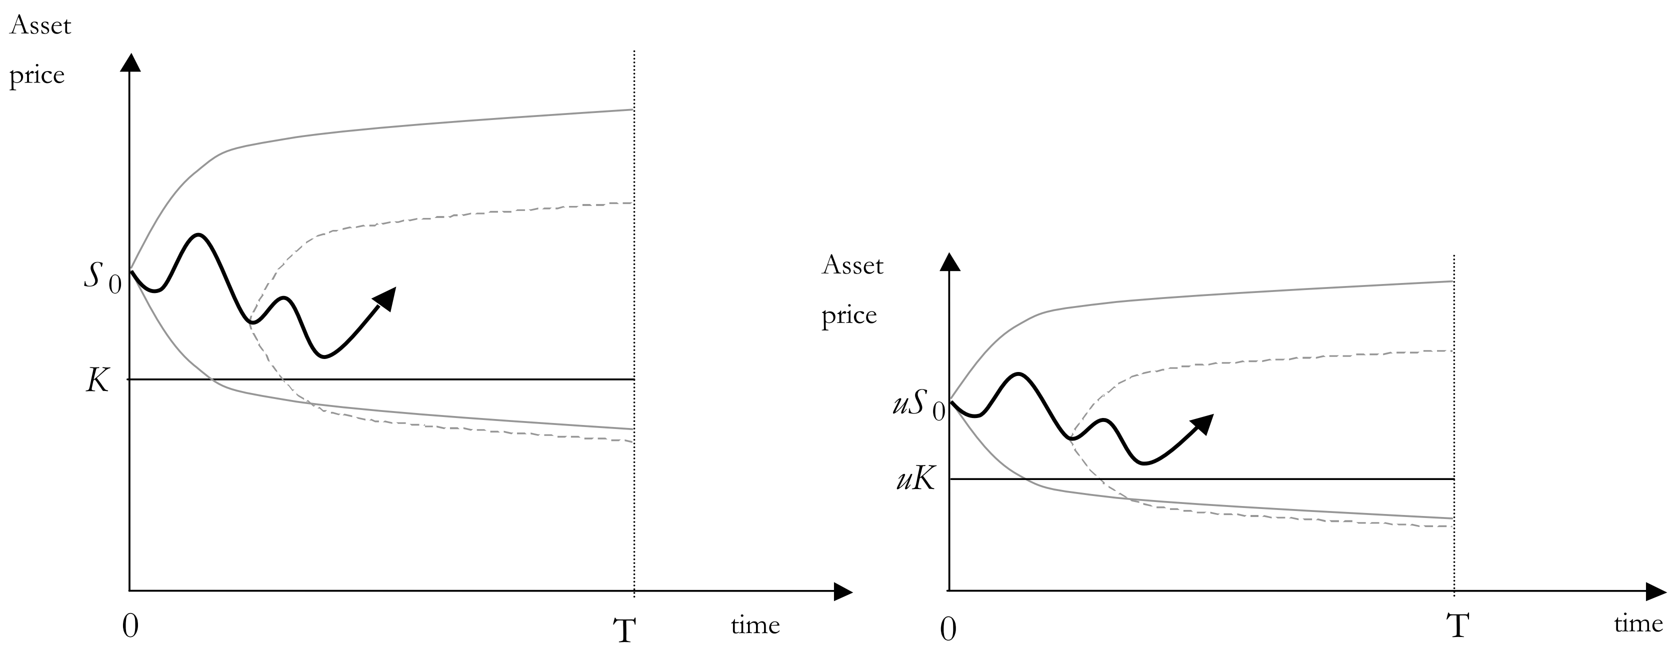

Figure 2: Invariance of Equivalent Local Volatility

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:

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

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.

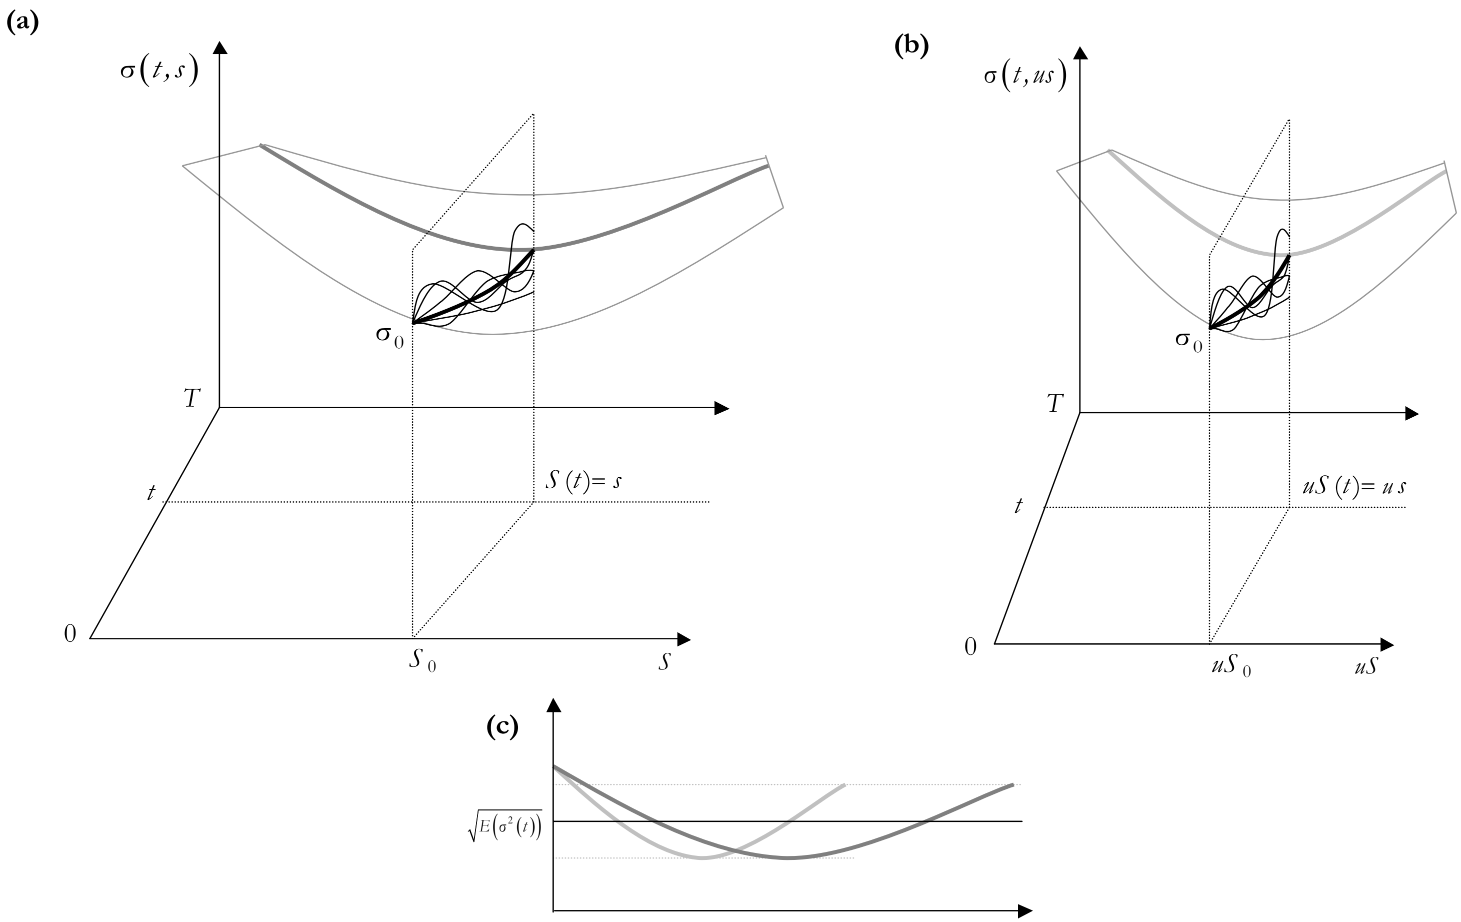

Figure 3: SIV versus MV Hedge Ratios in the Heston Model

The Heston (1993) model is scale-invariant and so the partial price sensitivity 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 in the same way of partial and total price sensitivities. 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 partial 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 partial 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, SIV 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.

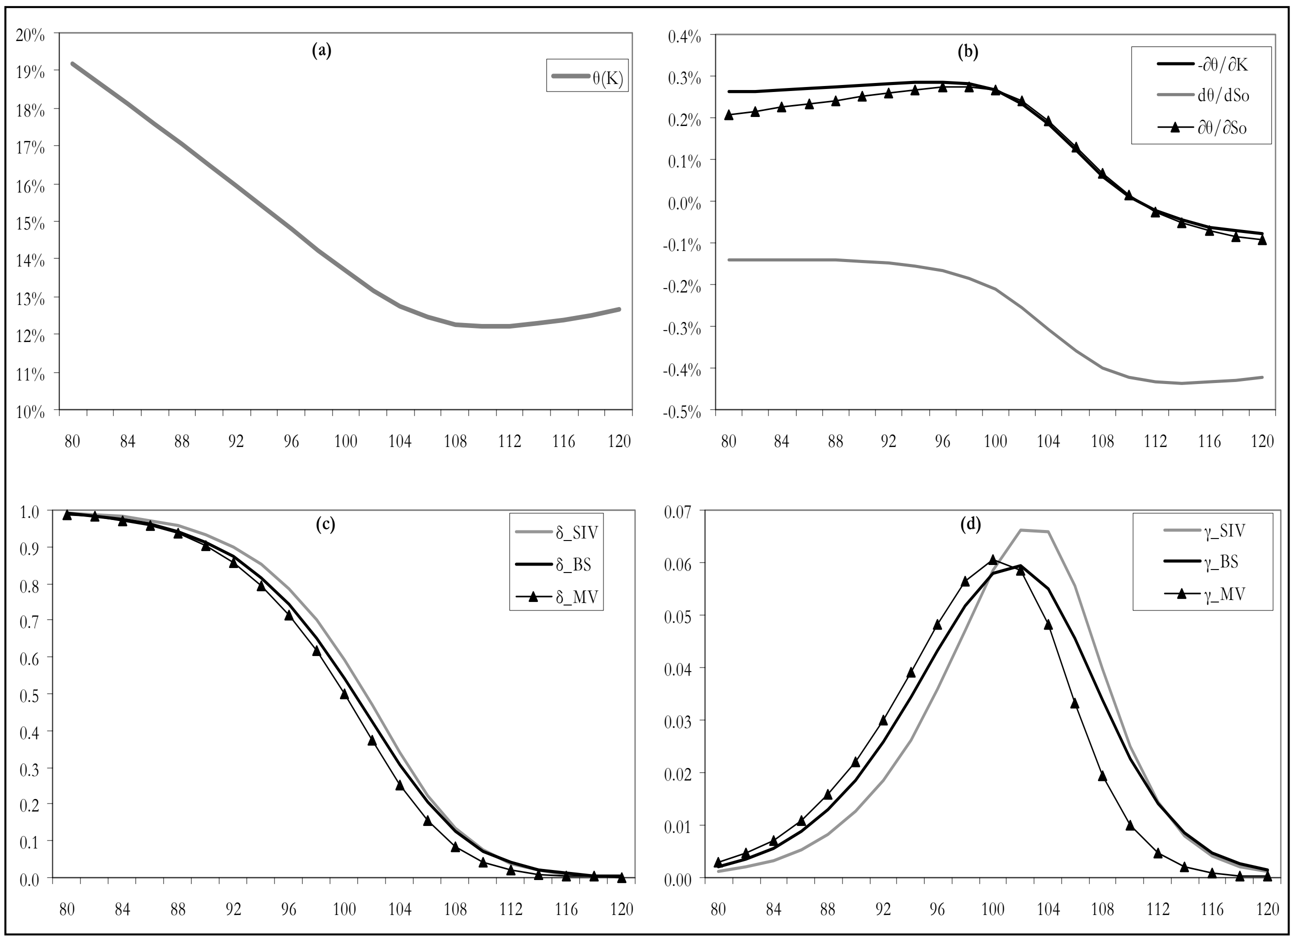

Figure 4: The model's delta and gamma by moneyness on May 21st 2004.

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)

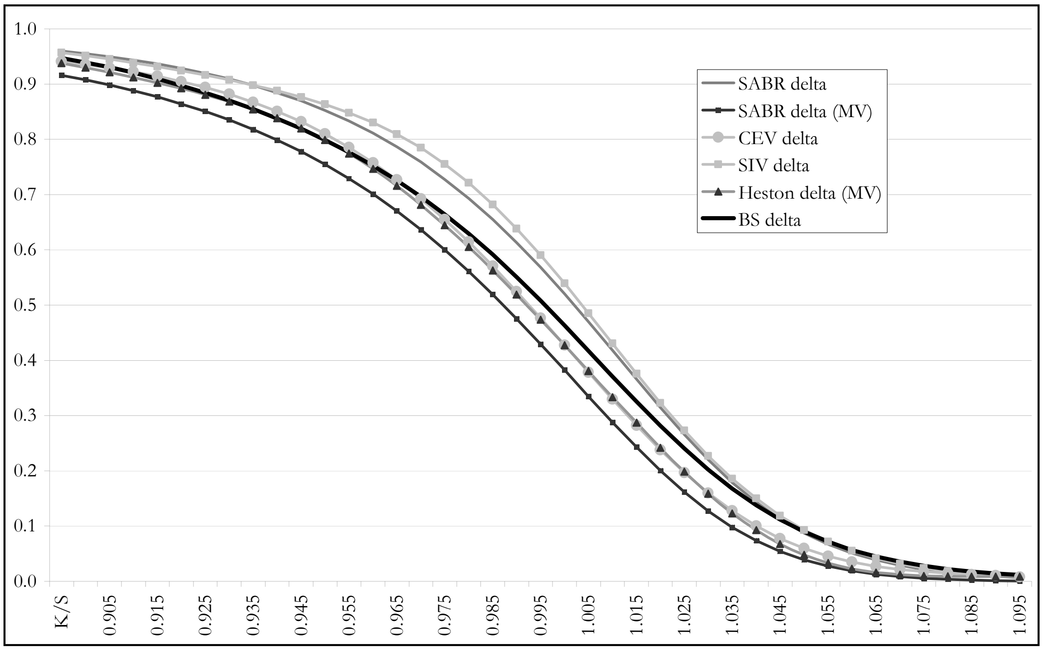

(b)

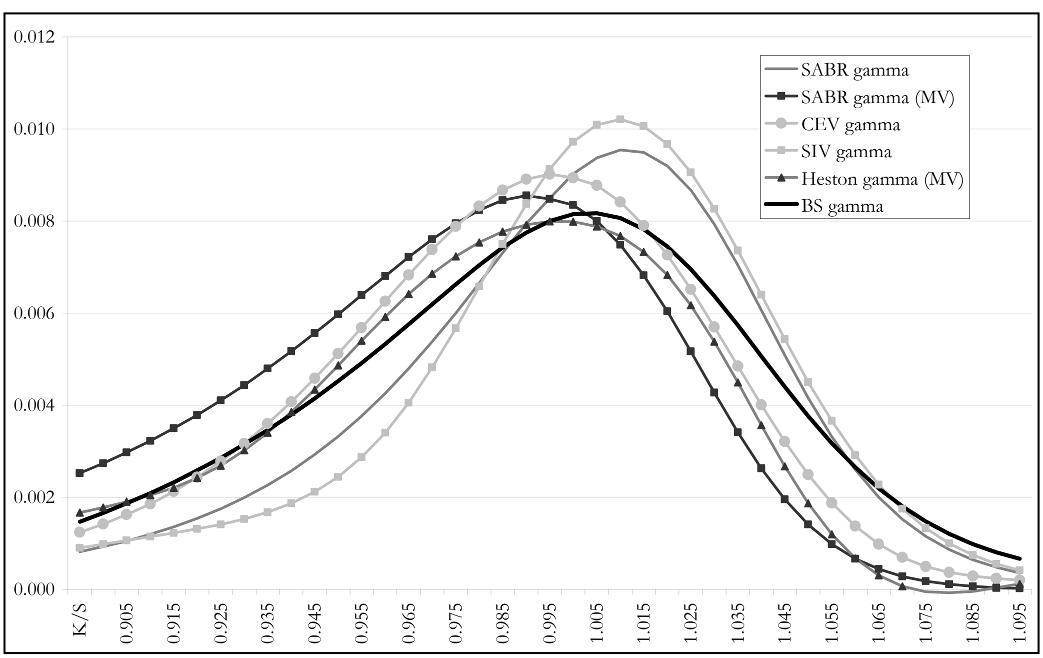

Figure 5: Standard Deviation of Hedging P&L

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 for each model 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 moneyness for the delta hedging strategy (Figure 5-a) and the delta-gamma hedging strategy (Figure 5-b). The SABR model (with $\beta = 0$ in this case) has noticeably high standard deviation for delta-gamma hedging OTM call options even when minimum variance hedge ratios are used.

(a) Delta Hedging

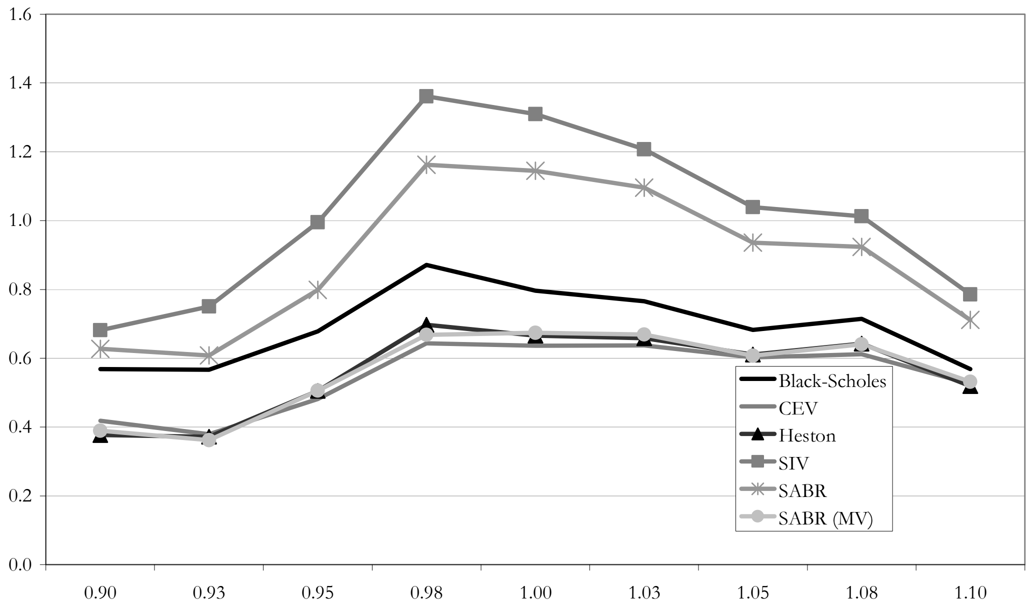

(b) Delta-Gamma Hedging

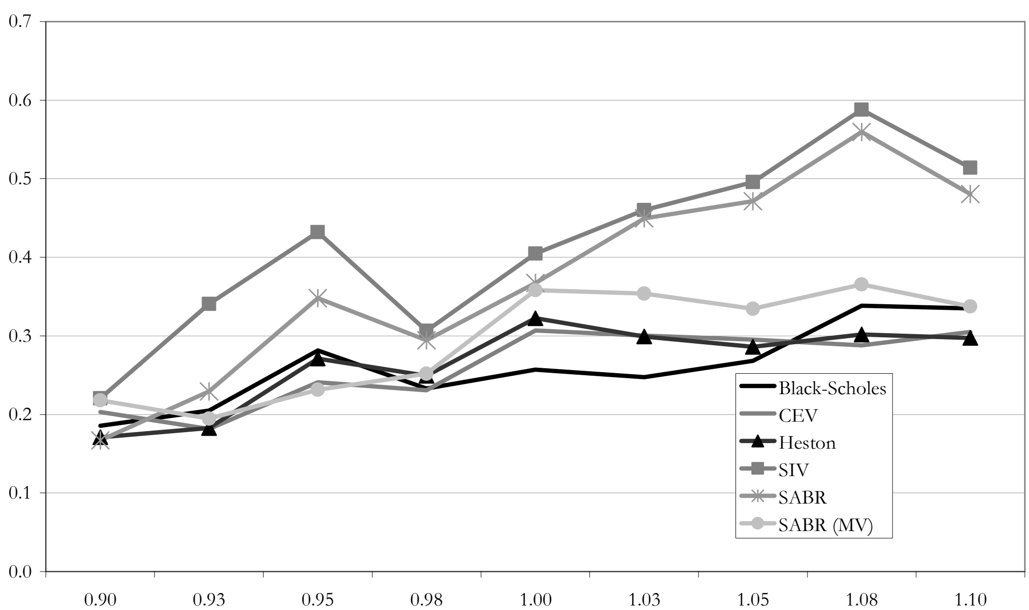

Table I: Sample Statistics for the Differences between Hedge Ratios

This table reports the differences between hedge ratios over all options and over all days in the hedging period, averaged by moneyness. All differences are significant at the 0.1% level, except for those marked with an asterisk. On average, SIV and SABR deltas are greater than the BS delta, while minimum variance deltas are smaller than the BS delta. SIV deltas are the largest while SABR(MV) deltas are the smallest. Differences between CEV and Heston deltas are not significant, and this is roughly independent of moneyness. The minimum variance gammas – SABR(MV), Heston and CEV – are greater/smaller than the BS gamma for ITM/OTM call options, while the sub-optimal SIV and SABR gammas are smaller/greater than the BS gamma for ITM/OTM call options. In effect, the differences between the sub-optimal gammas and the BS gamma have the opposite sign of the same difference for their MV counterparties. Finally, as opposed to the case for the delta, the CEV gamma is on average higher than the Heston gamma in our sample. On average, these statistics concord with Figure 4.

Differences between Deltas

K/S	0.850-0.875	0.875-0.900	0.900-0.925	0.925-0.950	0.950-0.975	0.975-1.000	1.000-1.025	1.025-1.050	1.050-1.075	1.075-1.100
SABR(MV) – BS	-0.022	-0.021	-0.027	-0.030	-0.041	-0.061	-0.085	-0.069	-0.049	-0.032
Heston – BS	-0.010	-0.010	-0.011	-0.017	-0.021	-0.032	-0.046	-0.040	-0.035	-0.027
CEV – BS	-0.013	-0.012	-0.015	-0.019	-0.019	-0.028	-0.048	-0.041	-0.029	-0.019
SABR – BS	0.027	0.028	0.031	0.042	0.054	0.061	0.037	0.012	0.002*	0.000*
SIV – BS	0.028	0.031	0.037	0.055	0.076	0.088	0.060	0.032	0.018	0.008
CEV – Heston	-0.003*	-0.002*	-0.004	-0.002*	0.002*	0.004*	-0.002*	-0.002*	0.006	0.008

Differences between Gammas

K/S	0.850-0.875	0.875-0.900	0.900-0.925	0.925-0.950	0.950-0.975	0.975-1.000	1.000-1.025	1.025-1.050	1.050-1.075	1.075-1.100
SABR(MV) – BS	0.00027	0.00044	0.00071	0.00087	0.00110	0.00139	-0.00059	-0.00168	-0.00156	-0.00118
Heston – BS	0.00021	0.00008*	0.00010*	0.00002*	0.00016	0.00035	-0.00087	-0.00122	-0.00116	-0.00112
CEV – BS	0.00052	0.00049	0.00048	0.00037	0.00017*	0.00093	0.00030	-0.00080	-0.00089	-0.00080
SABR – BS	-0.00053	-0.00058	-0.00070	-0.00089	-0.00107	-0.00017*	0.00133	0.00065	0.00022	0.00002*
SIV – BS	-0.00081	-0.00084	-0.00097	-0.00125	-0.00159	-0.00007*	0.00144	0.00096	0.00062	0.00029
CEV – Heston	0.00032	0.00041	0.00038	0.00035	0.00001*	0.00058	0.00117	0.00042	0.00027	0.00033
# options	11	110	146	184	242	319	301	248	133	36

Table II: 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 III: 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 II. 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

Table IV: Standard Deviation of the Daily P&L Aggregated by Moneyness of Option

This table reports the standard deviation of daily P&L for each model, aggregated over all options of a given moneyness and over all days in the hedging period, for the delta and delta-gamma hedging strategies, with daily rebalancing. According to this criterion the BS model performs best only for the delta-gamma hedging of near ATM options.

Delta Hedging

K/S	0.90-0.95		0.95-1.00		1.00-1.05		1.05-1.10		1.10-1.15	
Best	SABR_MV	0.3657	CEV	0.5740	CEV	0.6372	CEV	0.6051	Heston	0.5507
	Heston	0.3714	SABR_MV	0.5988	Heston	0.6629	SABR_MV	0.6178	CEV	0.5602
	CEV	0.3854	Heston	0.6161	SABR_MV	0.6729	Heston	0.6202	SABR_MV	0.5673
	BS	0.5652	BS	0.7876	BS	0.7844	BS	0.6921	BS	0.5917
	SABR	0.6099	SABR	1.0106	SABR	1.1251	SABR	0.9301	SABR	0.7077
Worst	SIV	0.7357	SIV	1.2055	SIV	1.2691	SIV	1.0283	SIV	0.7746

Delta-Gamma Hedging

K/S	0.90-0.95		0.95-1.00		1.00-1.05		1.05-1.10		1.10-1.15	
Best	Heston	0.1801	CEV	0.2358	BS	0.2531	Heston	0.2907	Heston	0.3134
	CEV	0.1853	SABR_MV	0.2431	CEV	0.3040	CEV	0.2923	CEV	0.3222
	SABR_MV	0.1984	BS	0.2561	Heston	0.3132	BS	0.2929	BS	0.3597
	BS	0.2012	Heston	0.2594	SABR_MV	0.3562	SABR_MV	0.3444	SABR_MV	0.3617
	SABR	0.2187	SABR	0.3202	SABR	0.4015	SABR	0.5022	SABR	0.4871
Worst	SIV	0.3214	SIV	0.3695	SIV	0.4271	SIV	0.5277	SIV	0.5175
# options	141		476		435		217		55	

