



Model-free hedge ratios and scale-invariant models

Carol Alexander ^{a,*}, Leonardo M. Nogueira ^{a,b}

^a ICMA Centre, The University of Reading, School of Business, P.O. Box 242, Reading RG6 6BA, United Kingdom

^b Banco Central do Brasil, SBS Quadra 3 Bloco B, Brasília, 70074-900, Brazil

Received 23 June 2006; accepted 30 November 2006

Available online 25 January 2007

Abstract

A price process is scale-invariant if and only if the returns distribution is independent of the price measurement scale. We show that most stochastic processes used for pricing options on financial assets have this property and that many models not previously recognised as scale-invariant are indeed so. We also prove that price hedge ratios for a wide class of contingent claims under a wide class of pricing models are model-free. In particular, previous results on model-free price hedge ratios of vanilla options based on scale-invariant models are extended to any contingent claim with homogeneous pay-off, including complex, path-dependent options. However, model-free hedge ratios only have the minimum variance property in scale-invariant stochastic volatility models when price–volatility correlation is zero. In other stochastic volatility models and in scale-invariant local volatility models, model-free hedge ratios are not minimum variance ratios and our empirical results demonstrate that they are less efficient than minimum variance hedge ratios.

© 2007 Elsevier B.V. All rights reserved.

JEL classification: G13; C14

Keywords: Scale invariance; Model-free; Hedging; Minimum variance; Stochastic volatility; Local volatility

* Corresponding author. Tel.: +44 (0) 1183 786431; fax: +44 (0) 1189 314741.

E-mail addresses: c.alexander@icmacentre.rdg.ac.uk (C. Alexander), l.nogueira@icmacentre.rdg.ac.uk (L.M. Nogueira).

1. Introduction

Do different option pricing models yield different hedge ratios? This important question is related to model error in option pricing models, an issue that has been addressed by Derman (1996), Green and Figlewski (1999), Cont (2006), Psychoyios and Skiadopoulos (2006) and others. Another challenging question, related to work by Bakshi et al. (1997, 2000) and Lee (2001), is whether minimum variance hedge ratios perform better than standard price hedge ratios in dynamic hedging within a stochastic volatility setting. This paper pursues the answer to these two questions by focussing on scale-invariant models and proving four main results.

A multitude of models for option pricing have been developed in recent years and the academic literature is enormous (see Jackwerth, 1999; Skiadopoulos, 2001; Bates, 2003; Psychoyios et al., 2003; Cont and Tankov, 2004, for comprehensive reviews). However, our first result implies that the vast majority of models share the common property of being scale-invariant. A price process is scale-invariant if and only if the asset price returns distribution is independent of the price measurement scale. The first result allows models to be classified as scale-invariant or otherwise without deriving the returns density. This is important because the returns distribution for many models is not known in analytic form. Thus it broadens the scale-invariant class to encompass models that have not previously been acknowledged as scale-invariant.

Two further results will prove that the price hedge ratios of virtually any claim are model-free, and any difference between the empirically observed hedge ratios can only be attributed to a different quality of the models' fit to market data. More precisely, the standard delta, gamma and higher order price hedge ratios are model-free in the class of scale-invariant models provided only that the claim's expiry pay-off is homogeneous of some degree in the price, strike and any other claim characteristic in the price dimension. Almost all claims in current use have such a homogeneity property.

Vanilla options (i.e. standard European and American calls and puts) have expiry pay-offs that are homogeneous of degree one in the underlying price and strike. Merton (1973) showed that when such options are priced under a scale-invariant process their prices at any time prior to expiry are also homogeneous of degree one. Our second result extends this property to other claims with homogeneous pay-offs: Suppose the expiry pay-off of a claim is homogeneous of degree k in the underlying price, strike and every other parameter in the price dimension (e.g. a barrier). Then, when priced under a scale-invariant process, the price at any time prior to expiry of the claim is also homogeneous of degree k . In other words, the prices of most path-dependent options, such as barriers, Asians, look-backs and forward starts, and the prices of options with pay-offs that are homogeneous of degree $k \neq 1$ such as binary options and power options, at any time prior to expiry, have the same degree of homogeneity as their pay-off functions when they are priced under a scale-invariant process.

Bates (2005) proved that if an option *price* is homogeneous of degree one in the underlying price and strike then its standard delta and gamma are model-free in the class of scale-invariant processes. Our third result extends this model-free property: to options with *pay-off* functions that are homogeneous of *any* degree k in the price dimension, to higher-order price hedge ratios, and to include other characteristics in the price dimension, such as barriers.

Our fourth result is related to minimum variance hedging. The minimum variance (MV) hedge ratio is that ratio which minimizes the variance of the hedged portfolio. See Bakshi

et al. (1997, 2000) for applications. Building on the work of Schweizer (1991), Frey (1997) and Lee (2001) we derive explicit expressions for the minimum variance delta and gamma of some standard option pricing models. We then show that model-free scale-invariant hedge ratios only have the MV property when there is zero correlation between the price and another (possibly stochastic) component in the model such as volatility or interest rates. Otherwise, to be MV the standard hedge ratio requires a simple adjustment that depends on this correlation.

An empirical study of standard European options on the S&P 500 index indicates that extending the definition of delta and gamma from simple partial derivatives to the MV hedge ratios mentioned above yields a major improvement in dynamic hedging performance. However, we find no significant difference between the performances of different MV hedges. Finally, our results are not conclusive about the superiority of MV hedge ratios over the Black–Scholes (1973) delta–gamma hedge.

The rest of this paper is structured as follows. Section 2 proves our first two results on the classification of scale-invariant option pricing models and the preservation of homogeneity by scale-invariant processes; Section 3 proves our third result on the model-free delta and gamma of European and American claims, knowing only that their expiry pay-off is homogeneous of degree k in the price dimension, and derives expressions for the minimum variance delta and gamma of some option pricing models; Section 4 presents the results of the empirical study and Section 5 concludes.

2. Scale-invariant models and their properties

Let S_t denote the price at time t of the contract underlying a contingent claim and denote the relative price at time t by $X_t = S_t/S_0$. A price process $S = (S_t)_{t \geq 0}$ is defined as *scale-invariant* if and only if the marginal distribution of X_t is independent of S_0 for all $t \geq 0$ that is

$$\frac{\partial \pi_t(x)}{\partial S_0} = 0 \quad \text{for all } t \geq 0, \quad (1)$$

where $\pi_t(x) = \frac{d}{dx} P(X_t < x)$ is the probability density of X_t .

In other words, S is scale-invariant if and only if the returns density is independent of the price dimension. Merton (1973) identified this ‘constant returns to scale’ property as a desirable feature for pricing options. He also showed that if the probability density of the underlying asset returns is invariant under scaling then the price of a standard American or European option scales with the underlying price. Put another way, it does not matter whether the asset price is measured in dollars or in cents – the relative value of an option should remain the same.¹

¹ Hoogland and Neumann (2001) consider scale invariance as a parallel to a change of numeraire, but we regard scale invariance as the invariance of the returns density under a change in the unit of measurement of the underlying price. This is not the same as a change of numeraire. The price of every asset in the economy changes if we change the numeraire, whilst for our purposes scale invariance refers only to a change in the unit for measuring the underlying price and everything else that is in the same dimension as this price, such as an option strike or barrier. A simple example is a stock split. After the split, the value of the stock and the strike price of any option on this stock will be scaled, but the prices of the remaining assets in the economy (e.g. bonds) are not changed. For a review of this and other general properties of option prices, including bounds for hedge ratios, refer to Merton (1973), Cox and Ross (1976), Bergman et al. (1996) and Bakshi and Madan (2002).

The theorem below shows that almost all of the models for pricing options on financial assets that are in common use are scale-invariant, however most interest rate models are not scale-invariant. The proof of this and other theoretical results are given in the [Appendix](#).

Theorem 1. *A price process $S = (S_t)_{t \geq 0}$ is scale-invariant if it is a semi-martingale and can be written in the form*

$$\frac{dS}{S} = \Theta' d\Lambda, \quad (2)$$

where Θ is a vector of random or deterministic coefficients that are independent of the unit of price measurement and $\Lambda = (\Lambda_t)_{t \geq 0}$ is a vector of factors driving the asset price that contains the time t , Wiener processes and/or jump processes.

It is easy to see that several classes of models are scale-invariant because they satisfy the general form (2). The price process does not even need to be Markovian. [Bates \(2005\)](#) observes that [Merton's \(1976\)](#) jump-diffusion and most stochastic volatility models, even with stochastic interest rates, are scale-invariant. [Theorem 1](#) allows further models to be classified as scale-invariant including: mixture diffusions (such as in [Brigo and Mercurio, 2002](#)), uncertain volatility models ([Avellaneda et al., 1995](#)), double jump models ([Naik, 1993](#); [Duffie et al., 2000](#); [Eraker, 2004](#)), and Lévy processes ([Schoutens, 2003](#)) if the drift and Lévy density are dimensionless. Option pricing models that are not scale-invariant include: models based on arithmetic Brownian motion and its extensions (e.g. [Cox et al., 1985](#)) that are commonly applied to interest rate options; deterministic volatility models in which the instantaneous volatility is a static function of the asset price (e.g. [Cox, 1975](#); [Dumas et al., 1998](#)); and the 'implied tree' local volatility models of [Dupire \(1994\)](#), [Derman and Kani \(1994\)](#) and [Rubinstein \(1994\)](#) where the diffusion coefficient in (2) depends on the price level. Hybrid models that mix local volatility with stochastic volatility or jumps (e.g. [Hagan et al., 2002](#); [Carr et al., 2004](#)) are typically not scale-invariant because of the local volatility component.

Now consider an arbitrary claim on S with expiry T and characteristics $\mathbf{K}' = (K_1, \dots, K_n)$ in the same unit of measurement of S , such as strikes and barriers. The claim may itself be a portfolio of other claims on S , e.g. a straddle, butterfly spread, etc. Without loss of generality we assume the claim characteristics are known constants and we omit variables such as interest rates, dividends and other model parameters because these are of lesser importance for price hedging. We therefore denote its price at time t , with $0 \leq t \leq T$, by $g(T, \mathbf{K}; t, S)$.

Theorem 2. *A price process is scale-invariant if and only if it preserves the homogeneity of a claim pay-off at expiry throughout the life of the claim.*

Many types of options have homogeneous pay-off functions. Pay-offs that are homogeneous of degree zero include the log-contract and binary options.² Power options are homogeneous of degree $k > 1$.³ But most claims have pay-off functions that are homo-

² The log contract pays $\ln(S_T/S_0)$ at expiry and a binary option pays $1_{\{S_T > K\}}$ for a call or $1_{\{K > S_T\}}$ for a put.

³ For instance those with pay-off $[(S_T - K)^k]^+$.

geneous of degree one, including standard options, cash-or-nothing and asset-or-nothing options and many path-dependent options such as look-backs, single and multiple barrier options, average-rate and average-strike options, forward start and cliquet options and compound options.⁴ Theorem 2 shows that when a scale-invariant process is used to value any of the claims mentioned above, the claim price at any point in time before expiry will be homogeneous and has the same degree of homogeneity as its payoff.

To classify option pricing models as scale-invariant or otherwise when neither the returns density nor the price process are known in analytic form, it is useful to consider the following corollary to Theorem 2. Let $\theta(T, K; t, S)$ denote the implied volatility of a standard European option with maturity T and strike K , and $\hat{\sigma}(T, K; t, S)$ denote the local volatility for future time T and price K , both seen from time t when $S_t = S$.

Corollary 1. *The following properties are equivalent for all T and K :*

- (i) S is generated by a scale-invariant process;
- (ii) $\theta(T, K; t, S) = \theta(T, uK; t, uS) \quad u \in \mathbb{R}^+$;
- (iii) $\hat{\sigma}(T, K; t, S) = \hat{\sigma}(T, uK; t, uS) \quad u \in \mathbb{R}^+$.

The corollary shows that a model is scale-invariant if and only if the implied volatility and the local volatility are both homogeneous functions of degree zero in S and K , a property that is sometimes called the ‘floating-smile’. Note that *all* scale-invariant models share these volatility characteristics, not just scale-invariant local volatility models. Applying Euler’s theorem to property (ii) shows that *every* scale-invariant model has the same implied volatility sensitivity to S

$$\theta_S(T, K; t, S) = -\left(\frac{K}{S}\right)\theta_K(T, K; t, S), \quad (3)$$

where θ_K is the slope of the implied volatility smile in the strike metric. We note that Bates (2005) also derived the identity (3).

3. Hedging with scale-invariant models

Bates (2005) showed that if at some time t , $0 \leq t \leq T$, the price of an option is homogeneous of degree one in S and K , then every scale-invariant process gives the same option delta and gamma at time t . The theorem in this section extends and generalises Bates’ result by showing that all price sensitivities of an arbitrary claim are model-free within

⁴ Pay-offs are defined as follows: Standard options: e.g. a vanilla call pays $(S_T - K)^+$; Cash-or-nothing options: $K1_{\{S_T > K\}}$ for a call; Asset-or-nothing options: $S_T 1_{\{S_T > K\}}$ for a call; Look-back options: e.g. $(S_T - S_{\min})^+$; Barrier options: e.g. $(S_T - K)^+ 1_{\{S_t < B, 0 \leq t \leq T\}}$ is a single barrier up-and-out call. Multiple barrier options are also homogeneous of degree one; Asian options: e.g. $(A_T - K)^+$ where A_T is an average of prices prior to and at expiry; Compound options: e.g. $(C(T_1, T_2) - K)^+$ where $C(T_1, T_2)$ is the value of a vanilla call at time T_1 with expiry date $T_2 > T_1$; Forward start options: e.g. $(S_{T_2} - S_{T_1})^+$, where the strike is set as the at-the-money strike at $T_1 < T_2$. Cliquet options are a series of forward start options and are therefore also homogeneous of degree one.

the class of scale-invariant process, provided only that the claim pay-off at expiry is homogeneous of degree k in the price dimension.⁵

3.1. The model-free property

The next theorem implies that if prices of claims of the same type are observable in the market, then so are the price hedge ratios of these claims. More precisely, any two scale-invariant models yield the same price hedge ratios for a claim with homogeneous pay-off and characteristics \mathbf{K} if the same claim prices are used to calibrate the models and if both models fit these prices exactly. A perfect fit to market prices is not always attainable in practice, but if two scale-invariant models fit the data reasonably well then no significant difference between the empirical hedging performances of the models should be observed.

Theorem 3. *Suppose the claim pay-off is homogeneous of degree k and that S is generated by a scale-invariant process. Then all partial derivatives of the claim price with respect to S at any time $t < T$ are given by linear combinations of $g = g(T, \mathbf{K}; t, S)$ and its partial derivatives with respect to \mathbf{K} , and in particular*

$$\begin{aligned} g_S &= S^{-1}(kg - \mathbf{K}'g_{\mathbf{K}}) \\ g_{SS} &= S^{-2}[\mathbf{K}'g_{\mathbf{K}\mathbf{K}}\mathbf{K} + (k - 1)(kg - 2\mathbf{K}'g_{\mathbf{K}})]. \end{aligned} \quad (4)$$

Applying the theorem to standard European options: if two scale-invariant models are calibrated to the same market prices, both models should give the same delta and the same gamma for the options, because the price sensitivities to \mathbf{K} can be computed directly from the market prices. The theorem also applies to path-dependent options: for instance, if two scale-invariant models are calibrated to the same market prices of barrier options both models should give the same price hedge ratios for these barrier options. Empirically, there will be differences between the delta and gamma obtained using the two models but this is due to the fact that the models do not fit market data equally well. On the other hand, if a model is calibrated to standard European calls and puts and then used to price and hedge path-dependent options such as barrier or cliquet options, the price and the price sensitivities of the path-dependent options will be model-dependent. When the prices of the path-dependent options are not observable in the market and are given by the model, both price and price hedge ratios are model-dependent.

3.2. Minimum variance hedge ratios

So far we have defined delta and gamma as the usual partial derivatives of the claim price with respect to the underlying price. However, when there are extra dynamic features in the model such as stochastic volatility or stochastic interest rates, these might not be the most efficient hedge ratios to use in a delta or delta–gamma hedging strategy. This sub-section

⁵ In Theorem 3 all claim prices and derivatives of these prices are functions of $(T, \mathbf{K}; t, S)$ but we have dropped this dependence for ease of notation. Also $(g_{\mathbf{K}})_{n \times 1}$ is the gradient vector of partial derivatives and $(g_{\mathbf{K}\mathbf{K}})_{n \times n}$ is the Hessian matrix of second partial derivatives of g with respect to \mathbf{K} , all evaluated at time t when $S = S_t$. Finally \mathbf{K}' denotes the transpose of \mathbf{K} .

investigates when the model-free hedge ratios of scale-invariant models given by **Theorem 3** and the minimum variance hedge ratios coincide.

We define the minimum variance (MV) delta, δ_{mv} , as the amount of the underlying asset at time t that minimizes the instantaneous variance of a delta-hedged portfolio, $\Pi = g - \delta_{mv} S$, or equivalently, that reduces the instantaneous covariance of the portfolio with the underlying asset price S to zero⁶

$$\langle d\Pi, dS \rangle = \langle dg - \delta_{mv} dS, dS \rangle = \langle dg, dS \rangle - \delta_{mv} \langle dS, dS \rangle = 0, \tag{5}$$

where $\langle \cdot; \cdot \rangle$ denotes the instantaneous covariance between two random variables. As before, we drop the dependence of Π , g and δ_{mv} on $(T, \mathbf{K}; t, S)$ for ease of notation.

In the Black–Scholes model, the MV delta is the same as the first partial derivative of the claim price with respect to S , but this is not the case when any model component such as the volatility or interest rate is correlated with the asset price. Suppose the spot volatility (or variance) is a continuous and stochastic process itself and there are no jumps. Then the dynamics of the claim price $g = g^{sv}(T, \mathbf{K}; t, S, \sigma)$ according to the stochastic volatility (SV) model are given by Itô’s formula as

$$dg = g_t dt + g_S dS + g_\sigma d\sigma + \frac{1}{2}g_{SS} dS^2 + \frac{1}{2}g_{\sigma\sigma} d\sigma^2 + g_{S\sigma} dS d\sigma, \tag{6}$$

where the subscripts of g denote partial differentiation. In a stochastic volatility model without jumps the MV delta, δ_{mv}^{sv} , is the ratio of the instantaneous covariance between increments in the claim price and the underlying price and the instantaneous variance of the increments in the underlying price. Therefore, since the quadratic terms in (6) are adapted processes of order dt ,

$$\delta_{mv}^{sv}(T, \mathbf{K}; t, S, \sigma) = \frac{\langle dg, dS \rangle}{\langle dS, dS \rangle} = \frac{\langle g_S dS + g_\sigma d\sigma, dS \rangle}{\langle dS, dS \rangle} = g_S + g_\sigma \frac{\langle d\sigma, dS \rangle}{\langle dS, dS \rangle}. \tag{7}$$

Intuitively, this resembles a total derivative of the claim price with respect to S , in which the total derivatives are defined as

$$\frac{dg}{dS} \equiv \frac{\langle dg, dS \rangle}{\langle dS, dS \rangle} \quad \text{and} \quad \frac{d\sigma}{dS} \equiv \frac{\langle d\sigma, dS \rangle}{\langle dS, dS \rangle} \tag{8}$$

and

$$\delta_{mv}^{sv} = \frac{dg}{dS} = g_S + g_\sigma \frac{d\sigma}{dS}. \tag{9}$$

Thus, the MV delta in a stochastic volatility model is the standard delta plus an additional term that is non-zero only when the two Brownian motions driving price and the volatility are correlated.

⁶ This is also known as local risk minimization, and has been studied extensively in the context of incomplete markets by Schweizer (1991), Bakshi et al. (1997), Bakshi et al. (2000), Frey (1997), Lee (2001) and others. The advantage of using minimum variance hedge ratios is their tractability and intuition. However, like other quadratic hedging strategies, minimum variance hedging treats losses and gains in a symmetric manner and one may prefer an alternative hedging strategy, such as super-hedging or utility maximization. Refer to Cont and Tankov (2004, chapter 10) for a review. Our definition of the minimum variance hedge ratio is consistent with the definition of the minimum variance futures hedge ratio suggested by Johnson (1960) and Ederington (1979).

The MV gamma, γ_{mv}^{sv} , can be derived by setting

$$\langle d\delta_{mv}^{sv} - \gamma_{mv}^{sv} dS, dS \rangle = 0 \Rightarrow \gamma_{mv}^{sv} = \frac{\langle d\delta_{mv}^{sv}, dS \rangle}{\langle dS, dS \rangle} \tag{10}$$

and applying Itô’s formula to $\delta_{mv}^{sv}(T, \mathbf{K}; t, S, \sigma)$ to obtain

$$\begin{aligned} \gamma_{mv}^{sv} &= \frac{d^2g}{dS^2} = (\delta_{mv}^{sv})_S + (\delta_{mv}^{sv})_\sigma \frac{\langle d\sigma, dS \rangle}{\langle dS, dS \rangle} \\ &= g_{SS} + \left(2g_{S\sigma} \frac{d\sigma}{dS} + g_{\sigma\sigma} \left(\frac{d\sigma}{dS} \right)^2 + g_\sigma \frac{d^2\sigma}{dS^2} \right), \end{aligned} \tag{11}$$

where the second-order total derivative on the right-hand side is interpreted as

$$\frac{d^2\sigma}{dS^2} \equiv \left(\frac{d\sigma}{dS} \right)_S + \left(\frac{d\sigma}{dS} \right)_\sigma \frac{d\sigma}{dS}. \tag{12}$$

We conclude that in stochastic volatility models with uncorrelated Brownian motions (e.g. Hull and White, 1987; Nelson, 1990; Stein and Stein, 1991; and others) the MV delta is equal to the standard delta and if these models are also scale-invariant, the MV delta is model-free and equal to the standard delta.

The same observation holds true for the gamma. But this is not true for stochastic volatility models with non-zero price–volatility correlation. For instance the Heston (1993) model

$$\begin{aligned} \frac{dS}{S} &= \mu dt + \sqrt{V} dB \\ dV &= a(m - V)dt + b\sqrt{V} dZ \quad \langle dB, dZ \rangle = \rho dt \end{aligned} \tag{13}$$

is scale-invariant.⁷ Hence

$$\begin{aligned} \delta_{mv}^{heston} &= g_S + g_V \frac{\langle b\sqrt{V} dZ, S\sqrt{V} dB \rangle}{\langle S\sqrt{V} dB, S\sqrt{V} dB \rangle} = g_S + g_V \left(\frac{\rho b}{S} \right) \\ \gamma_{mv}^{heston} &= g_{SS} + \frac{\rho b}{S} \left(\frac{\rho b}{S} g_{VV} + 2g_{SV} - \frac{1}{S} g_V \right) \end{aligned} \tag{14}$$

and the only model-dependent part of the hedge ratio is the second term on the right-hand side. In the case of equity options, when ρ is typically negative, the Heston MV delta is lower (greater) than the model-free delta if the vega g_V is positive (negative). This implies that the model-free delta over-hedges (under-hedges) equity options relative to the MV delta, and should be less efficient for pure delta hedging.

More generally, the MV delta and gamma account for the total effect of a change in the underlying price, including the indirect effect of the price change on the claim price via its effect on the volatility (or any other parameter that is correlated with the underlying price).

⁷ The variance process is correlated with the price process but it is independent of the *scale* of the price. Hence \sqrt{V} is dimensionless and the Heston model is scale-invariant. This follows from Theorem 1. The derivation of the returns density is not necessary to verify that the model is scale-invariant.

3.3. Local volatility hedge ratios

Now consider the hedge ratios for local volatility (LV) models, scale-invariant or otherwise. As these models do not introduce new sources of uncertainty, the dynamics of the claim price $g = g^{lv}(T, \mathbf{K}; t, S)$ are given by

$$dg = g_t dt + g_S dS + \frac{1}{2}g_{SS} dS^2 \tag{15}$$

(c.f. (6) for stochastic volatility models) and the local volatility hedge ratios are the standard (partial derivative) hedge ratios

$$\begin{aligned} \delta^{lv}(T, \mathbf{K}; t, S) &= \frac{\langle dg, dS \rangle}{\langle dS, dS \rangle} = \frac{\langle g_S dS, dS \rangle}{\langle dS, dS \rangle} = g_S \\ \gamma^{lv}(T, \mathbf{K}; t, S) &= \frac{\langle d\delta^{lv}, dS \rangle}{\langle dS, dS \rangle} = g_{SS}. \end{aligned} \tag{16}$$

These are the minimum variance hedge ratios for any local volatility model in which the instantaneous volatility is a static function of t and S (e.g. Cox, 1975; Dumas et al., 1998; and others) and in particular for implied tree models (e.g. Derman and Kani, 1994; Rubinstein, 1994). But note that such models are not scale-invariant.

In scale-invariant local volatility models, the instantaneous volatility is dimensionless and is typically a function of S/S_0 . This implies that the local volatility surface is static with respect to S/S_0 rather than with respect to S . The surface ‘floats’ with the asset price. As a result, the hedge ratios in (16) are not the minimum variance ratios because movements in the local volatility surface are correlated with movements in the underlying asset. The solution is to treat the local volatility model as a stochastic volatility model with perfect price–volatility correlation, as follows.

In the stochastic volatility case, the second source of randomness from the volatility process motivates the adjustment to the hedge ratios shown in (9) and (11); but in local volatility models there is just one source of randomness. Nevertheless, because the instantaneous volatility $\sigma(t, S)$ in a local volatility model is a deterministic function of t and S , it is also a continuous process and it has dynamics given by Itô’s formula as

$$d\sigma = \sigma_t dt + \sigma_S dS + \frac{1}{2}\sigma_{SS} dS^2 = (\sigma_t + \frac{1}{2}\sigma^2 S^2 \sigma_{SS}) dt + \sigma_S dS \tag{17}$$

which can be interpreted as a stochastic volatility model with perfect correlation between the instantaneous volatility and the underlying asset price.

Now using (9) and (11) the MV local volatility hedge ratios of the claim price $\hat{g} = g^{lv}(T, \mathbf{K}; t, S, \sigma)$ are

$$\begin{aligned} \delta_{mv}^{lv} &= \hat{g}_S + \hat{g}_\sigma \sigma_S \\ \gamma_{mv}^{lv} &= \hat{g}_{SS} + (2\hat{g}_{S\sigma} \sigma_S + \hat{g}_{\sigma\sigma} (\sigma_S)^2 + \hat{g}_\sigma \sigma_{SS}). \end{aligned} \tag{18}$$

The difference between (18) and the hedge ratios (16) is that in (18) the partial derivatives of the claim price with respect to S are computed while keeping the volatility σ constant. If σ is an explicit parameter of the model the partial derivatives \hat{g}_σ , $\hat{g}_{S\sigma}$ and $\hat{g}_{\sigma\sigma}$ are well-defined. Otherwise, it may be possible to re-parameterize the model in terms of this. This distinction is important because the hedge ratios from scale-invariant local volatility models

are model-free, by [Theorem 3](#), and they are different from the MV hedge ratios (18) just as the Heston hedge ratios are different from the MV ratios (14) when there is price–volatility correlation.

4. Empirical results

This section compares the hedging performance of a selection of option pricing models using the standard delta and gamma and the MV hedge ratios. On testing the model-free hedge ratios from different scale-invariant models no significant difference between the model's performances was found. We have therefore used the [Heston \(1993\)](#) model as a representative scale-invariant model. Its delta and gamma are model-free but if the price–volatility correlation is non-zero the MV hedge ratios (14) will be different from the model-free hedge ratios. The CEV model ([Cox, 1975](#)) is included because it is not scale-invariant and hence has the potential to generate significantly different results. Its standard hedge ratios are not model-free, but they are equal to the MV hedge ratios given by (16). Finally, the Black–Scholes (BS) hedge ratios are used as a benchmark.

4.1. Data

Bloomberg data on the June 2004 European call options on the S&P 500 index, i.e. daily close prices from 16 January 2004 to 15 June 2004 (101 business days) for 34 different strikes (from 1005 to 1200), have been applied in this study. Only the strikes within $\pm 10\%$ of the current index level were used for the model's calibration each day but all strikes were used for the hedging strategies. Implied volatilities are computed from mid prices (i.e. the average of bid and ask option prices). Options whose mid prices were below the intrinsic value or unrealistic were discarded. Time series of daily USD Libor rates were downloaded from the British Bankers Association (BBA) website for several maturities and used as a proxy for the risk-free rate. Linear interpolation was applied to produce a continuous function of the Libor rate with respect to time to maturity. This procedure was repeated for every date in the sample. Dividend yields were calculated by inverting the arbitrage-free pricing formula for a futures contract, i.e. $F = Se^{(r-q)(T-t)}$, where for all $t < T$, S and F are the close values of the S&P 500 index and of the S&P 500 futures with expiry in June 2004, respectively. The calculation of the dividend yield is not exact since the S&P 500 futures market closes 15 min later than the spot market. However the impact from the measurement error of the dividend yield is negligible in equity markets, and for short maturities in particular. During the period, no trend was observed for the S&P 500 index: the average daily return was only 0.02% with an annualised standard deviation (volatility) of 11.96%.

The delta hedge strategy consists of one delta-hedged short call on each available strike, 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 δ (delta) of the underlying asset, where δ is determined by both the model and the option's characteristics. The portfolio is rebalanced daily, assuming zero transaction costs, stopping on 2nd June because from then until the expiry date the fit to the smile worsened considerably for the models considered here. The delta–gamma hedge strategy again consists of a short call on each strike, but this time an amount of the 1125 option, which is closest to at-the-money 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 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 (profit and loss) with 1324 observations is obtained.

4.2. Calibrated hedge ratios

Each model was calibrated daily by minimizing the root-mean-square-error between the model implied volatilities and the market implied volatilities of the options used in the calibration set. We used the closed-form solution for the Heston model based on Fourier transforms (Lewis, 2000), chose a risk-aversion parameter of zero and set the long-term volatility at 12%.⁸ The calculation of the CEV hedge ratios is based on the non-central chi-square distribution result of Schroder (1989). For the BS model, the deltas and gammas are obtained directly from the market data and there is no need for model calibrations.

The deltas and gammas of each model, whilst changing daily, exhibit some strong patterns when they are plotted by strike or by moneyness: the same shapes emerge day after day. In Fig. 1a and b 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 S&P 500 call options. In Fig. 1a, the scale-invariant model-free delta is greater than the BS delta for all but the very high strikes. So if the BS model over-hedges in presence of the skew (as shown by Coleman et al., 2001) then scale-invariant models should perform worse than the BS model. A different picture emerges when MV hedge ratios are used. In the CEV model, where the MV hedge ratios are the same as the standard hedge ratios (see Section 3.3), and in the Heston model, where the MV hedge ratios are the model-free hedge ratios adjusted according to (14), the MV deltas are generally lower than the BS deltas. Another pattern is observed in Fig. 1b for the gammas. The model-free gammas are lower than the BS gamma for in-the-money calls and greater than the BS gamma for out-of-the-money calls (except for very deep out-of-the-money calls) while the opposite is observed when MV gammas are considered. So partial price sensitivities will under-hedge/over-hedge the gamma risk for in-the-money/out-of-the-money calls respectively, relative to the BS hedges.

4.3. Distribution of hedging profit and loss

Table 1 reports the sample statistics of the aggregate daily P&L for each model, over all options and over all days in the hedging period. The models are ordered by the standard deviation of the daily P&L. Small skewness and excess kurtosis in the P&L distribution are 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

⁸ In Lewis (2000) pricing formula, a risk-aversion parameter of 0 implies a logarithmic utility for the investor, whilst risk neutrality requires a parameter of 1. The investor's risk aversion is irrelevant for the calculation of the standard delta and gamma in the Heston model (because these are model-free) but it may influence the MV hedging performance. Nevertheless, the calibration of the model under different assumptions for the risk-aversion parameter did not produce significantly different MV hedge ratios (results available from Authors on request). Hence the risk-aversion parameter appears to be of lesser importance for hedging than the correlation coefficient.

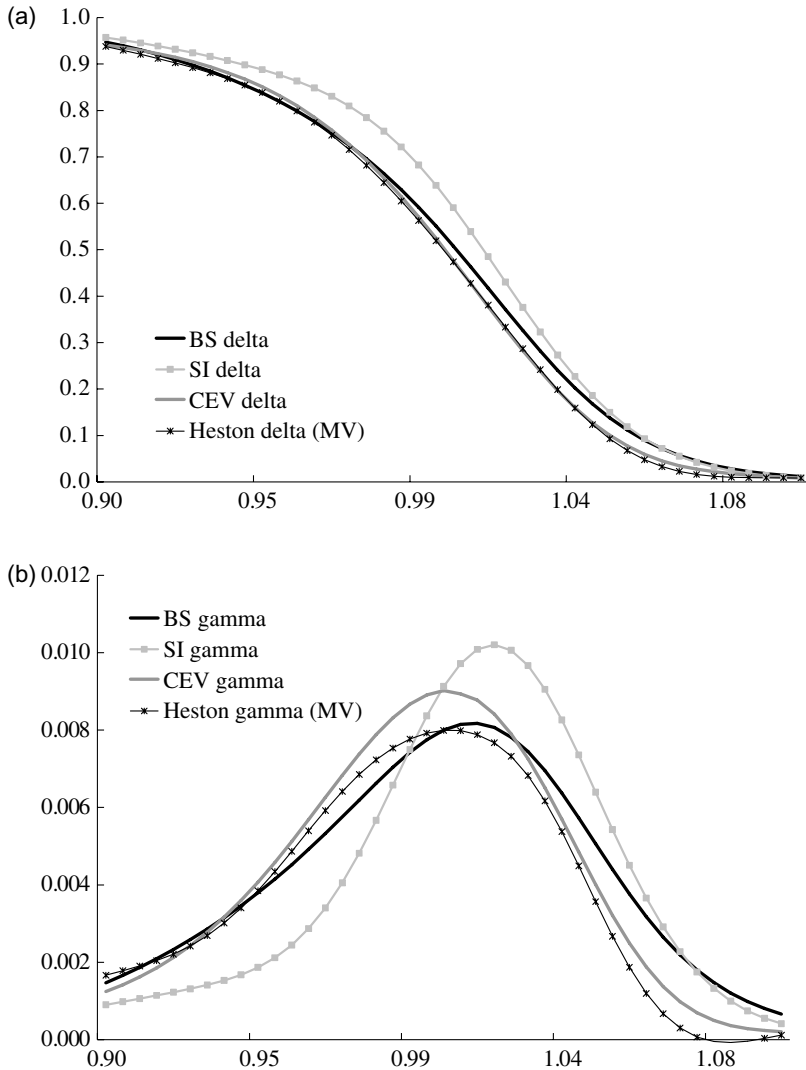


Fig. 1. The models' delta and gamma by moneyness on May 21st 2004 for S&P 500 call options: (a) shows the minimum variance (MV) delta of the Heston model, the model-free delta of scale-invariant (SI) models, and the deltas of the Black–Scholes (BS) and CEV models (for which the standard deltas are also MV). (b) shows the corresponding gammas. In each figure, the hedge ratios are drawn as functions of K/S and May 21st was chosen as a day when all the hedge ratios exhibited their typical pattern. The SI deltas are greater than the BS deltas in general, whilst the minimum variance deltas (CEV and Heston (MV)) are typically lower than the BS deltas. The SI gammas are lower than the BS gamma for low strike options and greater than the BS gamma for high strike options (except for exceptionally high strikes) while the opposite is observed when MV gammas are considered.

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 S&P 500 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 S&P 500

Table 1
Sample statistics of the aggregate daily P&L for delta hedging

Model	Mean	Std. Dev.	Skewness	Excess Kurtosis	R^2
<i>(a) Delta hedging</i>					
CEV	0.1462	0.5847	−0.3424	0.7820	0.113
Heston (MV)	0.1370	0.6103	−0.5704	1.6737	0.152
BS	0.1401	0.7451	−0.7029	2.0370	0.412
SI	0.1373	1.1788	−0.5928	1.4834	0.693
<i>(b) Delta–gamma hedging</i>					
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
SI	0.0428	0.4548	0.0208	4.0123	0.060

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 and delta–gamma hedging strategies with daily rebalancing. The models are ordered by the standard deviation of the daily P&L. Small skewness and excess kurtosis are desirable. We also 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 S&P 500 returns. The R from this regression, reported in the last column of the table, is small when the hedge is effective. The models are BS (Black and Scholes, 1973), CEV (constant elasticity of variance of Cox, 1975), Heston (MV) (the Heston model using minimum variance hedge ratios) and SI (using the model-free hedge ratios of scale-invariant models).

returns. The lower the R^2 from this regression, reported in the last column, the more effective the hedge.

According to these criteria, the best delta hedges are obtained from the MV hedge ratios, irrespective of the underlying model used. The MV deltas yield lower standard deviations than the BS delta, and these also have P&L that are closest to being normally distributed according to the observed skewness and excess kurtosis. Conversely, the model-free deltas perform worse than the BS delta. Apart from this, the positive mean P&L for delta hedging is a result of the short volatility exposure and gamma effects, since we have only rebalanced daily (see also Bakshi et al., 1997, and Lee, 2001). The delta–gamma hedging results in part (b) of Table 1 show a mean P&L that is close to zero. On adding a gamma hedge it is remarkable that the BS model performance improves considerably, whilst the other models ranked more or less as before.

One possible explanation for the superiority of the BS model in Table 1b is that the same hedging strategy is used to gamma hedge or vega hedge vanilla options: the ratio of the gammas is equal to the ratio of the vegas in the BS model. This is evidence that most of the imperfections of the BS model can be dealt with by hedging the movements in implied volatility. In fact, Bakshi et al. (1997) also find that vega hedging with the BS model performs well except for low strike in-the-money call options.

Results on hedged portfolio P&L standard deviation by moneyness, averaged over all days in our sample are given in Table 2. This table shows that the apparent superiority of the BS model for delta–gamma hedging is due to its success at hedging the strikes slightly higher than at-the-money. This may be linked to our finding in Fig. 1 that the BS gamma is similar to the MV gammas for near-the-money options. For out-of-the-money calls, the MV hedge ratios from the Heston model give the lowest standard deviation of hedged portfolio P&L. Hedging performance is particularly bad when the model-free hedge ratios are used.

Table 2
Standard deviation of the daily P&L aggregated by moneyness of option

K/S	0.90–0.95		0.95–1.00		1.00–1.05		1.05–1.10		1.10–1.15	
<i>(a) Delta hedging</i>										
Best	Heston (MV)	0.3714	CEV	0.5740	CEV	0.6372	CEV	0.6051	Heston (MV)	0.5507
	CEV	0.3854	Heston (MV)	0.6161	Heston (MV)	0.6629	Heston (MV)	0.6202	CEV	0.5602
Worst	BS	0.5652	BS	0.7876	BS	0.7844	BS	0.6921	BS	0.5917
	SI	0.7357	SI	1.2055	SI	1.2691	SI	1.0283	SI	0.7746
<i>(b) Delta–gamma hedging</i>										
Best	Heston (MV)	0.1801	CEV	0.2358	BS	0.2531	Heston (MV)	0.2907	Heston (MV)	0.3134
	CEV	0.1853	BS	0.2561	CEV	0.3040	CEV	0.2923	CEV	0.3222
	BS	0.2012	Heston (MV)	0.2594	Heston (MV)	0.3132	BS	0.2929	BS	0.3597
Worst	SI	0.3214	SI	0.3695	SI	0.4271	SI	0.5277	SI	0.5175
# Options	141		476		435		217		55	

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 Black–Scholes (BS) model performs best only for the delta–gamma hedging of near at-the-money options. The model-free hedge ratios of scale-invariant (SI) models perform worst irrespective of the option moneyness or hedging strategy. The minimum variance hedge ratios (CEV and Heston (MV)) perform best and only a small difference between their hedging performances is observed. Table 3 shows that this difference is not significant.

The hedging performance of the Heston model has also been considered by Bakshi et al. (1997) and Nandi (1998), among others, and their findings agree with the results reported above.⁹ Nandi (1998) investigates the importance of the correlation coefficient in the Heston model and concludes that the model's delta–vega hedging performance is significantly improved when the correlation coefficient is not constrained to be zero. That paper also finds that, after taking into account the transactions costs (bid–ask spreads) in the index options market and using S&P 500 futures to hedge, the stochastic volatility model outperforms the Black–Scholes model only if correlation is not constrained to be zero.

Bakshi et al. (1997) consider minimum variance delta hedging and ‘delta-neutral’ hedging (using as many hedging instruments as there are sources of risk, except for jump risk) and compare the hedging performance of models that include stochastic volatility, jumps and/or stochastic interest rates. They find that a stochastic volatility model such as Heston (1993) is adequate for price hedging. In fact, 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 out-of-the-money options, but for other options stochastic volatility is the most important factor to model.

4.4. Testing for differences between the models

Fig. 2a and b plot the cumulative distribution functions of the hedging P&L, taken over all options and over all days in the sample. Fig. 2a depicts the P&L from delta hedging only and Fig. 2b depicts the P&L from delta–gamma hedging. In both charts, there are two distinct groups: the MV hedging strategies (CEV and Heston (MV)) and the hedging strategies based on the (model-free) hedge ratios of any scale-invariant (SI) model that fits the market option prices. The former group is more efficient because it produces a P&L distribution that is less dispersed around the mean. The BS model lies in between the two groups in (a) and very close to the MV hedges in (b). The P&L for delta–gamma hedging with the SI models is also slightly shifted to the right. These findings are consistent with Table 1, which reports the moments of the same distributions.

Applying a Kolmogorov–Smirnov test (Massey, 1951; Siegel, 1988) to these distribution functions yields the results in Table 3. The null hypothesis is that the two P&L distributions are the same and the Kolmogorov–Smirnov statistic is asymptotically χ^2 distributed with two degrees of freedom. Significant values at the 10%, 5% or 1% levels are marked with one, two or three asterisks, respectively. The results confirm our theoretical findings. There are very significant differences between the P&L from MV deltas and gammas and the P&L from the model-free deltas and gammas. However, no significant difference is found between the two MV strategies for delta and for delta–gamma hedging. Both CEV and Heston models provide an effective delta or

⁹ Bakshi et al. (2000) start from the general SVSI-J model (Bakshi et al., 1997) and, by fixing some of the model parameters, they investigate the performance of alternative models for pricing and hedging options of different maturities. The SVSI-J model is scale-invariant. As a result, the standard deltas for the specific models considered there (SV, SVSI and SVJ) are model-free and should be equal *if* the models fit the same option prices.




Fig. 2. Cumulative distribution functions of the hedging P&L, taken over all options and over all days in the sample: In both charts there are two distinct groups: the minimum variance (MV) hedging strategies (CEV and Heston (MV)) and the non-MV hedging strategies based on the model-free hedge ratios of scale-invariant (SI) models. The former group is more efficient because it produces a P&L that is less dispersed. The BS model lies in between the two groups in (a) and very close to the MV hedges in (b). (a) Delta Hedge P&L c.d.f. (b) Delta-Gamma Hedge P&L.

delta-gamma hedge for S&P 500 call options. Finally, the differences between the BS P&L and the P&L from the MV hedge ratios are significant for delta hedging but not for delta-gamma hedging.

Table 3
Kolmogorov–Smirnov test results

	BS	SI	CEV	Heston (MV)
<i>(a) Delta hedge P&L distribution functions</i>				
BS	–	29.889***	5.114*	4.923*
SI	29.889***	–	52.664***	51.297***
CEV	5.114*	52.664***	–	1.232
Heston (MV)	4.923*	51.297***	1.232	–
<i>(b) Delta–gamma hedge P&L distribution functions</i>				
BS	–	35.212***	1.232	2.327
SI	35.212***	–	33.293***	32.409***
CEV	1.232	33.293***	–	0.742
Heston (MV)	2.327	32.409***	0.742	–

This table reports Kolmogorov–Smirnov statistics for the null hypothesis that two P&L cumulative distribution functions are the same. The test statistic is χ^2 distributed with two degrees of freedom. Significant values at 10%, 5% or 1% levels are marked with one, two or three asterisks, respectively. The hedging performance of scale-invariant (SI) models is significantly different from the performance of the other models. No significant difference is found between the hedging performances of the CEV and Heston (MV) models. The differences between the BS P&L and the P&L from the MV hedge ratios (CEV and Heston (MV)) are significant for delta hedging but not for delta–gamma hedging.

The similarity in the performance of MV hedges is certainly intriguing as these hedge ratios were not expected to be model-free.¹⁰ Since the CEV and Heston models have been calibrated to the same implied volatility smile we do expect them to produce roughly the same local volatility surface at the calibration time, as follows from the forward equation (Dupire, 1996; Derman and Kani, 1998). Yet each model assumes different underlying price dynamics, so both the option price dynamics and the local volatility dynamics will differ from one model to another. Thus it is not intuitively obvious why the MV hedge ratios should be the same for both models. If true, this would add an important constraint to the permissible dynamics of local volatility, a result that is left to further research.

5. Summary and conclusions

Merton (1973) identified that scale invariance leads to the homogeneity of vanilla option prices. More recently Bates (2005) proved that it 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.

This paper uses the scale-invariant property to address two challenging questions: are there significant differences between the price hedge ratios of these models and are such hedge ratios optimal for dynamic hedging? To answer these questions we have extended the work of Bates (2005), who examined a limited set of models applied only to vanilla options and did not consider the optimality of partial derivatives as hedge ratios. We have shown that the scale-invariant property is common to the vast majority of models in the

¹⁰ In an earlier version of this paper we also considered the hedging performance of the SABR model of Hagan et al. (2002) and found that the MV hedge ratios in this model produced hedging P&L distributions that were not significantly different from the other MV hedging P&L distributions. However the standard hedge ratios in the SABR model (which are not model-free, as the model is not scale-invariant) performed poorly.

option pricing literature, that model-free results extend to all claims with homogeneous pay-off functions, and that model-free hedge ratios only have the minimum variance property for scale-invariant stochastic volatility models if price–volatility correlation is zero.

Our results show that to classify a model as scale-invariant or otherwise, one does not need to know the returns density. In fact one does not even need an explicit price process. Moreover, scale invariance preserves the homogeneity of any contingent claim pay-off throughout the life of the claim. In fact, for any claim with homogeneous pay-off, a model is scale-invariant if and only if the claim price is homogeneous at all times. Then we prove that all partial derivatives of the claim price with respect to the underlying price are given by linear combinations of the claim price and its derivatives with respect to the claim characteristics. Thus scale invariance implies that price hedge ratios will be model-free for any claim with a homogeneous pay-off and claim prices that are observable in the market.

We then showed how minimum variance (MV) hedge ratios require an adjustment to the model-free delta and gamma of scale-invariant models whenever there is a non-zero correlation between the underlying price and any other stochastic component of the model. Empirical results on S&P 500 index options showed that, whilst the standard (model-free) hedge ratios of scale-invariant models perform worse than the BS model, MV hedge ratios provide better hedges on average. Our results also reveal a remarkable similarity in the performance of MV hedges, indicating that some model-free relationship may hold even for MV hedge ratios.

There remains much scope for empirical and theoretical research arising from the results in this paper: we have restricted the present study to local and stochastic volatility models but an extension to general semi-martingales is possible; and the behaviour of scale-invariant models under other hedging strategies, such as super-hedging, utility maximization or mean-variance hedging, remains to be explored.

Acknowledgement

We thank two anonymous referees for their excellent comments on an earlier version of this paper. Very useful comments are also gratefully acknowledged from Dr. Bruno Dupire, Bloomberg New York; 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.

Appendix

Proof of Theorem 1. From the definition of X_t and (1),

$$\begin{aligned} X_t = \frac{S_t}{S_0} &\iff dX_t = \frac{dS_t}{S_0} = \frac{S_t}{S_0} \frac{dS_t}{S_t} \iff \frac{dX_t}{X_t} = \frac{dS_t}{S_t} \\ &= \Theta' d\Lambda \iff X_T = X_0 + \int_0^T \Theta' X_t d\Lambda. \end{aligned} \quad (19)$$

Since $X_0 = 1$, X_T is independent of S_0 if Θ is dimensionless, i.e. homogeneous of degree zero in S . Hence, Θ is at most a function of the past history of X_t but not of S_0 or S_t explicitly. Finally, since Λ includes only the time t , Wiener processes and jump processes, the fact that S_t is a semi-martingale implies that Θ satisfies the regularity conditions for the coefficients of a semi-martingale and the integral in (19) is well-defined. \square

Proof of Theorem 2. First consider European-type claims whose pay-off at expiry, $G(S_T, \mathbf{K})$, is homogeneous of degree k , that is

$$G(uS_T, u\mathbf{K}) = u^k G(S_T, \mathbf{K}) \quad u \in \mathbb{R}^+ \tag{20}$$

We show that the process for S is scale-invariant if and only if

$$g(T, u\mathbf{K}; t, uS) = u^k g(T, \mathbf{K}; t, S) \quad \forall t \in [0, T]. \tag{21}$$

Define the numeraire N_t so that $Z_{t,T} = N_t/N_T$ is independent of S and \mathbf{K} . Also define the relative price $X_{t,T} = S_T/S_t$ so that a model is scale-invariant if and only if $X_{t,T}$ is dimensionless relative to S . It follows from martingale theory (Harrison and Kreps, 1979; Harrison and Pliska, 1981) that:

$$g(T, \mathbf{K}; t, S) = E^{\mathcal{Q}_N} \left[G(S_T, \mathbf{K}) \frac{N_t}{N_T} \middle| \mathfrak{F}_t \right] = E^{\mathcal{Q}_N} [G(S_t X_{t,T}, \mathbf{K}) Z_{t,T} | \mathfrak{F}_t] \quad t \in [0, T], \tag{22}$$

where the expectation is conditional on information up to time t , denoted by \mathfrak{F}_t , and is under the martingale measure \mathcal{Q}_N associated with the numeraire (see also Geman, 2005). Now apply the substitutions $S \mapsto uS$ and $\mathbf{K} \mapsto u\mathbf{K}$, and assume (20). As $Z_{t,T}$ and $X_{t,T}$ are invariant under scaling in S and \mathbf{K} , we have

$$\begin{aligned} g(T, u\mathbf{K}; t, uS) &= E^{\mathcal{Q}_N} [G(uS_t X_{t,T}, u\mathbf{K}) Z_{t,T} | \mathfrak{F}_t] = E^{\mathcal{Q}_N} [G(uS_T, u\mathbf{K}) Z_{t,T} | \mathfrak{F}_t] \\ &= u^k E^{\mathcal{Q}_N} [G(S_T, \mathbf{K}) Z_{t,T} | \mathfrak{F}_t] = u^k g(T, \mathbf{K}; t, S) \quad \forall t \in [0, T]. \end{aligned} \tag{23}$$

For the converse, suppose the pay-off function is homogeneous of degree k but that the model is not scale-invariant. Then the relative price $X_{t,T}$ is not dimensionless and scaling $S \mapsto uS$ implies $X_{t,T} \mapsto X_{t,T}^u$ where $X_{t,T}^u \neq X_{t,T}$ in general. Hence, there exists at least one time t at which

$$G(uS_t X_{t,T}^u, u\mathbf{K}) \neq G(uS_t X_{t,T}, u\mathbf{K}) \text{ almost surely} \tag{24}$$

so that, replacing into (22), we have

$$g(T, u\mathbf{K}; t, uS) \neq u^k g(T, \mathbf{K}; t, S) \tag{25}$$

and the claim price is not a homogeneous function of degree k . It follows that if the claim price at every time t is a homogeneous function of degree k , then the price process must be scale-invariant.

The above argument only applies to claims without the possibility of early exercise. The extension to American/Bermudan claims follows because if a European claim price is homogeneous of degree k at all times, then so is the American/Bermudan equivalent. At any time t before expiry, the claim is either exercised and its value equals the pay-off $G(S_t, \mathbf{K})$, which is homogeneous by assumption, or not exercised and the claim value follows the same p.d.e. as the European claim, which is homogeneous for all t . Thus, the American/Bermudan claim price is also homogeneous of degree k for all t . Conversely, if

the pay-off were homogeneous but the American/Bermudan price were not, then the European price could not be homogeneous because they are based on the same p.d.e., and the price process would not be scale-invariant.

For clarity **Theorem 2** supposed that the pay-off depends on the value of S_T only, yet the pay-off of a path-dependent claim can be a function of the whole path of S . This is not a problem since the martingale argument can also be applied to path-dependent claims. See e.g. Schweizer (1991). □

Proof of Corollary 1

(i) \iff (ii): The implied volatility is the volatility parameter in the Black and Scholes (1973) model that equates the Black–Scholes (BS) price C^{bs} to the price C of a standard European call or put option (Latané and Rendleman, 1976). That is:

$$C(T, K; t, S) = C^{bs}(T, K; t, S, \theta(T, K; t, S)). \tag{26}$$

Merton (1973) proved that scale invariance implies the price of a standard European option is homogeneous of degree one (this also follows from **Theorem 2**), hence

$$C(T, K; t, S) = SC\left(T, \frac{K}{S}; t, 1\right) \quad \text{and} \tag{27}$$

$$C^{bs}(T, K; t, S, \theta(T, K; t, S)) = SC^{bs}\left(T, \frac{K}{S}; t, 1, \theta(T, K; t, S)\right). \tag{28}$$

Now by (26)

$$C\left(T, \frac{K}{S}; t, 1\right) = C^{bs}\left(T, \frac{K}{S}; t, 1, \theta(T, K; t, S)\right). \tag{29}$$

Since $\theta(T, K; t, S)$ is implicitly defined in terms of K/S by (29), it is homogeneous of degree zero in S and K . Conversely, if the implied volatility is homogeneous of degree zero, then (26) implies that the European option price C will be homogeneous of degree one in S and K because the BS price C^{bs} is homogeneous of degree one. Thus, by **Theorem 2**, the process must be scale-invariant.

(i) \iff (iii): From Dupire’s equation (Dupire, 1996; Derman and Kani, 1998) we have

$$\hat{\sigma}^2(T, K; t, S) = 2(C_T + (r - q)KC_K + qC)/K^2C_{KK}, \tag{30}$$

where C_T , C_K and C_{KK} are partial derivatives of the price $C(T, K; t, S)$ of a standard European option with expiry T and strike K . Then, define $h(x) = C(T, x; t, 1)$ and using (27) it follows that for every $x = \frac{K}{S}$

$$\begin{aligned} C_T(T, K; t, S) &= Sh_T(x), \\ C_K(T, K; t, S) &= h_x(x), \\ C_{KK}(T, K; t, S) &= h_{xx}(x) \frac{1}{S} \end{aligned} \tag{31}$$

and replacing into (30),

$$\hat{\sigma}^2(T, K; t, S) = 2(h_T(x) + (r - q)xh_x(x) + qh(x))/x^2h_{xx}(x). \tag{32}$$

That is, the local volatility is a function of the moneyness K/S and not of K and S separately, hence it is homogeneous of degree zero. Conversely, if the local volatility is

homogeneous of degree zero, it follows from Theorem 1 that there is a scale-invariant local volatility model (an ‘equilibrium’ model, according to Derman and Kani, 1998) that fits all vanilla option prices and, from Theorem 2, these prices must be homogeneous of degree one. Hence, from Theorem 2 again, the original price process is scale-invariant.

Proof of Theorem 3. Since S is generated by a scale-invariant process, Theorem 2 yields

$$g(T, u\mathbf{K}; t, uS) = u^k g(T, \mathbf{K}; t, S) \quad \forall t \in [0, T]. \tag{33}$$

Differentiating (33) with respect to u and setting $u = 1$ we obtain

$$Sg_S(T, \mathbf{K}; t, S) + \mathbf{K}'g_{\mathbf{K}}(T, \mathbf{K}; t, S) = kg(T, \mathbf{K}; t, S) \tag{34}$$

which is the well-known Euler’s theorem for homogeneous functions. After re-arranging, this gives the expression for g_S in (4). For g_{SS} , we differentiate (33) twice with respect to u and set $u = 1$ to obtain

$$S^2g_{SS} + 2S\mathbf{K}'g_{\mathbf{K}S} + \mathbf{K}'g_{\mathbf{K}\mathbf{K}}\mathbf{K} = k(k - 1)g. \tag{35}$$

On differentiating (34) with respect to S we obtain

$$\mathbf{K}'g_{\mathbf{K}S} = (k - 1)g_S - Sg_{SS}. \tag{36}$$

Combining (34)–(36) gives the expression for g_{SS} in the theorem. Now assume $g_{S^m} = \sum_{i=0}^m \mathbf{A}_i g_{\mathbf{K}^i} \mathbf{B}_i$ for $m \geq 1$, where g_{S^m} denotes the m th partial derivative of g with respect to S and $(g_{\mathbf{K}^i})_{m^i}$ is the i -dimensional matrix of i th partial derivatives of g with respect to \mathbf{K} , and in particular we define $g_{\mathbf{K}^0} = g$. Note that $\mathbf{A}_i(S, \mathbf{K})$ and $\mathbf{B}_i(S, \mathbf{K})$ are known matrices at time t . It follows that

$$g_{S^{m+1}} = (g_{S^m})_S = \sum_{i=0}^m [(\mathbf{A}_i)_S g_{\mathbf{K}^i} \mathbf{B}_i + \mathbf{A}_i (g_{\mathbf{K}^i})_S \mathbf{B}_i + \mathbf{A}_i g_{\mathbf{K}^i} (\mathbf{B}_i)_S], \tag{37}$$

where

$$(g_{\mathbf{K}^i})_S = (g_S)_{\mathbf{K}^i} = S^{-1}(kg - \mathbf{K}'g_{\mathbf{K}})_{\mathbf{K}^i} = S^{-1}((k - i)g_{\mathbf{K}^i} - \mathbf{K}'g_{\mathbf{K}^{i+1}}) \tag{38}$$

so that we may write $g_{S^{m+1}} = \sum_{i=0}^{m+1} \tilde{\mathbf{A}}_i g_{\mathbf{K}^i} \tilde{\mathbf{B}}_i$ for some matrices $\tilde{\mathbf{A}}_i(S, \mathbf{K})$ and $\tilde{\mathbf{B}}_i(S, \mathbf{K})$. As m is arbitrary, we conclude that all partial derivatives with respect to S are linear combinations of $g_{\mathbf{K}^i}$. \square

References

- Avellaneda, M., Levy, A., Parás, A., 1995. Pricing and hedging derivative securities in markets with uncertain volatilities. *Applied Mathematical Finance* 2, 73–88.
- Bakshi, G., Cao, C., Chen, Z., 1997. Empirical performance of alternative option pricing models. *Journal of Finance* 52, 2003–2049.
- Bakshi, G., Cao, C., Chen, Z., 2000. Pricing and hedging long-term options. *Journal of Econometrics* 94, 277–318.
- Bakshi, G., Madan, D., 2002. Average rate claims with emphasis on catastrophe loss options. *Journal of Financial and Quantitative Analysis* 37 (1), 93–115.
- Bates, D.S., 2003. Empirical option pricing: A retrospection. *Journal of Econometrics* 116, 387–404.
- Bates, D.S., 2005. Hedging the smirk. *Financial Research Letters* 2 (4), 195–200.

- Bergman, Y., Grundy, B., Wiener, Z., 1996. General properties of option prices. *Journal of Finance* 51 (5), 1573–1610.
- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637–659.
- Brigo, D., Mercurio, F., 2002. Lognormal-mixture dynamics and calibration to market volatility smiles. *International Journal of Theoretical and Applied Finance* 5 (4), 427–446.
- Carr, P., Geman, H., Madan, D., Yor, M., 2004. From local volatility to local Lévy models. *Quantitative Finance* 4 (October), 581–588.
- Coleman, T., Kim, Y., Li, Y., Verma, A., 2001. Dynamic hedging with a deterministic local volatility function model. *Journal of Risk* 4 (1), 63–89.
- Cont, R., 2006. Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical Finance* 16 (3), 519–547.
- Cont, R., Tankov, P., 2004. *Financial Modelling with Jump Processes*. Chapman & Hall/CRC Press.
- Cox, J.C., 1975. Notes on Option Pricing I: Constant Elasticity of Variance Diffusions. Working Paper, Stanford University.
- Cox, J.C., Ross, S.A., 1976. The valuation of options for alternative stochastic processes. *Journal of Financial Economics* 3, 145–166.
- Cox, J.C., Ingersoll, J.E., Ross, S.A., 1985. A theory of the term structure of interest rates. *Econometrica* 53, 385–408.
- Derman, E., 1996. Model risk. *Risk* 9 (5).
- Derman, E., Kani, I., 1994. Riding on a smile. *Risk* 7 (2), 32–39.
- Derman, E., Kani, I., 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.
- Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68 (6), 1343–1376.
- Dupire, B., 1994. Pricing with a smile. *Risk* 7 (1), 18–20.
- Dupire, B., 1996. A unified theory of volatility. In: Carr, P. (Eds.), Working Paper, now published in *Derivatives Pricing: The Classic Collection*, Risk Books, 2004.
- Dumas, B., Fleming, F., Whaley, R., 1998. Implied volatility functions: Empirical tests. *Journal of Finance* 53 (6), 2059–2106.
- Ederington, L.H., 1979. The hedging performance of the new futures markets. *Journal of Finance* 34 (1), 157–170.
- Eraker, B., 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *Journal of Finance* 59 (3), 1367–1403.
- Frey, R., 1997. Derivative asset analysis in models with level-dependent and stochastic volatility. *CWI Quarterly* 10 (1), 1–34.
- Geman, H., 2005. From measure changes to time changes in asset pricing. *Journal of Banking and Finance* 29, 2701–2722.
- Green, T.C., Figlewski, S., 1999. Market risk and model risk for a financial institution writing options. *Journal of Finance* 54 (4), 1465–1499.
- Hagan, P.S., Kumar, D., Lesniewski, A.S., Woodward, D.E., 2002. Managing smile risk. *Wilmott Magazine* (September), 84–108.
- Harrison, J.M., Kreps, D., 1979. Martingales and arbitrage in multiperiod securities market. *Journal of Economic Theory* 20, 381–408.
- Harrison, J.M., Pliska, S., 1981. Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and their Applications* 11, 381–408.
- 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., Neumann, C.D.D., 2001. Local scale-invariance and contingent claim pricing. *International Journal of Theoretical and Applied Finance* 4 (1), 1–21.
- Hull, J., White, A., 1987. The pricing of options on assets with stochastic volatilities. *Journal of Finance* 42 (2), 281–300.
- Jackwerth, J.C., 1999. Option-implied risk-neutral distributions and implied binomial trees: A literature review. *The Journal of Derivatives* 7, 66–82.
- Johnson, L.L., 1960. The theory of hedging and speculation in commodity futures. *Review of Economic Studies* 27, 139–151.

- Latané, H.A., Rendleman, R.J., 1976. Standard deviations of stock price ratios implied in option prices. *Journal of Finance* 31 (2), 369–381.
- Lee, R.W., 2001. Implied and local volatilities under stochastic volatility. *International Journal of Theoretical and Applied Finance* 4 (1), 45–89.
- Lewis, A., 2000. *Option Valuation under Stochastic Volatility with Mathematica Code*, first ed. Finance Press.
- Massey Jr., F.J., 1951. The Kolmogorov–Smirnov test for goodness of fit. *Journal of the American Statistical Association* 46, 68–78.
- 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.
- Nandi, S., 1998. How important is the correlation between returns and volatility in a stochastic volatility model? Empirical evidence from pricing and hedging in the S&P 500 index options market. *Journal of Banking and Finance* 22, 589–610.
- Nelson, D.B., 1990. ARCH models as diffusion approximations. *Journal of Econometrics* 45, 7–38.
- Psychoyios, D., Skiadopoulos, G., 2006. Volatility options: Hedging effectiveness, pricing, and model error. *Journal of Futures Markets* 26 (1), 1–31.
- Psychoyios, D., Skiadopoulos, G., Alexakis, P., 2003. A review of stochastic volatility processes: Properties and implications. *Journal of Risk Finance* 4 (3), 43–60.
- Rubinstein, M., 1994. Implied binomial trees. *Journal of Finance* 49 (3), 771–818.
- Schoutens, W., 2003. *Lévy Processes in Finance*, first ed. John Wiley & Sons.
- Schroder, M., 1989. Computing the constant elasticity of variance option pricing formula. *Journal of Finance* 44 (1), 221–229.
- Schweizer, M., 1991. Option hedging for semimartingales. *Stochastic Processes and their Applications* 37, 339–363.
- Siegel, S., 1988. *Nonparametric Statistics for the Behavioral Sciences*, second ed. McGraw-Hill.
- Skiadopoulos, G., 2001. Volatility smile-consistent option models: A survey. *International Journal of Theoretical and Applied Finance* 4 (3), 403–438.
- Stein, E.M., Stein, J.C., 1991. Stock price distributions with stochastic volatility: An analytic approach. *Review of Financial Studies* 4, 727–752.