

Stochastic Volatility Correction to Black-Scholes

Jean-Pierre Fouque* George Papanicolaou† K. Ronnie Sircar‡

July 1999; revised August 1999. This version January 2000.

Stochastic volatility models have become popular for derivative pricing and hedging in the last ten years as the existence of a nonflat implied volatility surface (or term-structure) has been noticed and become more pronounced, especially since the 1987 crash. This phenomenon, which is well-documented¹, stands in empirical contradiction to the consistent use of a classical Black-Scholes (constant volatility) approach to pricing options and similar securities. However, it is clearly desirable to maintain as many of the features as possible that have contributed to this model's popularity and longevity, and the natural extension pursued both in the literature and in practice has been to modify the specification of volatility in the stochastic dynamics of the underlying asset price model.

There are many stories behind why we should model volatility to be a random process. For example, it could simply represent estimation uncertainty, or it can arise as a friction from transaction costs, or it could simulate non-Gaussian (heavy-tailed) returns distributions. In other words, stochastic volatility is a far-reaching extension of the Black-Scholes lognormal model, describing a much more complex market.

Any extended model must also specify what data it is to be calibrated from. The pure Black-Scholes procedure of estimating from historical stock data *only* is not possible in an incomplete market if one takes the view (as we shall) that the market selects a unique risk-neutral derivative pricing measure, from a family of possible measures, which reflects its degree of "crash-o-phobia". This pricing measure is reflected in traded at-the-money European options prices, so, as is common practice, this "smile data" is used for calibration.

Parameter estimation and stability of the estimates in time presents the major mathematical and practical challenge here. Without a formula for option prices under a particular stochastic volatility model, estimating the risk-neutral parameters is computationally intensive (we have to run a tree or simulations at each step in an iterative search procedure). Often models are chosen so that there is a closed-form solution, and this usually means taking the volatility to be independent of the Brownian motion driving the stock price, whereas

*Department of Mathematics, NC State University, Raleigh NC 27695-8205, fouque@math.ncsu.edu.

†Department of Mathematics, Stanford University, Stanford CA 94305, papanico@math.stanford.edu

‡Department of Mathematics, University of Michigan, Ann Arbor MI 48109-1109, sircar@umich.edu.

Work supported by NSF grant DMS-9803169.

¹See for example, Jackwerth, J. and Rubinstein, M., 1996, *Recovering Probability Distributions from Contemporaneous Security Prices*, Journal of Finance 51(5), pages 1611-1631, and Rubinstein, M., 1985, *Nonparametric Tests of Alternative Option Pricing Models*, Journal of Finance 40(2), pages 455-480.

common experience (and empirical evidence) suggests a negative correlation: when volatility goes up, stock prices tend to go down.

Now suppose that we have estimated the parameters. How good will these estimates be tomorrow? It is possible to have a very tight fit over a short time, but often these break down significantly thereafter. This is certainly a problem with estimates of volatility surfaces², where (unlike in stochastic volatility modeling), volatility is modeled to be a function of time and stock price with no independent randomness.

We present here a new approach to stochastic volatility that has the following features:

- It requires that volatility be mean-reverting, but, other than that, does not depend in an essential way on how the volatility is modeled.
- It translates the slope and intercept of the implied volatility skew into information about the correlation between volatility and stock price shocks and the market's volatility risk premium.
- It simplifies enormously the parameter estimation problem.
- It gives a recipe for pricing (and hedging) other derivatives in a stochastic volatility environment by identifying their *effective* approximating derivative security in a constant volatility environment. This includes barriers, Asians and Americans.
- It produces parameter estimates from implied volatility skews that are stable.

This is achieved by exploiting the mean-reverting behaviour of volatility and the much-noted observation that volatility is persistent.

Framework

The stock price $(S_t)_{t \geq 0}$ satisfies

$$dS_t = \mu S_t dt + \sigma_t S_t dW_t,$$

where $(\sigma_t)_{t \geq 0}$ is the volatility process. To incorporate the correlation with the Brownian motion (W_t) (which leads to the implied volatility skew in these models), it is convenient to take (σ_t) to be a diffusion process too, although it can have jumps as well. To fix ideas, we shall write volatility as a positive function of a mean-reverting Gaussian (Ornstein-Uhlenbeck) process: $\sigma_t = f(Y_t)$, where

$$dY_t = \alpha(m - Y_t)dt + \beta \left(\rho dW_t + \sqrt{1 - \rho^2} dZ_t \right),$$

with Z an independent Brownian motion (the source of the additional randomness) and

²See Dumas, B., Fleming, J. and Whaley, R., 1998, *Implied Volatility Functions: Empirical Tests*, Journal of Finance 53(6), pages 2059-2106 for an empirical study of this and, for a mathematical analysis, Lee, R., 1999, *Local Volatilities under Stochastic Volatility*, to appear in International Journal of Theoretical and Applied Finance.

$$\begin{aligned}
\alpha &= \text{the rate of mean-reversion,} \\
m &= \text{the long-run mean of } Y, \\
\beta &= \text{the “v-vol”,} \\
\rho &= \text{the correlation coefficient.}
\end{aligned}$$

Our objective is to analyze the effect of stochastic volatility in the basic Black-Scholes model. Therefore, we assume that these parameters, as well as the rate of return μ are, for simplicity, constant. For the same reason we assume that $f(Y_t)$ does not depend on t explicitly. In fact we do not need to specify f in detail since the mean reversion asymptotics give results that are insensitive to all but a few general features of f . The precise model for the process (Y_t) driving the volatility does not matter either, so long as it is an ergodic process like the Ornstein-Uhlenbeck process above.

Three observable quantities emerge from the asymptotics and they are the only ones that must be calibrated from historical data and the term structure of volatility. These quantities are complicated functions of the primitive model parameters, the function f and the market price of volatility risk γ introduced below, which need not be calibrated separately. This is a new approach to the study of stochastic volatility models.

Volatility Persistence

It is often noted in empirical studies of stock prices that volatility is persistent or bursty - for days at a time it is high and then, for a similar length of time, it is low. However, over the lifetime of a derivative contract (a few months), there are many such periods, and looked at on this timescale, volatility is fluctuating fast, but not as fast as the rapidly changing stock price.

In terms of our model, we say that the volatility process is fast mean-reverting relative to the yearly timescale, but slow mean-reverting by the tick-tick timescale. Since the derivative pricing and hedging problems we study are posed over the former period, we shall say that volatility exhibits fast mean-reversion without explicitly mentioning the longer timescale of reference.

The rate of mean-reversion is governed by the parameter α , in annualized units of years⁻¹. Fast mean-reversion means that α is in fact large and that $\beta^2/(2\alpha)$, the variance of the invariant distribution of the OU process, is a stable $\mathcal{O}(1)$ constant. As an illustration, Figure 1 shows simulated volatility paths for the model above in which $\alpha = 1$ on the top and $\alpha = 200$ below. In practice the volatility process is not directly observable. In fact the true observable is the de-measured returns process

$$\frac{dS_t}{S_t} - \mu dt = \sigma_t dW_t,$$

at discrete times. In Figure 2 we show simulated trajectories of this returns process corresponding to these volatility trajectories in the first two graphs. We observe that the *size of the fluctuation* in the returns process of the top picture is relatively constant over time, while for the large α picture, in the middle, it is changing a lot. This is exactly what we call fast mean-reverting or persistent stochastic volatility.

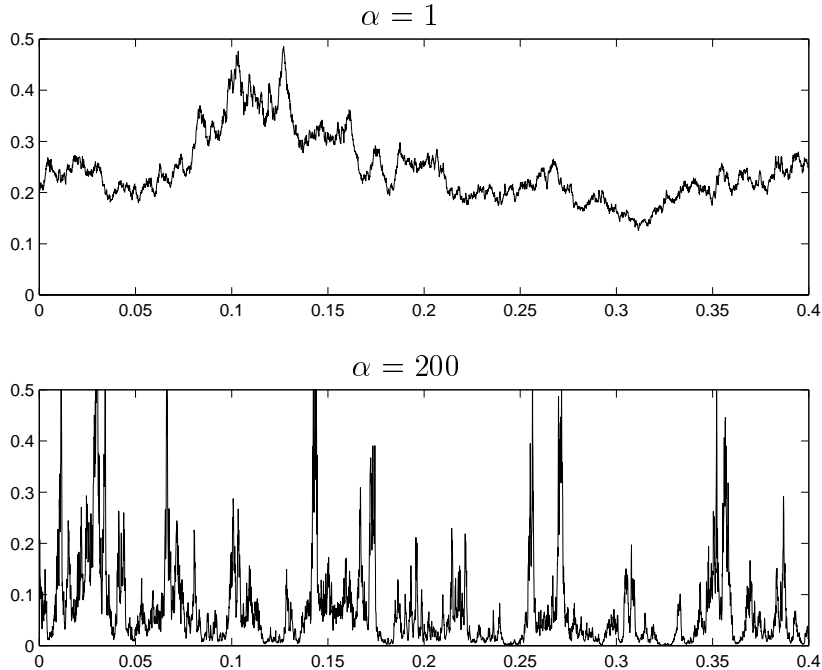


Figure 1: *Simulated volatility for small and large rates of mean-reversion for the OU model, with the choice $f(y) = e^y$. Note how volatility “clusters” in the latter case.*

In the bottom graph of Figure 2, we show the returns process for the S&P 500 over the first five months of 1996. By comparing with the top two paths, we note how the structure of the returns fluctuation size resembles more the one from a *fast* mean-reverting stochastic volatility model (middle), than one from a model with slow mean-reversion (top). A detailed analysis of the mean reverting structure of the S&P 500 is presented in a forthcoming paper³.

In the methodology we describe next we will not need the precise *value* of α , only that it is large.

Derivative Pricing

We start with European derivatives with terminal payoff $h(S_T)$. The *no arbitrage* price P_t depends on the present stock price S_t and the present level of the volatility-driving process Y_t . It is given by the risk-neutral expected discounted payoff

$$P_t = \mathbb{E}^{Q(\gamma)} \{ e^{-r(T-t)} h(S_T) | S_t, Y_t \},$$

where $Q(\gamma)$ represents probabilities in the risk-neutral world and γ is the volatility risk premium⁴. In this world,

$$dS_t = rS_t dt + \sigma_t S_t dW_t^*,$$

³Fouque, Papanicolaou, Sircar and Solna, 1999, *Mean-Reversion of S&P500 Volatility*.

⁴A stochastic volatility model is an incomplete market model, and as such there is a whole family of risk-neutral pricing measures, unlike in the constant volatility case when there is only one. It can be shown from Girsanov’s theorem, that each possible measure in this Itô framework corresponds to a choice of the volatility risk premium, also known as the market price of volatility risk. See, for example, Hull, J. and

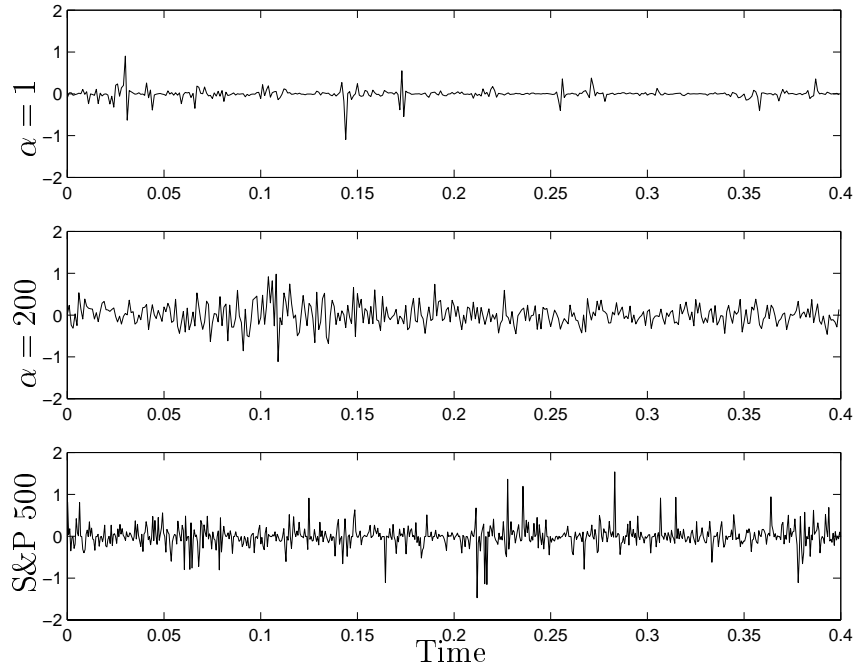


Figure 2: The top and middle graphs show simulated returns for small and large rates of mean-reversion for the OU model, with the choice $f(y) = e^y$. The bottom graph shows 1996 S&P 500 returns computed from half-hourly data.

$$\begin{aligned} \sigma_t &= f(Y_t), \\ dY_t &= \left[\alpha(m - Y_t) - \beta \left(\rho \frac{(\mu - r)}{f(Y_t)} + \gamma(Y_t) \sqrt{1 - \rho^2} \right) \right] dt + \beta \left(\rho dW_t^* + \sqrt{1 - \rho^2} dZ_t^* \right), \end{aligned}$$

where (W^*, Z^*) are independent Brownian motions under the Equivalent Martingale Measure $Q(\gamma)$. The interest rate r is assumed to be constant and known. As for the other model parameters, we assume a market price of volatility risk $\gamma(Y_t)$ driven only by the volatility-driving process (Y_t) . This can be validated *a posteriori* when the formula presented below is fitted to implied volatility data. The new function $\gamma(y)$ is included in the asymptotics and not directly estimated (using other derivatives for instance).

The main result of an asymptotic analysis of this problem⁵ says that when volatility persists, we can approximate the derivative price P_t in the stochastic volatility environment by pricing a more complicated (path-dependent) security in the Black-Scholes constant volatility environment. The payoff structure of the new security depends on the Black-Scholes pricing formula for the original one and accounts appropriately for volatility risk.

The procedure is as follows:

1. Let $(\bar{S}_t)_{t \geq 0}$ be the Black-Scholes lognormal model:

$$d\bar{S}_t = r\bar{S}_t dt + \bar{\sigma}\bar{S}_t dW_t^*,$$

White, A., 1987, *The Pricing of Options on Assets with Stochastic Volatilities*, Journal of Finance 42(2), pages 281-300.

⁵See Fouque, Papanicolaou and Sircar, 1999, *Mean-Reverting Stochastic Volatility*, to appear in International Journal of Theoretical & Applied Finance, for details of the calculation.

with $\bar{\sigma}$ the usual constant historical volatility. Price the derivative with this model. That is, find

$$P_{BS}(t, S) = \mathbb{E}^* \{ e^{-r(T-t)} h(\bar{S}_T) | \bar{S}_t = S \}.$$

2. The stochastic volatility price P_t is well-approximated by the price *in the Black-Scholes model* \tilde{P}_t of the path-dependent contract with the same terminal payoff $h(\bar{S}_T)$ and a payout rate $H(t, \bar{S}_t)$ between times t and T , given by

$$H(t, S) = V_2 S^2 \frac{\partial^2 P_{BS}}{\partial S^2}(t, S) + V_3 S^3 \frac{\partial^3 P_{BS}}{\partial S^3}(t, S),$$

where V_2 and V_3 are constants related to the original parameters (α, β, m, ρ) and the functions f and γ . They contain information about the market, but are not specific to any derivative contract. That is, \tilde{P} is given by the **modified Black-Scholes formula**

$$\tilde{P}_t = \mathbb{E}^* \left\{ e^{-r(T-t)} h(\bar{S}_T) - \int_t^T e^{-r(u-t)} H(u, \bar{S}_u) du | \bar{S}_t \right\}.$$

The path-dependent payment stream $H(t, S)$ is computed from the second and third order derivatives of the classical Black-Scholes price P_{BS} . It may be positive or negative and accounts dynamically for volatility randomness in a robust model-independent way. It also accounts for the market price of volatility risk effectively selected by the market.

3. There is a simple explicit formula whenever there is a formula for the Black-Scholes price of the contract $P_{BS}(t, S)$, for example calls, puts, binaries. It is given by

$$\tilde{P}_t = P_{BS}(t, S_t) - (T - t)H(t, S_t).$$

The Black-Scholes price is corrected by a term containing the **Gamma** of the contract and a term containing its third derivative which we name **Epsilon** since it is related to a small correction.

The error is of order less than $1/\sqrt{\alpha}$ which is small when mean-reversion is fast. Notice that this approximation does *not* depend on the present level of volatility, which is not directly observable and usually difficult to estimate.

Calibration

Where this simplification is extremely useful is in parameter estimation. We first estimate in the usual way the historical volatility $\bar{\sigma}$ which can be related to the model parameters: it is the square root of the average of f^2 with respect to the invariant distribution of the OU process Y_t . We are not using this explicit relation since we are not aiming at estimating the model parameters. The quantities V_2 and V_3 are easily calibrated by using the full term structure of the volatility smile across strikes and maturities. If we take $h(S_T) = (S_T - K)^+$, a call option, compute the approximation to the stochastic volatility price described above and then work out the implied volatility I , we obtain the simple formula

$$I = a \frac{\log(K/S)}{(T-t)} + b.$$

The parameters V_2 and V_3 are related to the smile parameters a, b and the mean volatility $\bar{\sigma}$ by

$$\begin{aligned} V_2 &= \bar{\sigma} \left((\bar{\sigma} - b) - a \left(r + \frac{3}{2} \bar{\sigma}^2 \right) \right), \\ V_3 &= -a \bar{\sigma}^3. \end{aligned}$$

It turns out that a has the same sign as ρ , so that a downward sloping skew indicates a negative correlation. In the particular uncorrelated case, $\rho = 0$, we deduce that, at this order of approximation, the smile gives only a constant correction to the historical volatility $\bar{\sigma}$. In that case, one can verify that the corrected price itself satisfies a Black-Scholes equation with the corrected volatility $\sqrt{\bar{\sigma}^2 + 2\bar{\sigma}(b - \bar{\sigma})}$ used effectively by the market (b is close to $\bar{\sigma}$ for fast mean-reverting stochastic volatility markets). This can be thought of as a pure kurtosis effect, as pointed out to us by a referee. The general case $\rho \neq 0$ is more interesting since it reflects the skew effect. The formulas above are derived in the paper cited in footnote 5.

So, the calibration procedure is as follows:

1. Fit near-the-money implied volatilities for several maturities, to a straight line in the composite variable called the *log-moneyness-to-maturity-ratio* (LMMR)

$$\text{LMMR} := \frac{\log \left(\frac{\text{Strike Price}}{\text{Stock Price}} \right)}{\text{Time to Maturity}}.$$

Estimate the slope a and the intercept b . Since $\text{LMMR} = 0$ when stock price = strike price, b is exactly the at-the-money implied volatility.

2. Estimate $\bar{\sigma}$, the historical volatility from stock price returns, and compute V_2 and V_3 using the formulas above.
3. Price any other European by pricing the adjusted claim (with the payout rate H for that contract) in a Black-Scholes world using the explicit formula given above. A similar procedure that needs only the market-describing parameters V_2 and V_3 holds for Asian, American⁶ and barrier⁷ options.

We stress again that this is *not* model specific: it does not depend on a particular choice of the functions f or γ or a particular ergodic driving diffusion Y , in the sense that many such choices will lead to the three observable quantities $\bar{\sigma}, a, b$ (or $(\bar{\sigma}, V_2, V_3)$) with no need to estimate the parameters (α, β, m, ρ) separately: only the V 's, which contain these, are needed. Nor is the present value Y_t required.

⁶See Fouque, Papanicolaou and Sircar, 1999, *From the Implied Volatility Skew to a Robust Correction to Black-Scholes American Option Prices*, preprint.

⁷See Fouque, Papanicolaou and Sircar, 1999, *Financial Modeling in a Fast Mean-Reverting Stochastic Volatility Environment*, Asia-Pacific Financial Markets 6(1), pages 37-48.

Stability of Parameter Estimates

We have undertaken, in the reference in footnote 3, a detailed empirical study of high-frequency S&P 500 index data to establish that volatility reverts slowly to its mean compared to the tick-by-tick scale fluctuations, but it reverts fast when looked at over the longer time scale of months. The key conclusion of this study is that while the rate of mean-reversion (in units $years^{-1}$) is large, it is a difficult parameter to estimate precisely, being the reciprocal of the correlation time of a *hidden* Markov process. However, the asymptotic derivatives theory does not need the value of α , only that it be large.

We have also tested *a posteriori* the feasibility of the theory-predicted LMMR linefit for actual implied volatility data. We show in Figure 3 daily estimates of the slope and intercept coefficients \hat{a} and \hat{b} from fitting Black-Scholes implied volatilities from observed S&P 500 European call option prices:

$$I^{obs}(t, S; K, T) = \hat{a} \left(\frac{\log(K/S)}{T-t} \right) + \hat{b},$$

across strikes K and maturities T .

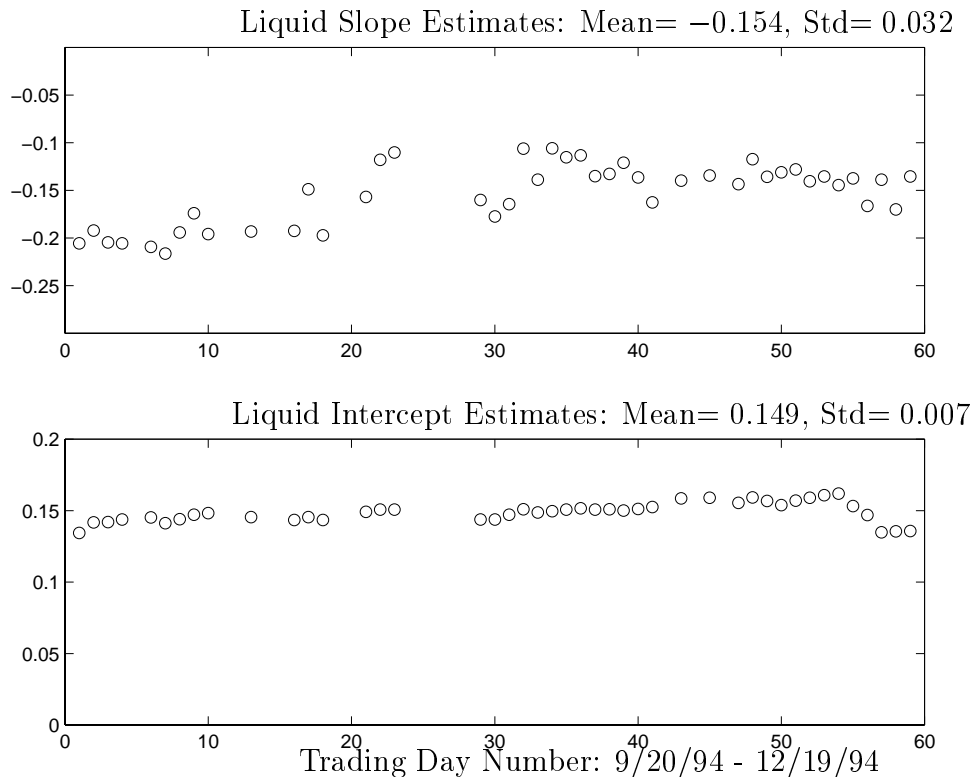


Figure 3: *Daily fits of S&P 500 European call option implied volatilities to a straight line in LMMR, excluding days when there is insufficient liquidity (16 days out of 60).*

We observe from the results that the slope coefficients \hat{a} are small. This strongly supports the fast mean-reverting hypothesis and validates use of the asymptotic formula as the full skew formula shows that a is a term of order $1/\sqrt{\alpha}$. We also find that the estimates \hat{a} and

\hat{b} over a sixty-day period are stable as attested to by the standard deviations of the daily estimates reported in the Figure.

Hedging

There is a related theory for hedging which we do not describe here in detail⁸. It relies on the parameter V_3 estimated from the smile to account for the effect of correlation or leverage. The amount of stock to hold is given by a correction to the Black-Scholes Delta by a combination of the Gamma, Epsilon and the fourth derivative which we call **Kappa**:

$$\frac{\partial P_{BS}}{\partial S}(t, S_t) - \frac{V_3(T-t)}{S_t} \left(4S_t^2 \frac{\partial^2 P_{BS}}{\partial S^2} + 5S_t^3 \frac{\partial^3 P_{BS}}{\partial S^3} + S_t^4 \frac{\partial^4 P_{BS}}{\partial S^4} \right) (t, S_t).$$

This strategy removes a bias in the shortfall of a Black-Scholes hedge used in a random volatility market. The new average hedging error (measured, for example, as the expected shortfall) is of order $1/\alpha$.

Conclusion

We have presented a modified Black-Scholes pricing and hedging methodology to account for persistent stochastic volatility. It is effective in identifying the important components of the implied volatility skew, from which we calibrate, and gives a recipe for pricing and hedging more exotic securities. As the implied volatility fits show, the calibrated group parameters are quite stable over time.

It dynamically captures the main skewness of the observed implied volatility surface, independent of specific modeling of volatility. This, plus the simplicity of the procedure, makes the results extremely suitable for practice.

⁸The hedging problem and a presentation from first principles of the work discussed here plus application to interest-rate derivatives appears in the forthcoming book *Derivatives in Financial Markets with Stochastic Volatility* by the authors of this article.