Financial Modeling in a Fast Mean-Reverting Stochastic Volatility Environment

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

November 17, 1998

Abstract

We present a derivative pricing and estimation methodology for a class of stochastic volatility models that exploits the observed "bursty" or persistent nature of stock price volatility. Empirical analysis of high-frequency S&P 500 index data confirms that volatility reverts slowly to its mean in comparison to the tick-by-tick fluctuations of the index value, but it is *fast* mean-reverting when looked at over the time scale of a derivative contract (many months). This motivates an asymptotic analysis of the partial differential equation satisfied by derivative prices, utilizing the distinction between these time scales.

The analysis yields pricing and implied volatility formulas, and the latter provides a simple procedure to "fit the skew" from European index option prices. The theory identifies the important group parameters that are needed for the derivative pricing and hedging problem for European-style securities, namely the average volatility and the slope and intercept of the implied volatility line, plotted as a function of the log-moneyness-to-maturity-ratio. The results considerably simplify the estimation procedure.

The remaining parameters, including the growth rate of the underlying, the correlation between asset price and volatility shocks, the rate of mean-reversion of the volatility and the market price of volatility risk are not needed for the asymptotic pricing formulas for European derivatives, and we derive the formula for a knock-out barrier option as an example. The extension to American and path-dependent contingent claims is the subject of future work.

Contents

*Department of Mathematics, North Carolina 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.

1 Introduction

This article summarizes a flexible methodology for stochastic volatility modeling which has the following features.

- It applies to a large class of volatility processes that are driven by an ergodic process that is tending rapidly (in a sense to be explained below) to its invariant measure.
- It incorporates a nonzero volatility risk premium that models the market's "crash-o-phobia", and a nonzero correlation between volatility and asset price shocks that explains the much-observed skew or leverage effect.
- An asymptotic analysis that exploits volatility clustering yields a simple pricing (and hedging) theory for European-style and some path-dependent contingent claims whose implementation requires solution of a PDE problem that is a minor extension of the corresponding classical Black-Scholes PDE problem for that security. In particular, where the Black-Scholes theory produces an explicit formula, so does the new theory.
- The parameters needed for the theory are easily "read from the skew". That is, calibration from near-the-money European option implied volatilities is simple and direct. The difficult-to-estimate volatility risk premium, correlation parameter, and persistence-time of the volatility are not explicitly needed. Further, the theory does not need estimation of today's volatility level.
- The theory can be extended to give a good approximation of the stochastic volatility corrected law of the risk-neutral asset price process that can be used to simulate, for example to price path-dependent and, in principle, American securities.

We outline the main results of this approach and cite references for the mathematical details and empirical motivation. We also present the pricing formula for a knock-out barrier option that can be used directly after calibration from the observed European-option skew. We conclude with a summary of ongoing and future work.

1.1 Motivation for Stochastic Volatility

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-documented in, for example, [6, 9], 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 in the literature and in practice has been to modify the specification of volatility in the stochastic dynamics of the underlying asset price model.

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". Thus at least *some* derivative data has to be used to price other derivatives, and much recent work uses *only* derivative data to estimate all the model parameters so that the assumed relationship between the dynamics of derivative prices and the dynamics of the underlying is not exploited at all.

We also refer the reader to recent surveys of the stochastic volatility literature such as [4, 5].

1.2 Mean-Reverting Diffusion Model

While the general asymptotic theory [2] can be given for volatility processes driven by any ergodic stochastic process with a unique invariant measure (*eg.* Markov chains, jump processes), it is convenient to present the analysis for a diffusion driving process, as is done in [1, 3]. The analysis in [10] is independent of specific modeling of the volatility process, but results in bands for option prices that describe potential volatility risk while obviating the need to estimate the risk premium. However, the market in at- and near-the-money European options is liquid and its historical data can be used to estimate this premium. We attempt this with a parsimonious model that is complex enough to reflect an important number of observed volatility features:

1. volatility is positive;

2. volatility is mean-reverting, but persists;

3. volatility shocks are negatively correlated with asset price shocks. That is, when volatility goes up, stock prices tend to go down and *vice-versa*. This is often referred to as leverage, and it at least partially accounts for a skewed distribution for the asset price that lognormal or zero-correlation stochastic volatility models do not exhibit.

1.3 Model

We present the results for models in which stock prices are conditionally lognormal, and the volatility process is a positive increasing function of a mean-reverting Ornstein-Uhlenbeck

(OU) process. That is,

$$\frac{dX_t}{X_t} = \mu dt + f(Y_t)dW_t, \tag{1}$$

$$dY_t = \alpha(m - Y_t)dt + \beta d\hat{Z}_t, \tag{2}$$

$$\hat{Z}_t := \rho W_t + \sqrt{1-\rho^2}Z_t,$$

where W and Z are independent Brownian motions, and ρ is the correlation between price and volatility shocks, with $|\rho| < 1$.

The solution to (2) is

$$Y_t = m + (Y_0 - m)e^{-\alpha t} + \beta \int_0^t e^{-\alpha(t-s)} d\hat{Z}_s, \tag{3}$$

and, given Y_0, Y_t is Gaussian,

$$Y_t - Y_0.e^{-\alpha t} \sim \mathcal{N}\left(m\left(1-e^{-\alpha t}\right), \nu^2\left(1-e^{-2\alpha t}\right)\right), \tag{4}$$

where $\nu^2 := \beta^2/(2\alpha)$. Thus Y has a unique invariant distribution, namely $\mathcal{N}(m, \nu^2)$, and is a simple building-block for a large class of stochastic volatility models described by choice of $f(\cdot)$. We call these models mean-reverting because the volatility is a monotonic function of a process Y whose drift pulls it towards the mean value m. The volatility is correspondingly pulled towards $f(m)$ *approximately*.

1.4 Fast mean reversion

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}. In [3], we present empirical evidence from S&P 500 data that α is in fact large and that ν^2 is a stable $\mathcal{O}(1)$ constant, so that our large-α option pricing formulas of Section 2 can be used.

As an illustration, Figure 1 shows sample stock price paths for the model (1-2) in which $\alpha = 1$ and $\alpha = 50$. Since, from (4), $\frac{1}{\alpha}\log 2$ is the time for the expected distance to the mean to halve, $\alpha = 1$ corresponds to 0.7 of a year (roughly 8 months), and $\alpha = 50$ corresponds to about half a week. Alternatively, under the invariant distribution $\mathcal{N}(m, \nu^2)$, the covariance of Y_s and Y_{s+t} is $\nu^2 e^{-\alpha t}$ and α^{-1} is the correlation time of the OU process. For $\alpha = 1$ this correlation time is a year while for $\alpha = 50$ it is about a week.

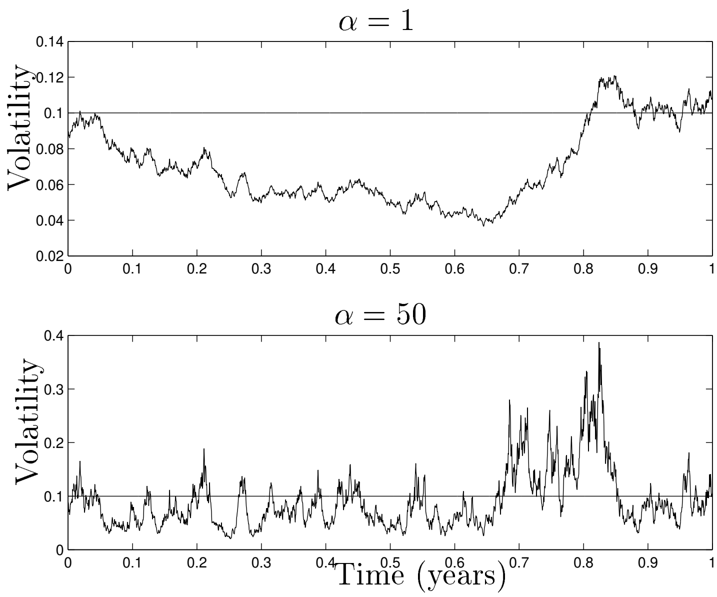

Figure 1: The top figure shows a simulated path of $f(Y_t) = e^{Y_t}$ with $\alpha = 1$, and the bottom one shows a path with $\alpha = 50$. In both cases, $\nu^2 = 0.25, (E\{e^{2Y_t}\})^{\frac{1}{2}} = 0.1$. Note how volatility "clusters" in the latter case.

1.5 Derivative Pricing

We are interested in pricing European-style derivative contracts on the underlying stock. When volatility is supposed constant, the classical Black-Scholes theory applies; when it is modelled as a stochastic process as here, the derivative price $C(t,x,y)$ is given by

$$C(t,x,y) = E^{Q(\gamma)}_{t,x,y}\{h(X_T)\}, \tag{5}$$

where $E^{Q(\gamma)}_{t,x,y}$ denotes the expectation given that $X_t = x$, $Y_t = y$, and under an Equivalent Martingale Measure (EMM) $Q(\gamma)$. The payoff function of the derivative is $h(x)$. Under such an EMM the discounted stock price is a martingale. By standard no-arbitrage pricing theory (see for example [7]), there is more than one possible EMM because the market is incomplete (the volatility is not a traded asset); the nonuniqueness is denoted by the dependence of Q on γ, the market price of volatility risk.

We shall assume that γ is constant because it has to be estimated from market data, at- and near-the-money call option prices in [3]. Most studies take $\gamma = 0$ for simplicity, but we take the view that the market selects a pricing measure identified by a particular γ which will be shown to occur in a simple manner in our pricing and implied volatility formulas, hence considerably simplifying estimation of its contribution to the observed skew. This can then be used to price more complicated derivatives in a consistent manner.

In [3], we analyze the PDE corresponding to (5) in the presence of fast mean-reversion:

$$\begin{aligned} C_t + \frac{1}{2}f(y)^2x^2C_{xx} + \rho\beta x f(y)C_{xy} + \frac{1}{2}\beta^2C_{yy} + & \\ r(xC_x - C) + (\alpha(m-y) - \beta\lambda(y))\,C_y &= 0, \quad (6)\\ C(T,x,y) &= h(x) \quad (7) \end{aligned}$$

where

$$\lambda(y) := \rho \frac{(\mu - r)}{f(y)} + \gamma\sqrt{1-\rho^2}. \tag{8}$$

There is also a (left) boundary condition

$$C(t, L, y) = g(t) \tag{9}$$

which does not, in general, depend on y. For example, for a European call, $L = 0, g = 0$. We also require that the solution not be "too singular" as $x \to \infty$; for example, linear growth is permissible. This is sufficient to identify a unique solution (see [11] for details). A barrier option is discussed in Section 4.

To summarize, the stochastic volatility model studied here is described by the five parameters $(m, \nu, \alpha, \rho, \gamma)$ which are, respectively, the mean m and the standard deviation ν of the invariant distribution of the driving OU process, the rate of mean reversion α, the skewness ρ, and the market price of volatility risk[1] γ. The last parameter cannot be estimated from historical asset price data. As we shall see in the next section, not all of these are needed for the pricing theory.

2 Main Result

1. When the rate of mean-reversion α is large (volatility persistence), the implied volatility curve from European call options is well-approximated by a straight line in the composite variable labelled the *log-moneyness-to-maturity-ratio* (LMMR)

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

   That is, if C^{call} is the stochastic volatility call option price satisfying (6-7) with $h(x) = (x-K)^+$, then I defined by

$$C^{\text{call}} = C^{BS}(I),$$

   where C^{BS} is the Black-Scholes formula, is given by

$$I = a\frac{\log(K/x)}{(T-t)} + b + \mathcal{O}(\alpha^{-1}).$$

   The parameters a and b are easily estimated as the slope and intercept of the linefit.

2. The price C^h of any other derivative satisfying a problem of type (6,7,9), for example binary options, barrier options, is given by

$$C^h = C_0 + C_1 + \mathcal{O}(\alpha^{-1}),$$

[1] A detailed study of possible ways to define this concept, along with other results, is given in [8].

where $C_0(\overline{\sigma})$ is the solution to the corresponding Black-Scholes problem with constant volatility $\overline{\sigma}$, and $C_1(t,x)$ solves

$$\mathcal{L}^{BS}(\overline{\sigma})C_1 = Vx^3\frac{\partial^3 C_0}{\partial x^3} + Wx^2\frac{\partial^2 C_0}{\partial x^3},$$

with

$$\mathcal{L}^{BS}(\overline{\sigma}) := \frac{\partial}{\partial t} + \frac{1}{2}\overline{\sigma}^2 x^2 \frac{\partial^2}{\partial x^2} + r\left(x\frac{\partial}{\partial x} - \cdot\right), \tag{10}$$

$$V := \overline{\sigma}^3 a, \tag{11}$$

$$W := a\overline{\sigma}^3 - \overline{\sigma}(b - \overline{\sigma}), \tag{12}$$

and $\overline{\sigma}$ is the long-run historical asset price volatility. The terminal condition is $C_1(T,x) = 0$ and the boundary condition is zero also: $C_1(t,L) = 0$.

The table below then distinguishes the model parameters from the parameters that are actually needed for the theory. The latter can be written as groupings of the former by the formulas given in [3], but for practical purposes, there is no need to do so. We pursue this in [3] for empirical completeness.

Model Parameters	Parameters that are needed
Growth rate of stock μ	Mean historical volatility of stock $\overline{\sigma}$
Long-run mean volatility m	
Rate of mean-reversion of volatility α	Slope of implied volatility linefit a
Volatility of volatility β	
Correlation between shocks ρ	Intercept of implied volatility linefit b
Volatility risk premium γ	

The three parameters on the right-side of the table are easily estimated and found to be quite stable from S&P 500 data in [3].

2.1 Empirical validation of fast mean-reversion

We have undertaken in [3] an extensive 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 an extremly 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.

A brief description of our validation procedure is as follows.

- We identify and use segments where the volatility process can be considered stationary. These turn out to be between one and a half to six months in length for the 1994 and 1995 datasets we have looked at, and across these segments, nonstationary effects would have to be taken into consideration. In fact studies of daily closing data, that is *low frequency* prices, often identify this timescale as characteristic of volatility persistence. However the intraday data highlights a shorter timescale of local volatility fluctuations.

- We extract the rate of mean-reversion from the Lorenz part of the spectrum of the logarithm of the squared de-meaned returns process. Such a spectral analysis is suited to the high-frequency data that we have, and provides a convenient graphical tool for picking off an order estimate for α. We find that the correlation time of the process is on the order of one to two days. Thus mean-reversion is fast over the timescale of months.

- We validate both the OU mean-reverting model and the estimation of the fast rate of mean-reversion by bootstrap, that is, comparison with spectra of simulated data. The method separates the intrinsic variability over segments of the model parameters from their statistical variability. Note that we do not expect parameters of the volatility process to be constant across the segments of stationarity.

3 Results of fitting the skew

To test the feasibility of the theory-predicted LMMR linefit for actual implied volatility data, we estimate in [3] 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:

$$I^{obs}(t,x;K,T) = \hat{a}\left(\frac{\log(K/x)}{T-t}\right) + \hat{b}. \tag{13}$$

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 in [3] shows that a is a term of order $1/\sqrt{\alpha}$. We also find that the estimates $\hat{a}$ and $\hat{b}$ within the segments of stationarity are relatively stable.

The following table separates the needed parameters, whose estimates are fairly stable, from the ones presented only for completeness, whose estimates have a high degree of uncertainty. The figures are for the 1994 S&P 500 .

Segment (length)	$\hat{\bar{\sigma}}$	$\hat{a}$	$\hat{b}$	$\widehat{\nu^2}$	$\hat{\mu}$	$\hat{\alpha}$	$\hat{\rho}$	$\hat{\gamma}$
1 (6 months)	0.1015	-0.1009	0.1410	0.9153	-0.1428	∼155	∼-0.11	∼-4.97
2 ($1\frac{1}{2}$ months)	0.0994	-0.1270	0.1286	0.7835	0.4170	∼155	∼-0.20	∼2.36
4 (3 months)	0.1030	-0.0888	0.1457	1.0794	-0.0695	∼155	∼-0.065	∼-4.07

Notice that the market price of volatility risk estmate $\hat{\gamma}$ inherits the variability of the growth rate estimate $\hat{\mu}$, but that, just as the Black-Scholes theory did not depend on μ, the fast mean-reverting stochastic volatility theory only depends upon a stable grouping of the γ and other parameters.

4 Example: Pricing a Barrier Option

We briefly sketch, for illustrative purposes, the extension of the pricing theory to more exotic securities by outlining the calculation for a knock-out barrier call option which gives the holder the right to buy the underlying asset on expiration date T for strike price K unless the asset price has hit the barrier H at any time before T, in which case the contract expires worthless. In what follows, we shall assume $H < K$.

In the stochastic volatility environment, the price $B(t,x,y)$ of the barrier option satisfies (6,7,9) with $h(x) = (x-K)^+$, and boundary condition $B(t,H) = 0$. Our fast mean-reverting approximation is

$$B = B^{(0)}(t,x) + B^{(1)}(t,x) + \mathcal{O}(\alpha^{-1}),$$

where $B^{(0)}$ is the Black-Scholes barrier price with constant volatility parameter $\overline{\sigma}$. The stochastic volatility correction $B^{(1)}$ satisfies the PDE problem

$$\mathcal{L}^{BS} B^{(1)} = V x^3 B^{(0)}_{xxx} + W x^2 B^{(0)}_{xx}, \text{ in } x > H, t < T \tag{14}$$

with zero terminal and boundary conditions. The operator $\mathcal{L}^{BS}$ is defined in (10), and the coefficients V and W are estimated from the historical volatility $\overline{\sigma}$ and the slope and intercept of the skew fit through the expressions (11) and (12).

Following the exposition in [11], $B^{(0)}(t,x)$ is obtained by the method of images and given by

$$B^{(0)}(t,x) = C^{BS}(t,x) - \left(\frac{x}{H}\right)^{1-k} C^{BS}(t, H^2/x),$$

where $C^{BS}(t,x)$ is the Black-Scholes formula *for a vanilla call option*, with the volatility parameter $\overline{\sigma}$, and $k := 2r/\overline{\sigma}^2$. The right-hand side of (14) is then given by

$$\begin{aligned} F(t,x) \quad := \quad & V x^3 C^{BS}_{xxx}(t,x) + W x^2 C^{BS}_{xx}(t,x) - \left(\frac{x}{H}\right)^{1-k} \left(W \frac{H^4}{x^2} C^{BS}_{xx}(t, H^2/x) \right. \\ & \left. - V \frac{H^6}{x^3} C^{BS}_{xxx}(t, H^2/x) + q(t, H^2/x) \right), \end{aligned} \tag{15}$$

with

$$\begin{aligned} q(t,x) &:= \theta C^{BS}(t,x) + \kappa x C^{BS}_{x}(t,x) + \chi x^2 C^{BS}_{xx}(t,x), \\ \theta &:= k(k-1)(W - V(k+1)), \\ \kappa &:= 2kW - 3k(k+1)V, \\ \chi &:= -3(k+1)V. \end{aligned}$$

Motivated by the translation and reflection invariance of the spatial part of the Black-Scholes operator $\mathcal{L}^{BS}$ in logarithmic co-ordinates moving at the drift rate r, we define the mirror operator $\mathcal{M}$ by

$$\mathcal{M} g(t,x) = \left(\frac{x}{H}\right)^{1-k} g(t, H^2/x).$$

Then the method of images says that the solution to $\mathcal{L}^{BS} B^{(1)} = F(t,x)$ in $x > H$ is given by solving

$$\mathcal{L}^{BS} v(t,x) = F(t,x) - \mathcal{M} F(t,x),$$

in $x > 0$ and restricting the solution to $x > H$.

From (15), we then only have to solve

$$\mathcal{L}^{BS} v(t,x) = 2W(x^2 C_{xx}^{BS}(t,x) - \mathcal{M}(x^2 C_{xx}^{BS}(t,x))) + q(t,x) - \mathcal{M}q(t,x),$$

on the full domain $x > 0, t < T$. Since the right-hand side is a function minus its mirror, it can be shown that we can ignore the mirror terms, solve and then subtract the mirror of the solution. Thus we need to solve

$$\mathcal{L}^{BS} u(t,x) = 2W x^2 C_{xx}^{BS}(t,x) + q(t,x),$$

with zero terminal and boundary conditions.

A convenient expression for the contribution to the solution from the first forcing term is obtained by explicit computation from the Black-Scholes formula and the Green's function for $\mathcal{L}^{BS}$. This computation appears in [1] or [3] since it is part of the solution for the regular call option. The second part of the solution can be written in terms of derivatives of the Black-Scholes formula with respect to r and σ, by noticing that

$$\begin{aligned}
\mathcal{L}^{BS}((T-t)C^{BS}(t,x)) &= -C^{BS}(t,x), \\
\mathcal{L}^{BS} C_r^{BS}(t,x) &= C^{BS}(t,x) - x C_x^{BS}(t,x), \\
\mathcal{L}^{BS} C_\sigma^{BS}(t,x) &= -\overline{\sigma} x^2 C_{xx}^{BS}(t,x).
\end{aligned}$$

Using these, we find

$$u = -2W \frac{x e^{-d_1^2/2}}{\overline{\sigma}\sqrt{2\pi}} \sqrt{T-t} - (\theta + \kappa)(T-t)C^{BS} - \kappa C_r^{BS} - \frac{\chi}{\overline{\sigma}} C_\sigma^{BS},$$

where d_1 is in standard Black-Scholes notation

$$d_1 = \frac{\log(x/K) + (r + \frac{1}{2}\overline{\sigma}^2)(T-t)}{\overline{\sigma}\sqrt{T-t}}.$$

Finally, $v(t,x) = u(t,x) - \mathcal{M}u(t,x)$ and $B^{(1)}(t,x)$ is the restriction of v to $x > H$:

$$B^{(1)}(t,x) = u(t,x) - \left(\frac{x}{H}\right)^{1-k} u(t, H^2/x).$$

The separate components of the formula are easily computed in closed form and the skew-calibrated parameters a and b inserted through V and W. There is of course no dependence on the unseen value of today's volatility $f(y)$.

5 Future directions

1. The estimation tools outlined here can now be used to validate a fast mean-reverting model for other high-frequency datasets. We are presently preparing an empirical study of S&P 500 index data from other years, as well as foreign exchange rate data.

2. The asymptotic approximation of the derivative prices can be improved to give a probability law that approximates the full risk-neutral pricing law. The full theory [2] will depend on more global features of the stochastic volatility model than just the parameters a and b, but it will be applicable to short-maturity and far-from-the-money contracts which are outside the region of validity of the present theory.

3. We are working on an asymptotic simplification of the American option pricing problem under stochastic volatility, which currently must be solved numerically.

4. The problem of computing optimal hedging strategies under constraints when volatility is random is unsolved. For example, to optimize the probability of a successful hedge with just the underlying given an initial cash input would require solving a degenerate Hamilton-Jacobi-Bellman equation. We are looking at simplifying this problem with separation of scales asymptotics.

References

[1] J.-P. Fouque, G. Papanicolaou, and K.R. Sircar. Asymptotics of a two-scale stochastic volatility model. In *Equations aux derivees partielles et applications, in honour of Jacques-Louis Lions*, pages 517–525. Gauthier-Villars, May 1998.

[2] J.-P. Fouque, G. Papanicolaou, and K.R. Sircar. General theory of fast mean-reverting stochastic volatility. *In preparation*, 1998.

[3] J.-P. Fouque, G. Papanicolaou, and K.R. Sircar. Mean-Reverting Stochastic Volatility. *Preprint*, November 1998.

[4] R. Frey. Derivative asset analysis in models with level-dependent and stochastic volatility. Technical report, Department of Mathematics, ETH Zűrich, November 1996.

[5] E. Ghysels, A. Harvey, and E. Renault. Stochastic volatility. In G. Maddala and C. Rao, editors, *Statistical Methods in Finance*, volume 14 of *Handbook of Statistics*, chapter 5, pages 119–191. North Holland, Amsterdam, 1996.

[6] J. Jackwerth and M. Rubinstein. Recovering probability distributions from contemporaneous security prices. *J. Fin.*, 51(5):1611–1631, 1996.

[7] I. Karatzas. *Lectures on the Mathematics of Finance.* CRM Monograph Series, American Mathematical Society, 1996.

[8] R. Lee. Local volatilities in stochastic volatility models. *Preprint*, November 1998.

[9] M. Rubinstein. Nonparametric Tests of Alternative Option Pricing Models. *J. Finance*, XL(2):455–480, June 1985.

[10] K.R. Sircar and G.C. Papanicolaou. Stochastic volatility, smile and asymptotics. *Applied Mathematical Finance*, to appear, 1998.

[11] P. Wilmott, J. Dewynne, and S. Howison. *Mathematics of Financial Derivatives: A Student Introduction.* Cambridge University Press, 1996.