BAYESIAN STATISTICS 7, pp. 000–000
J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid,
D. Heckerman, A. F. M. Smith and M. West (Eds.)
© Oxford University Press, 2003

Bayesian Inference for Derivative Prices

NICHOLAS G. POLSON
University of Chicago, USA
ngp@gsb.uchicago.edu

JONATHAN R. STROUD
University of Pennsylvania, USA
stroud@wharton.upenn.edu

SUMMARY

This paper develops a methodology for parameter and state variable inference using both asset and derivative price information. We combine theoretical pricing models and asset dynamics to generate a joint posterior for parameters and state variables and provide an MCMC simulation strategy for inference. There are several advantages of our inferential approach. First, more precise parameter estimates are obtained when both asset and derivative price information are used. Secondly, we provide a diagnostic tool for model misspecification based on agreement of the state and parameter estimates with and without derivative price information. Furthermore, the time series properties of the state variables can also be used to evaluate model fit. We illustrate our methodology using daily equity index options on the Standard and Poor's (S&P 500) index from 1998–2002.

Keywords: DERIVATIVE PRICING; MCMC; STOCHASTIC VOLATILITY; SMOOTHING; OPTION PRICING; LEVERAGE EFFECT.

1. INTRODUCTION

Bayesian inference using both asset and derivative price information presents a number of challenges. While many authors have considered the problem of parameter inference in asset pricing models, few have considered the problem of also incorporating derivative price information. The advantages of using additional derivative price information is twofold. First, more precise parameter and state variable estimates are obtained. Second, these estimates can be used to provide model misspecification diagnostics.

Financial asset pricing theory determines derivative prices as a conditional expectation of a discounted payoff given the current state variables and parameter values. The expectation is taken with respect to a risk-neutral pricing kernel which is derived from the underlying asset price dynamics by a transformation of the underlying parameters. For closed form derivative pricing under a wide range of models we refer the reader to Duffie, Pan and Singleton (2000). See Johannes and Polson (2002) for further discussion. The joint inference problem, then, is to combine this information with the usual asset pricing model for the dynamics of the underlying asset and state variables.

Our empirical work focuses on the commonly used leverage stochastic volatility model for equity index options. This model uses a number of features for explaining asset and derivative prices: an unobserved mean reverting stochastic volatility state variable and a vector of parameters including a correlation between the underlying asset returns and volatility. Closed form derivative pricing is available and this includes an extra parameter governing the market price of volatility risk.

Figure 1 provides motivation for the additional market price of volatility risk parameter by comparing implied volatilities with historical filtered volatilities (see Polson, Stroud and Müller, 2002) using daily S&P 500 equity index return data for 1990–2001. The implied volatilities are consistently higher than the filtered ones due to a number of reasons. The main factor is that option pricing accounts for the market price of volatility risk. Option pricing is determined as a conditional expectation under the risk neutral measure and the market price of volatility risk effectively raises the level of volatility under which pricing occurs. This is discussed further in Section 2.2.

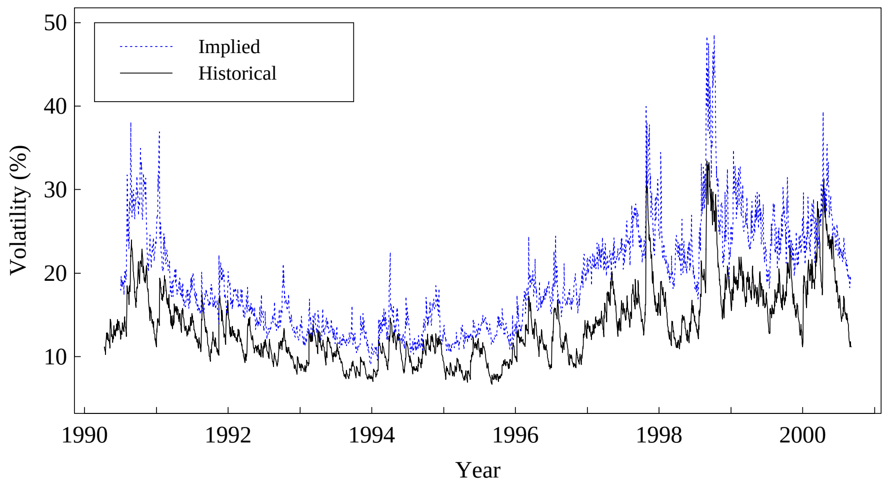

Figure 1 *Implied and historical volatilities for the S&P 500 Index.*

Related literature on combining derivative and asset price information is contained in Chernov and Ghysels (2000) and Johannes and Polson (2002). For option pricing applications, see Nandi and Heston (2000) and Eraker (2002) for S&P 500 options data, Florentina, Leon and Rubin (2002) for the IBEX 35 Spanish index, Lin, Strong and Xu (2002) for the FTSE 100. Forbes, Martin and Wright (2002) provides a simulation study and Lewis (2000) a current review of option pricing methods under stochastic volatility.

The main reason for combining information is to provide more precise parameter estimates. In addition, it provides a diagnostic tool for model misspecification. This is due to the fact that, theoretically, the estimated parameters implicit in derivatives prices should be consistent with those using asset price data alone. This insight also yields a model misspecification diagnostic by examining the consistency of the time series properties of the estimated state variables.

The rest of the paper is outlined as follows. Section 2 defines the basic problem and provides the necessary likelihood functions for combining derivative and asset price information. We explicitly study the leverage stochastic volatility model and show how to perform inference using MCMC methods. Section 3 provides an application to S&P 500 equity index options. Finally, Section 4 concludes.

2. BAYESIAN INFERENCE FOR DERIVATIVE PRICING

This section describes the general problem in three stages. First, we describe the inference problem when both asset and derivative prices are observed. Second, we provide an example of inference for equity index options using a leverage stochastic volatility model. Finally, we describe how to implement an MCMC simulation strategy to provide inference for unobserved state variables and parameters given both asset and derivative prices.

2.1. *The Problem*

To describe the basic problem, consider observations on both asset and derivative prices. Let Y^S denote the asset price returns and let Y^D denote the panel of derivative prices. The dynamics of the asset prices Y^S are determined via a likelihood $p(Y^S \mid X, \Theta)$ where X denotes a state variable and Θ a parameter vector. We assume that these evolve according to a model for the asset price dynamics which we denote by the joint distribution $p(X, \Theta) = p(X \mid \Theta)p(\Theta)$.

Inference using only the asset price information Y^S then leads to a posterior distribution for state and parameters given by

$$p(X, \Theta \mid Y^S) \ \propto \ p(Y^S \mid X, \Theta)\, p(X, \Theta)$$

The corresponding marginal distributions $p(\Theta \mid Y^S)$ and $p(X \mid Y^S)$ provide the usual inferential summaries.

The problem of incorporating the additional derivative price information can be described as follows. First, from a theoretical asset pricing model, the derivative price is determined as a function of the state variable X_t, parameter Θ and current asset price Y_t^S. We denote this function by $F(Y_t^S, X_t, \Theta, \Lambda)$, where Λ denotes additional market price of risk parameters. These parameters are used in derivative pricing, but do not affect the asset price or state dynamics. We assume that the observed derivative price Y_t^D is equal to the theoretical value plus a pricing error:

$$Y_t^D \ = \ F(Y_t^S, X_t, \Theta, \Lambda) \ + \ \epsilon_t^D$$

This leads to a likelihood for the observed derivative prices denoted by $p(Y^D \mid Y^S, \Theta, \Lambda)$.

In order to calculate $F(Y_t^S, X_t, \Theta, \Lambda)$ we need the risk-neutral dynamics of the Y^S process. This is determined by transforming the original state dynamics with the market price of risk parameters. See Duffie *et al.* (2000) and Johannes and Polson (2002) for a number of examples involving equity and interest rate models. More specifically, $F(Y_t^S, X_t, \Theta, \Lambda)$ is given by conditional expectation under $\mathrm{E}^Q(\cdot|X_t, \Theta, \Lambda)$, namely

$$F(Y_t^S, X_t, \Theta, \Lambda) \ = \ E^Q \left(e^{-\int_t^T r_s ds} g(S_T) \mid X_t, \Theta, \Lambda \right)$$

where $g(S_T)$ denotes the payout function for the underlying derivative.

The joint inference problem is then the computation of the joint posterior distribution of state variables and parameters. The parameters include those driving the state evolution Θ and those that are used to determine the risk neutral pricing measure Λ. Our data $Y = (Y^D, Y^S)$ represents a panel of derivative prices in Y^D and the underlying series Y^S and the joint posterior of interest is given by

$$p(X, \Theta, \Lambda \mid Y) \ \propto \ p(Y \mid X, \Theta, \Lambda)\, p(X, \Theta)\, p(\Lambda)$$

where $p(Y \mid X, \Theta, \Lambda)$ denotes the likelihood function, $p(X, \Theta)$ is as described above, and $p(\Lambda)$ denotes a prior on the market price of risk parameters.

The likelihood can be further decomposed as

$$p(Y \mid X, \Theta, \Lambda) = p(Y^D \mid Y^S, X, \Theta, \Lambda)\, p(Y^S \mid X, \Theta)$$

where $p(Y^S \mid X, \Theta)$ describes the evolution of the underlying asset price. We now describe an application to equity index options using a leverage stochastic volatility model.

2.2. *Stochastic Volatility Option Pricing*

A common approach to equity option pricing is to use a leverage stochastic volatility model due to Heston (1993). Suppose that the underlying equity price S_t evolves according to a stochastic volatility model with square-root dynamics for variance V_t, and correlated errors. The logarithmic asset price $Y_t^S = \log S_t$ then satisfies

$$dY_t^S = \mu(V_t)dt + \sqrt{V_t}dW_t^s$$

$$dV_t = \kappa(\theta - V_t)dt + \sigma_v\sqrt{V_t}dW_t^v$$

for some drift $\mu(V_t)$. Here (W_t^s, W_t^v) are a pair of Brownian motions with correlation ρ. This correlation, or so-called leverage effect (Black, 1976), is important to explain the empirical fact that volatility increases faster as equity prices drop. [1] The parameters (κ, θ) govern the speed of mean reversion and the long-run mean of volatility and σ_v measures the volatility of volatility.

Equity option prices Y_t^D are determined as a function of $F(Y_t^S, V_t, \Theta, \Lambda)$. Consider a derivative written on the underlying equity with a payoff $g(S_T)$ at some future time T. Let r_s denote the instantaneous short rate. Using standard arbitrage arguments (Duffie, 2000 and Duffie *et al.* 2000), the price of a call option is given by the conditional expectation

$$F(Y_t^S, V_t, \Theta, \Lambda) = E^Q\left(e^{-r\tau}(S_T - K)_+ \mid V_t, \Theta, \Lambda\right)$$

where the option payout function is $(S_T - K)_+$ at time T and exercise price K; E^Q is taken with respect to the risk-neutral dynamics; and $h(x)_+ = \max\{h(z), 0\}$. The instantaneous short rate r is assumed constant over the time interval $\tau = T - t$. The risk-neutral dynamics are given by the evolution

$$dS_t = r_t S_t dt + \sqrt{V_t}S_t dW_t^s$$

$$dV_t = \kappa^*(\theta^* - V_t)dt + \sigma_v\sqrt{V_t}dW_t^v$$

where

$$\kappa^* = \kappa + \lambda_v \text{ and } \theta^* = \frac{\kappa}{\kappa + \lambda_v}\theta$$

Here λ_v is the market price of volatility risk. Heston (1993) motivates this choice of a market price of volatility risk which is proportional to V_t, namely $\lambda(V_t, t) = \lambda_v V_t$. Typically the market price of risk parameter λ_v is negative which leads to option pricing being implemented with θ transformed higher to $(\kappa/\kappa + \lambda_v)\,\theta$.

[1] While this model has been shown to provide very feasible dynamics for asset prices and for patterns in option prices (Heston, 1993) there is still further empirical evidence for jumps in returns (Bates, 1996, 2000; Bakshi, Cao and Chen, 1997; and Eraker, 2002) as well as volatility (see Eraker, Johannes and Polson, 2002).

One of the advantages of this flexible model is that analogous to Black Scholes there is a closed form pricing equation. Specifically, there exists a pair of probabilities $P_j(V_t, \Theta, \Lambda), j = 1, 2$, such that the call price $F(Y_t^S, V_t, \Theta, \Lambda)$ is given by:

$$F(Y_t^S, V_t, \Theta, \Lambda) \;=\; S_t P_1(V_t, \Theta, \Lambda) - K e^{-r\tau} P_2(V_t, \Theta, \Lambda)$$

Specifically, $P_j(V_t, \Theta, \Lambda) = Pr_j\left(\ln(S_T/K) \geq 0 \,|\, V_t, \Theta, \Lambda\right)$, where these probabilities [2] can be determined by inverting a characteristic function

$$P_j(V_t, \Theta, \Lambda) \;=\; \frac{1}{2} + \frac{1}{\pi} \int_0^\infty Re\left[\frac{e^{-i\phi \ln(K)} f_j(V_t, \Theta, \Lambda)}{i\phi}\right] d\phi$$

where

$$f_j(V_t, \Theta, \Lambda) \;=\; e^{C(\tau;\phi) + D(\tau;\phi)V_t + i\phi \ln(S_t)},$$

and the coefficients $C(\tau;\phi)$ and $D(\tau;\phi)$ are defined in Heston (1993). The correlation and volatility of volatility parameters affect the skewness and kurtosis of the underlying return distribution and Heston (1993) describes how the correlation ρ affects the probabilities $P_j(V_t, \Theta, \Lambda)$. For example, as is typically the case, when the correlation is negative then the underlying risk neutral distribution has a skewed left tail which in turn decreases the price of out-of-the-money call options while increasing the out-of-the-money put options.

2.3. *MCMC Implementation*

This section describes how to sample from the joint distribution of state variables and parameters using MCMC methods. First we describe how to calculate the necessary likelihood functions with and without the use of derivatives prices. Then we illustrate how to use MCMC methods on the stochastic volatility Heston model described in Section 2.2. First, we use an Euler time discretisation of the continuous time model and write the observed daily log returns, $Y_t^S = \log(S_{t+1}/S_t)$, as evolving according to

$$Y_t^S = \mu + \sqrt{V_t}\epsilon_t^s$$
$$V_{t+1} - V_t = \kappa(\theta - V_t) + \sigma_v \sqrt{V_t}\epsilon_t^v.$$

For simplicity we assume that the drift μ is constant and that the unobserved state variables consists of a vector of volatilities $V = (V_0, \ldots, V_T)$ and the parameters are given by $\Theta = (\kappa, \theta, \sigma_v, \rho)$.

Incorporating the derivative price information leads to an observation equation consisting of a panel of N prices for T time periods of the form $Y_{i,t}^D = F_i(Y_t^S, V_t, \Theta, \Lambda) + \sigma_D \epsilon_{i,t}^D$ for $i = 1, \ldots, N$ and $t = 1, \ldots, T$. This gives an extra likelihood term

$$p(Y^D \,|\, Y^S, V, \Theta, \Lambda) \propto {\sigma_D}^{-NT} \exp\left(-\frac{1}{2\sigma_D^2} \sum_{i=1}^{N} \sum_{t=1}^{T} \left(Y_{i,t}^D - F_i(Y_t^S, V_t, \Theta, \Lambda)\right)^2\right).$$

Combining this information with the asset dynamics gives a joint posterior distribution

$$p(X, \Theta, \Lambda \,|\, Y) \;\propto\; p(Y^D \,|\, Y^S, X, \Theta, \Lambda)\; p(Y^S \,|\, X, \Theta)\; p(X, \Theta)\; p(\Lambda)$$

[2] The probabilities P_j are determined under the respective measures $dY_t^S = [r \pm \frac{1}{2}V_t]dt + \sqrt{V_t}dW_T^S$ and $dV_t = (a_j - b_j v_t)dt + \sigma_v\sqrt{V_t}dW_T^V$ where a_j and b_j are defined in Heston, 1993.

Inference is then determined by sampling from the joint posterior distribution of the parameters and volatility states. For example, in the leverage SV model, we sample in a block algorithm defined by the set of conditionals

$$\begin{aligned} \text{Volatility states:} \quad & p(V_t \,|\, V_{-t}, \sigma_D^2, \mu, \lambda_v, \kappa, \theta, \rho, \sigma_v, Y) \\ \text{Stock price drift:} \quad & p(\mu \,|\, V, \sigma_D^2, \lambda_v, \kappa, \theta, \rho, \sigma_v, Y) \\ \text{Pricing variance:} \quad & p(\sigma_D^2 \,|\, V, \mu, \lambda_v, \kappa, \theta, \rho, \sigma_v, Y) \\ \text{Market price of risk:} \quad & p(\lambda_v \,|\, V, \mu, \sigma_D^2, \kappa, \theta, \rho, \sigma_v, Y) \\ \text{Volatility parameters:} \quad & p(\kappa, \theta \,|\, V, \mu, \sigma_D^2, \lambda_v, \rho, \sigma_v, Y) \\ \text{Correlation and volatility of volatility:} \quad & p(\rho, \sigma_v \,|\, V, \mu, \sigma_D^2, \lambda_v, \kappa, \theta, Y) \end{aligned}$$

In constract to the problem of inference with only asset prices, many of these conditionals are unavailable in closed form when option prices are included due to the nonlinear pricing function $F(Y_t^S, V_t, \Theta, \Lambda)$ in the derivative likelihood. We therefore use a Metroplis-Hastings algorithm to sample the parameters that are affected by this, namely $(\lambda_v, \kappa, \theta, \rho, \sigma_v)$, and the volatilities V_t. To update the parameters, we used the block scheme indicated above with a normal proposal distribution centered at the current value with an estimated posterior covariance determined from a short pilot sample and a scale factor chosen to give acceptance probabilities around 30–50%. For the volatilities we used a univariate truncated normal centered at the current value with a variance chosen to yield acceptance rates of around 50%. Finally, the full conditional distributions for μ and σ_D^2 are known in closed form, being normal and inverse-gamma distributions respectively. We found that the algorithm converges fairly rapidly when the pricing error variance is fixed. However, we also found that convergence is much slower when σ_D^2 is unknown. In this case, we used a large Monte Carlo sample size of $M = 100,000$ and we found that this provides reasonable convergence diagnostics.

For our prior specification $p(\Theta, \Lambda)$ we use diffuse priors [3] where appropriate and normal priors for the transformed parameters $(\mu, \kappa, \kappa\theta)$. Inference is performed using the marginal posterior distribution $p(\Theta \,|\, Y)$. We use posterior means and standard deviations throughout as summaries.

3. APPLICATION: S&P 500 EQUITY OPTIONS

To illustrate our methodology we use daily data from the S&P 500 equity index returns and option prices from 1998–2002. The option prices are specified in the form of implied volatilities for different levels of the option Δ, specifically .25, .50, and .75, corresponding to out-of-the-money, at-the-money, and in-the-money options, respectively. For the purpose of illustration, we focus on short-term options with 1 month to expiration, giving a panel of $N = 3$ options at each time period.

For our option pricing model we use the leverage stochastic volatility model described above. The goal of our analysis is to provide more precise parameter estimates and to show how model diagnostics can be provided based on agreement of the posteriors with and without derivative prices.

Table 1 provides posterior estimates and standard deviations (in parentheses) for the mean annualized volatility $E(\sqrt{V_t})$, the speed of mean reversion κ, the volatility of

[3] Specifically, in our application we use the following prior distributions: $\mu \sim \mathrm{N}(0, .001^2)$ for the drift parameter; $\kappa\theta \sim \mathrm{N}(.000004, .00002^2)$ and $\kappa \sim \mathrm{N}(.02, .02^2)$ for the long run volatility mean and speed of mean reversion; $\rho \sim \mathrm{Un}(-1, 1)$ for the leverage effect; $\sigma_v^{-2} \sim \mathrm{Ga}(2, .0001)$ for the volatility of volatility; $\lambda_v \sim \mathrm{N}(0, .03^2)$ for the market price of risk parameter; and $V_0 \sim \mathrm{N}(.0002, .00002^2)$ for the initial volatility.

volatility σ_v, the leverage effect ρ, the drift μ, and the market price of risk parameter λ_v. We consider two scenarios for the pricing errors: first $\sigma_D = 5\%$, representing a typical bid-ask spread; and secondly, we leave σ_D unknown. The posterior mean for the pricing error was 3%. Our results show that there are little differences between these two cases.

Table 1 *Posterior Estimates of Model Parameters With and Without Derivative Information.*

Parameter	Y^S	Y^S, Y^D	Y^S, Y^D
		$\sigma_D = .05$	σ_D unknown
$E(\sqrt{V_t})$	21.60	21.90	21.82
	(1.46)	(1.15)	(0.86)
κ	0.045	0.046	0.055
	(0.012)	(0.005)	(0.005)
σ_v	0.047	0.048	0.055
	(0.004)	(0.002)	(0.001)
ρ	-0.70	-0.78	-0.74
	(0.08)	(0.02)	(0.01)
μ	.0014	0.0011	0.0010
	(.0004)	(.0004)	(0.0004)
λ_v	–	-0.010	-0.017
	–	(0.005)	(0.005)

We computed inference with and without derivative prices. As theory would suggest, our parameter estimates are very similar in both cases. However, the use of the derivative data provides more precise parameter estimates in some cases by a factor of eight. For example, the leverage effect is notoriously hard to estimate with asset price data alone (see Jacquier, Polson and Rossi, 2002). We find that, with equity price data alone, the posterior mean is $E(\rho \mid Y^S) = -0.70$ and the posterior standard deviation is $SD(\rho \mid Y^S) = .08$. With the additional derivative pricing information, we find that the estimates become $E(\rho \mid Y^S, Y^D) = -.74$ and $SD(\rho \mid Y^S, Y^D) = .01$. [4]

Figure 2 shows the posterior histograms for the four parameters $(\kappa, \theta, \sigma_v, \rho)$ with and without derivative price information. For the derivative pricing case we assumed a 5% pricing error. Notice that, as theory would suggest, the posterior means are in close agreement for the two cases, suggesting no model misspecification. As described above, the most significant reduction for the posterior SDs is for the leverage effect ρ.

Second, we provide estimates of the underlying state variables V_t. Comparing the state variables with and without derivative prices provides a useful model diagnostic. If the theoretical model is a true representation of the derivative prices and asset dynamics then these estimates should agree. Differences in sub-periods can be suggestive of periods of market stress or liquidity and credit risk events.

[4] These estimates are also consistent with Nandi's (1998) estimates for the S&P 500 of -0.79. There are a number of other estimates for different datasets and financial series; for example, Bakshi *et al.* (1997) finds an estimate of -0.64 and Eraker (2002) finds a lower estimate of -0.54 due to the inclusion of jumps.

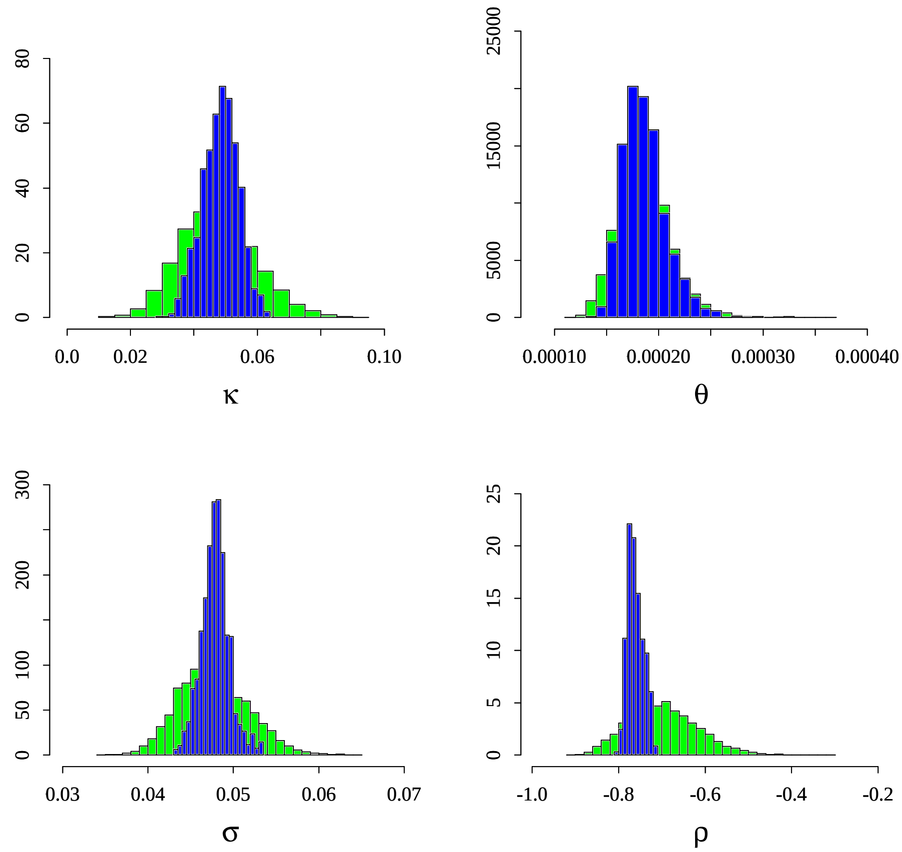

Figure 2 *Posterior histograms of* Θ, *leverage model,* $\sigma_D = .05$. *Without option data (light); with option data (dark).*

To illustrate this, Figure 3 compares the estimated state variables $\hat{V}^S$ and $\hat{V}^D$ with and without the additional derivative price information. Notice that although there is close agreement for most of the period, there are a number of discrepancies. For example, the period from Aug 26–Oct 20, 1998 shows clear differences in the estimated state volatilities. Although this could be due to estimation error, this suggests the need for other state variables such as jumps to explain the external market factors occurring in this period, such as the Russian GKO bond and LTCM crisis.

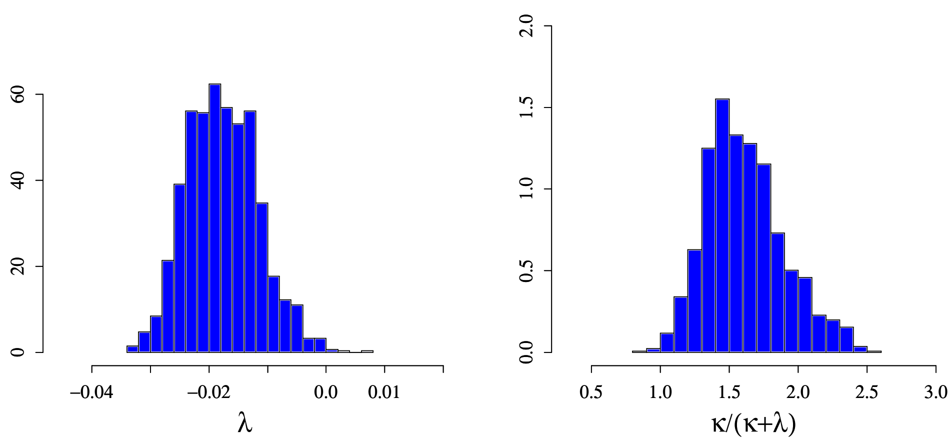

Figure 3 *Posterior histograms of market price of risk parameters, leverage model, unknown* σ_D.

Figure 4 shows the posterior distributions for the market price of volatility risk, $p(\lambda_v \mid Y)$. It also shows the transformed parameter $\kappa/(\kappa+\lambda_v)$, which corresponds to the factor multiplies that multiplies the long-run mean parameter θ in going from the asset price dynamics to the risk-neutral pricing measure. Not surprisingly, there is strong evidence that λ_v is less than zero, and the posterior mean of $\kappa/(\kappa+\lambda_v)$ is equal to 1.65, explaining the empirical fact that the mean level of the implied volatilities is higher than that of the estimated volatilities.

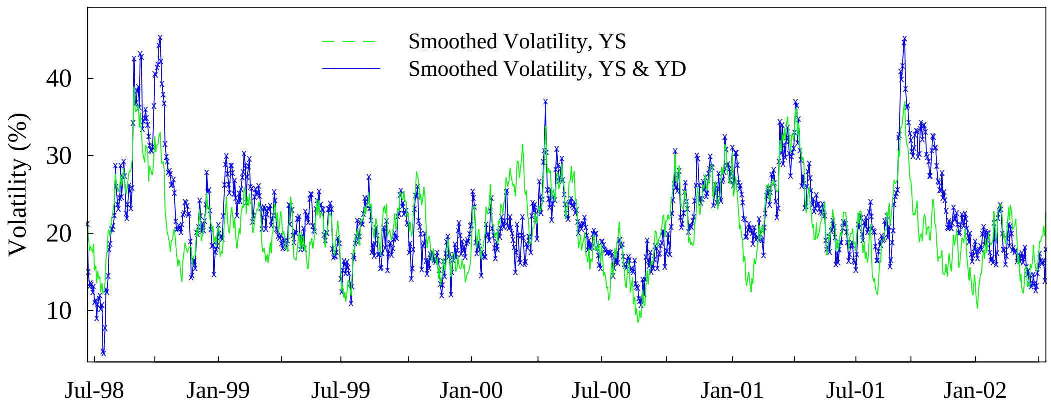

Figure 4 *Annualized estimated volatilities,* $\hat{V}_t^S$ *and* $\hat{V}_t^D$*, for the leverage model with* $\sigma_D = .05$.

To provide an illustration of model misspecification we consider a stochastic volatility model without a leverage effect, *i.e.*, $\rho = 0$. Figure 5 illustrates the parameter posterior distributions $p(\Theta \mid \rho = 0, Y^S)$ and $p(\Theta \mid \rho = 0, Y^S, Y^D)$ under the no leverage model with and without the extra derivative price information. If the model is correctly specified the estimates should be identical, the only difference being the tighter standard errors when the derivative prices are added. Notice the dramatic differences in parameter estimates. This clearly indicates that the model without a leverage effect is misspecified.

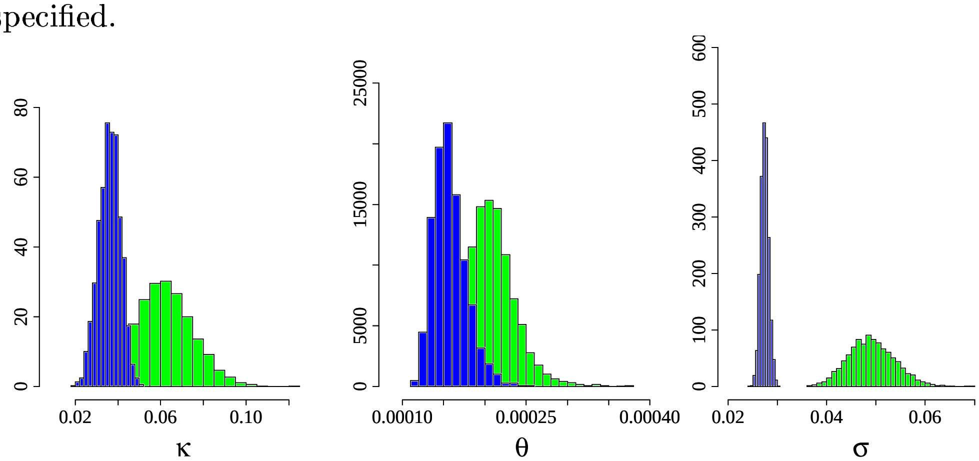

Figure 5 *Posterior histograms of* Θ*, no leverage model,* $\sigma_D = .05$*. Without option data (light); with option data (dark).*

4. DISCUSSION

This paper develops techniques for Bayesian inference that combines information from derivative and asset price data. Combining such information is increasingly common in practice as it provides more precise parameter estimates. It has the additional feature of drawing inference about market price of risk parameters which are useful in derivative pricing. Furthermore, the estimated parameters and state variables can be used to provide model misspecification diagnostics. These can be identified by finding discrepancies in the estimated parameters with and without the additional derivative price information, and also by comparing the time series properties of the estimated state variables.

There are a number of interesting directions for further study. First these methods can clearly be extended to multi-factor models and can be applied in a number of different situations such as term-structure and exchange rates modelling. Secondly, extending this approach to incorporate sequential state and parameter learning is an important direction. One difficulty that arises in practice is that the theoretical asset pricing framework needs to be extended to allow for sequential parameter learning and estimation risk. Recent work in nonlinear filtering methods such as Johannes, Polson and Stroud (2002) and Polson *et al.* (2002) are applicable for sequential state and parameter learning.

REFERENCES

Bakshi, G., Cao, C. and Chen, Z. (1997). Emprirical performance of alternative option pricing models. *Journal of Finance*, **5**, 2003–2049.

Bates, D. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *Review of Financial Studies*, **9**, 69–107.

— (2000). Post-'87 Crash fears in the S&P 500 Futures option market. *Journal of Econometrics*, **94**, 181–238.

Black, F. (1976). Studies of Stock Price Volatility Changes. *1976 Meetings of the Business and Economics Statistics Section, Proceedings of the American Statistical Association*, 177–181.

Chernov, M. and Ghysels, E. (2000). A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purposes of option evaluation. *Journal of Financial Economics*, **56**, 407–458.

Duffie, D. (2000). *Dynamic Asset Pricing Theory.* Princeton, NJ: Princeton University Press.

Duffie, D., Pan, J. and Singleton, K. (2000). Transform anaplysis and asset pricing for affine jump-diffusions. *Econometrica*, **68**, 1343–1376.

Eraker, B. (2002). Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *Tech. Rep.*, Department of Economics, Duke University.

Eraker, B., Johannes, M. and Polson, N. (2002). The impact of jumps in returns and volatility. In press, *Journal of Finance.*

Florentini, G., Leon, A. and Rubin, G. (2002). Estimation and empirical performance of Heston's stochastic volatility model: the case of a thinly traded market. *Journal of Empirical Finance*, **9**, 225–255.

Forbes, C., Martin, G. and Wright, J. (2002). Bayesian estimation of a stochastic volatility model using option and spot prices. *Tech. Rep.*, Department of Econometrics and Business Statistics, Monash University.

Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, **6**, 327–343.

Heston, S. and Nandi, S. (2000). A closed-form GARCH option pricing valuation model. *Review of Financial Studies*, **13**, 585–625.

Jacquier, E., Polson, N. and Rossi, P. (2002). Bayesian analysis of stochastic vol models with fat tails and correlated errors. To appear in *Journal of Econometrics.*

Johannes, M. and Polson, N. (2002). MCMC methods for financial econometrics. To appear in *Handbook of Financial Econometrics* (eds. Y. Aït-Sahalia and L. Hansen). Princeton University Press.

Johannes, M., Polson, N. and Stroud, J. (2002). Nonlinear filtering for Stochastic Differential Equations with Jumps. *Tech. Rep.*, Graduate School of Business, University of Chicago.

Lewis, A. (2000). *Option Valuation under Stochastic Volatility.* Financial Press.

Lin, N. Y.-N., Strong, N. and Xu, X. (2002). Pricing FTSE 100 options under stochastic volatility. *Journal of Futures Markets*, **21**, 127–211.

Nandi, S. (1998). How important is the correlation between assets and volatility in a stochastic volatility model? Empirical evidence from pricing and hedging in the S&P index options market. *Journal of Banking and Finance*, **22**, 589–610.

Polson, N., Stroud, J. and Müller, P. (2002). Practical filtering with sequential parameter learning. *Tech. Rep.*, Graduate School of Business, University of Chicago.