

Empirical Performance of Alternative Option Pricing Models

GURDIP BAKSHI, CHARLES CAO, and ZHIWU CHEN*

ABSTRACT

Substantial progress has been made in developing more realistic option pricing models. Empirically, however, it is not known whether and by how much each generalization improves option pricing and hedging. We fill this gap by first deriving an option model that allows volatility, interest rates and jumps to be stochastic. Using S&P 500 options, we examine several alternative models from three perspectives: (1) internal consistency of implied parameters/volatility with relevant time-series data, (2) out-of-sample pricing, and (3) hedging. Overall, incorporating stochastic volatility and jumps is important for pricing and internal consistency. But for hedging, modeling stochastic volatility alone yields the best performance.

IN THE LAST TWO DECADES, option pricing has witnessed an explosion of new models that each relax some of the restrictive Black-Scholes (BS) (1973) assumptions. Examples include (i) the stochastic-interest-rate option models of Merton (1973) and Amin and Jarrow (1992); (ii) the jump-diffusion/pure jump models of Bates (1991), Madan and Chang (1996), and Merton (1976); (iii) the constant-elasticity-of-variance model of Cox and Ross (1976); (iv) the Markovian models of Rubinstein (1994) and Ait-Sahalia and Lo (1996); (v) the stochastic-volatility models of Heston (1993), Hull and White (1987a), Melino and Turnbull (1990, 1995), Scott (1987), Stein and Stein (1991), and Wiggins (1987); (vi) the stochastic-volatility and stochastic-interest-rates models of Amin and Ng (1993), Bailey and Stulz (1989), Bakshi and Chen (1997a,b), and Scott (1997); and (vii) the stochastic-volatility jump-diffusion models of Bates (1996a,c), and Scott (1997). This list is by no means exhaustive, yet already overwhelming to anyone who has to choose among the alternatives. To make matters worse, the number of possible option pricing models is virtually infinite. Note that every option pricing model has to make three basic assump-

* Bakshi is at the University of Maryland, College Park. Cao is at Pennsylvania State University, University Park. Chen is at The Ohio State University, Columbus. This paper subsumes the previous one under the title "Option Pricing and Hedging Performance with Stochastic Volatility and Stochastic Interest Rates." We thank Sanjiv Das, Ranjan D'Mello, Jin-Chuan Duan, Helyette Geman, Eric Ghysels, Frank Hatheway, Steward Hodges, Ravi Jagannathan, Andrew Karolyi, Bill Kracaw, Dilip Madan, Victor Ng, Louis Scott, René Stulz, Stephen Taylor, Siegfried Trautmann, Alex Triantis, Alan White, and the anonymous referee. We gratefully acknowledge comments by seminar participants at the 1996 European Finance Association Meetings, 1997 Western Finance Meetings, the Chinese University of Hong Kong, the Hong Kong University of Science and Technology, The Ohio State University, and the University of New Orleans. Any remaining errors are our responsibility alone.

tions: the underlying price process (the distributional assumption), the interest rate process, and the market price of factor risks. For each of the assumptions, there are many possible choices. For instance, the underlying price can follow either a continuous-time or a discrete-time process. Among possible continuous-time processes, it can be Markov or non-Markov, a diffusion or a nondiffusion, a Poisson or a non-Poisson jump process, a mixture of jump and diffusion components with or without stochastic volatility and with or without random jumps. For the term structure of interest rates, there are similarly many choices. While the search for that perfect option pricing model can be endless, we are tempted to ask: What do we gain from each generalized feature? Is the gain, if any, from a more realistic feature worth the additional complexity or implementational costs? Can any of the relaxed assumptions help resolve known empirical biases associated with the Black-Scholes formula, such as the volatility smiles (e.g., Rubinstein (1985, 1994))? As a practical matter, that perfectly specified option pricing model is bound to be too complex for applications. Ultimately, it is a choice among misspecified models, made perhaps based on (i) "which is the least misspecified?" (ii) "which results in the lowest pricing errors?" and (iii) "which achieves the best hedging performance?" These empirical questions must be answered before the potential of recent advances in theory can be fully realized in practical applications.

The purpose of the present article is to fill in this gap and conduct a comprehensive empirical study on the relative merits of competing option pricing models.¹ To this goal, we first develop in closed form an implementable option pricing model that admits stochastic volatility, stochastic interest rates, and random jumps, which will be abbreviated as the *SVSI-J model*. The setup is rich enough to contain almost all the known closed form option formulas as special cases, including (i) the Black-Scholes (BS) model, (ii) the stochastic-interest-rate (SI) model, (iii) the stochastic-volatility (SV) model, (iv) the stochastic-volatility and stochastic-interest-rate (SVSI) model, and (v) the stochastic-volatility random-jump (SVJ) model. The constant-volatility jump-diffusion models of Bates (1991) and Merton (1976) are special cases of the SVJ. Consequently, we concentrate our efforts on the SVSI-J and the five models just described.

Besides the obvious normative reasons, a common motivation for these new models is the abundant empirical evidence that the benchmark BS formula exhibits strong pricing biases across both moneyness and maturity (i.e., the "smile") and that it especially underprices deep out-of-the-money puts and calls (see Bates (1996b) for an insightful review). Such evidence is clearly indicative of implicit stock return distributions that are negatively skewed with higher kurtosis than allowable in a BS log-normal distribution. Guided by this implication, the

¹ A few existing studies investigate the internal consistency of implied parameters (Bates (1991, 1996a,c)), and the pricing or the hedging performance (e.g., Bakshi, Cao, and Chen (1997), Cao (1993), Dumas, Fleming, and Whaley (1995), Madan and Chang (1996), Nandi (1996), and Rubinstein (1985)), of alternative stochastic-volatility models. Cao studies a stochastic-volatility model using currency options; Nandi investigates the pricing and single-instrument-hedging performance using the S&P 500 futures. In this article we address the empirical issues from different perspectives and under alternative models.

search for alternative models has mostly focused on finding the “right” distributional assumption. The SV model, for instance, offers a flexible distributional structure in which the correlation between volatility shocks and underlying stock returns serves to control the level of skewness and the volatility variation coefficient serves to control the level of kurtosis. But, since volatility in the SV is modeled as a diffusion and hence only allowed to follow a continuous sample path, its ability to internalize enough short-term kurtosis and thus to price short-term options properly is limited (unless the variation coefficient of spot volatility is unreasonably high). The jump-diffusion models, on the other hand, assert that it is the occasional, discontinuous jumps and crashes that cause the negative implicit skewness and high implicit kurtosis to exist in option prices. The fact that such jumps and crashes are allowed to be discontinuous over time makes these models more flexible than the diffusion-stochastic-volatility model, in internalizing the desired return distributions, especially at short time horizons. Therefore, the random-jump and the stochastic-volatility features can in principle improve the pricing and hedging of, respectively, short-term and relatively long-term options. The inclusion of a stochastic term structure model in an option pricing framework is, however, intended to improve the valuation and discounting of future payoffs, rather than to enhance the flexibility of permissible return distributions. Thus, while the stochastic-interest-rate feature is not expected to help resolve the cross-sectional pricing biases, it should in principle improve the pricing fit across option maturity.

We implement every model by backing out, on each day, the spot volatility and structural parameters from the observed option prices of that day. This approach is common in the existing literature (e.g., Bates (1996b)), partly out of the consideration that historical data reflect what happened in the past whereas information implicit in option prices is forward-looking. Backing out the BS model’s volatility and other model’s parameters daily is indeed ad hoc since volatility in the BS and the structural parameters in the other models are assumed to be constant over time. But, as this internally inconsistent treatment is how each model is to be applied, we follow this convention so as to ensure each model an equal chance.

In judging the alternative models, we employ three yardsticks. First, are the implied structural parameters consistent with those implicit in the relevant times-series data (e.g., the implied-volatility time series, and the interest-rate time series)? Much of this part of the discussion is based on Bates’ (1996a,c) work where he studies the relative desirability of the SV versus the SVJ models, using stock index futures and currency options. The reasoning is that if an option model is correctly specified, its structural parameters implied by option prices will necessarily be consistent with those implicit in the observed time-series data. Second, out-of-sample pricing errors give a direct measure of model misspecification. In particular, while a more complex model will generally lead to better in-sample fit, it will not necessarily perform better out of sample as any overfitting may be penalized. Third, hedging errors measure how well a model captures the dynamic properties of option and underlying security prices. In other words, in-sample and out-of-sample pricing errors reflect a model’s *static* performance,

while hedging errors reflect the model's *dynamic* performance. As shown later, these three yardsticks serve distinct purposes.

Based on 38,749 S&P 500 call option prices from June 1988 to May 1991, we find that the SI and the SVSI-J models do not significantly improve the performance of the BS and the SVJ models, respectively. To keep the presentation manageable, we focus on the four models of distinct interest: the BS, the SV, the SVSI, and the SVJ. Our empirical investigation leads to the following overall conclusions. First, judged on internal parameter consistency, all models are misspecified, with the SVJ the least and the BS the most misspecified. This conclusion is confirmed from several different angles. For example, according to the Rubinstein (1985) type of implied-volatility graphs, the SVJ implied volatility smiles the least across moneyness levels, followed in increasing order by the SVSI, the SV, and the BS. Second, out-of-sample pricing errors are the highest for the BS, the second highest for the SV, and the lowest for the SVJ. Overall, stochastic volatility alone achieves the first-order pricing improvement and typically reduces the BS pricing errors by 25 percent to 60 percent. However, our evidence also confirms the conjectures that (i) adding the random-jump feature improves the fit of short-term options and that (ii) including the SI feature enhances the pricing fit of long-term options. After both stochastic volatility and random jumps are modeled, the remaining pricing errors no longer exhibit clear systematic biases (e.g., across moneyness).

Two types of hedging strategy are employed to gauge the relative hedging effectiveness. First, we examine minimum-variance hedges of option contracts that rely on the underlying asset as the single hedging instrument. As argued by Ross (1995), the need for this type of hedge may arise in contexts where a perfect delta-neutral hedge may not be feasible, either because of untraded risks or because of model misspecifications and transaction costs. In the presence of more than one source of risk, single-instrument hedges can only be partial. According to results from these type of hedges, the SV outperforms all the others, while the SVJ is second. Between the other two models, the BS hedges in-the-money calls better than the SVSI, but the SVSI is better in hedging out-of-the-money calls. This hedging result is surprising as one would expect the SVSI to perform at least as well as the BS, and the SVJ to do better than the SV.

Next, we implement a conventional delta-neutral hedge, in which as many hedging instruments as there are risk sources are used to make the net position completely risk-immunized (locally). For the case of the BS, this means that only the underlying stock will be employed to hedge a call. For the SV model, however, both the price risk and volatility risk affect the value of a call, implying that an SV-based delta-neutral hedge will need a position in the underlying stock and one in a second option contract. For the SVSI, its delta-neutral hedge will involve a discount bond (to control for interest rate risk) in addition to the underlying stock and a second option contract. When such internally consistent hedges are implemented, the hedging errors for the SV, the SVSI and the SVJ are about 50 percent to 65 percent lower than those of the BS model, if each hedge is rebalanced daily. Furthermore, changing the hedge rebalancing frequency affects the BS model's hedging errors dramatically, while only affecting the other models' performance

marginally. That is, after stochastic volatility is controlled for, the errors of a delta-neutral hedge seem to be relatively insensitive to revision frequency.² However, like in the single-instrument hedging case, once stochastic volatility is modeled, adding the SI or the random-jump feature does not enhance hedging performance any further.

Since the delta-neutral hedge for the BS does not use a second option contract whereas it does for the other three models, this may have biased the delta-neutral hedging results against the BS model. To examine this point, we also implement the ad hoc BS delta-plus-vega neutral strategy in which the underlying stock and an option contract are used to neutralize both delta risk and vega risk (of the BS model). It turns out that in hedging out-of-the-money and at-the-money calls, this BS delta-plus-vega neutral strategy performs no worse than the other models' delta-neutral hedges. Only in hedging deep in-the-money calls do the stochastic volatility models perform better than the BS delta-plus-vega neutral strategy. This is true regardless of hedge revision frequency. Overall, hedging performance is relatively insensitive to model misspecification, since even ad hoc hedges can result in similar errors.

The rest of the article proceeds as follows. Section I develops the option pricing models. Section II provides a description of the S&P 500 option data. In Section III we present an estimation procedure, discuss the estimated parameters, and evaluate the in-sample fit of each model. Section IV assesses the extent of each model's misspecification. Sections V and VI, respectively, present the out-of-sample pricing and the hedging results. Concluding remarks are offered in Section VII. Proof of pricing equations and most formulas are provided in the Appendix.

I. Option Pricing Models

The purpose of this section is to derive a closed-form jump-diffusion option pricing model that includes all those to be studied in the empirical sections as special cases. As such, it is then convenient to follow a standard practice and specify from the outset a stochastic structure under a risk-neutral probability measure. The existence of this measure is equivalent to the absence of free lunches, and it allows us to value future risky payoffs as if the economy were risk-neutral. First, under the risk-neutral measure, the underlying nondividend-paying stock price $S(t)$ and its components are, for any t , given by

$$\frac{dS(t)}{S(t)} = [R(t) - \lambda\mu_J]dt + \sqrt{V(t)}d\omega_S(t) + J(t)dq(t) \quad (1)$$

$$dV(t) = [\theta_v - \kappa_v V(t)]dt + \sigma_v \sqrt{V(t)}d\omega_v(t) \quad (2)$$

² This finding is in accord with Galai's (1983) results that in any hedging scheme it is probably more important to control for stochastic volatility than for discrete hedging (see Hull and White (1987b) for a similar, simulation-based result for currency options).

$$\ln[1 + J(t)] \sim N(\ln[1 + \mu_J] - 1/2 \sigma_J^2, \sigma_J^2), \quad (3)$$

where:

$R(t)$ is the time- t instantaneous spot interest rate;
 λ is the frequency of jumps per year;
 $V(t)$ is the diffusion component of return variance (conditional on no jump occurring);
 $\omega_S(t)$ and $\omega_v(t)$ are each a standard Brownian motion, with
 $\text{Cov}_t[d\omega_S(t), d\omega_v(t)] \equiv \rho dt$;
 $J(t)$ is the percentage jump size (conditional on a jump occurring) that is lognormally, identically, and independently distributed over time, with unconditional mean μ_J . The standard deviation of $\ln[1 + J(t)]$ is σ_J ;³
 $q(t)$ is a Poisson jump counter with intensity λ , that is, $\Pr\{dq(t) = 1\} = \lambda dt$ and $\Pr\{dq(t) = 0\} = 1 - \lambda dt$;
 κ_v , θ_v/κ_v , and σ_v are respectively the speed of adjustment, long-run mean, and variation coefficient of the diffusion volatility $V(t)$;
 $q(t)$ and $J(t)$ are uncorrelated with each other or with $\omega_S(t)$ and $\omega_v(t)$.

Under the assumed framework, the total return variance can be decomposed into two components:

$$\frac{1}{dt} \text{Var}_t \left(\frac{dS(t)}{S(t)} \right) = V(t) + V_J(t), \quad (4)$$

where $V_J(t) \equiv (1/dt) \text{Var}_t[J(t)dq(t)] = \lambda[\mu_J^2 + (e^{\sigma_J^2} - 1)(1 + \mu_J)^2]$ is the instantaneous variance of the jump component.

This stock-return distributional assumption, similar to the one in Bates (1996a) for currency prices, offers a sufficiently versatile structure that can accommodate most of the desired features. For instance, skewness in the distribution is controlled by either the correlation ρ or the mean jump μ_J , whereas the amount of kurtosis is regulated by either the volatility diffusion parameter σ_v or the magnitude and variability of the jump component. But the ability of the diffusion component $V(t)$ to generate enough short-run negative skewness or excess kurtosis is limited, as $V(t)$ can only follow a continuous sample path. On the other hand, the discontinuous jump process can internalize any skewness and kurtosis level even in the short run, especially when λ , μ_J , and σ_J are substantial. Therefore, these two forces capture different aspects of return distributions.

Next, to ensure proper discounting of future cash flows, we adopt a single-factor term structure model of the Cox, Ingersoll, and Ross (1985) type as it requires the estimation of only three structural parameters:

$$dR(t) = [\theta_R - \kappa_R R(t)]dt + \sigma_R \sqrt{R(t)} d\omega_R(t), \quad (5)$$

³ See, for example, Bates (1996a,c), Merton (1976), and Scott (1997) for a similar jump setup.

where κ_R , θ_R/κ_R , and σ_R are respectively the speed of adjustment, long-run mean, and volatility coefficient of the $R(t)$ process; and $\omega_R(t)$ is a standard Brownian motion, uncorrelated with any other process in the economy.⁴ Of course, we can add more factors to the term structure model and make the resulting bond price formulas more plausible, but that will also make the option pricing formula harder to implement.

It is important to realize that the exogenous valuation framework given above can be derived from a general equilibrium in which the volatility risk $V(t)$, interest rate risk $R(t)$, and jump risk $J(t)dq(t)$ are all rewarded. For instance, Bakshi and Chen (1997a) and Bates (1996a,c) provide such examples in which each risk factor earns a risk premium proportional to the factor itself. That is, the factor prices for $V(t)$ and $R(t)$ are respectively $b_v V(t)$ and $b_r R(t)$, for some constants b_v and b_r . These factor prices are implicitly reflected in equations (2) and (5) and adjusted through κ_v and κ_R , respectively. Therefore, factor risk premiums are not assumed to be zero in our framework. Rather, they have been internalized in the stochastic structure.

Consider first a zero-coupon bond that pays \$1 in τ periods from time t , and let $B(t, \tau)$ be its current price. Then,

$$B(t, \tau) = \mathbb{E}_Q \left\{ \exp \left(- \int_t^{t+\tau} R(u) du \right) \right\} = \exp[-\varphi(\tau) - \varrho(\tau)R(t)], \quad (6)$$

where

$$\varphi(\tau) = \frac{\theta_R}{\sigma_R^2} \left\{ (s - \kappa_R)\tau + 2 \ln \left[1 - \frac{(1 - e^{-s\tau})(s - \kappa_R)}{2s} \right] \right\},$$

$$\varrho(\tau) = \frac{2(1 - e^{-s\tau})}{2s - [s - \kappa_R](1 - e^{-s\tau})}, \quad s \equiv \sqrt{\kappa_R^2 + 2\sigma_R^2},$$

⁴ This assumption on the correlation between stock returns and interest rates is somewhat severe and likely counterfactual. To gauge the potential impact of this assumption on the resulting option model's performance, we initially adopt the following stock price dynamics:

$$\frac{dS(t)}{S(t)} = \mu(S, t)dt + \sqrt{V(t)}d\omega_S(t) + \sigma_{S,R}\sqrt{R(t)}d\omega_R(t),$$

with the rest of the stochastic structure remaining the same as given above. Under this more realistic structure, the covariance between stock price changes and interest rate shocks is $\text{Cov}_t[dS(t), dR(t)] = \sigma_{S,R}\sigma_R R(t)S(t)dt$, so bond market innovations can be transmitted to the stock market and vice versa. The obtained closed-form option pricing formula under this scenario would have one more parameter $\sigma_{S,R}$ than the one presented shortly, but when we implement this slightly more general model, we find its pricing and hedging performance to be indistinguishable from that of the SVSI model studied in this article. For this reason, we choose to set $\sigma_{S,R} = 0$. We could also make both the drift and the diffusion terms of $V(t)$ a linear function of $R(t)$ and $\omega_R(t)$. In such cases, the stock returns, volatility and interest rates would all be correlated with each other (at least globally), and we could still derive the desired equity option valuation formula. But, that would again make the resulting formula more complex while not improving its performance.

and $E_Q(\cdot)$ is the expectations operator with respect to the risk-neutral measure.

For a European call option written on the stock with strike price K and term-to-expiration τ , its time- t price $C(t, \tau)$ must, by a standard argument, solve

$$\begin{aligned} \frac{1}{2} VS^2 \frac{\partial^2 C}{\partial S^2} + [R - \lambda \mu_J] S \frac{\partial C}{\partial S} + \rho \sigma_V VS \frac{\partial^2 C}{\partial S \partial V} + \frac{1}{2} \sigma_V^2 V \frac{\partial^2 C}{\partial V^2} + [\theta_V - \kappa_V V] \frac{\partial C}{\partial V} \\ + \frac{1}{2} \sigma_R^2 R \frac{\partial^2 C}{\partial R^2} + [\theta_R - \kappa_R R] \frac{\partial C}{\partial R} - \frac{\partial C}{\partial \tau} - RC \\ + \lambda E\{C(t, \tau; S(1+J), R, V) - C(t, \tau; S, R, V)\} = 0. \end{aligned} \quad (7)$$

subject to $C(t + \tau, 0) = \max\{S(t + \tau) - K, 0\}$. In the Appendix it is shown that

$$C(t, \tau) = S(t) \Pi_1(t, \tau; S, R, V) - KB(t, \tau) \Pi_2(t, \tau; S, R, V), \quad (8)$$

where the risk-neutral probabilities, Π_1 and Π_2 , are recovered from inverting the respective characteristic functions (see Bates (1996a,c), Heston (1993), and Scott (1997) for similar treatments):

$$\begin{aligned} \Pi_j(t, \tau; S(t), R(t), V(t)) \\ = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty Re \left[\frac{e^{-i\phi \ln[K]} f_j(t, \tau, S(t), R(t), V(t); \phi)}{i\phi} \right] d\phi, \end{aligned} \quad (9)$$

for $j = 1, 2$, with the characteristic functions f_j respectively given in equations (A10) and (A11) of the Appendix. The price of a European put on the same stock can be determined from the put-call parity.

The option valuation model in equation (8) has several distinctive features. First, it applies to economies with stochastic interest rates, stochastic volatility, and jump risk. It contains most existing models as special cases. For example, we obtain (i) the BS model by setting $\lambda = 0$ and $\theta_R = \kappa_R = \sigma_R = \theta_V = \kappa_V = \sigma_V = 0$; (ii) the SI model by setting $\lambda = 0$ and $\theta_V = \kappa_V = \sigma_V = 0$; (iii) the SV model by setting $\lambda = 0$ and $\theta_R = \kappa_R = \sigma_R = 0$; (iv) the SVSI model by setting $\lambda = 0$; and (v) the SVJ model by letting $\theta_R = \kappa_R = \sigma_R = 0$, where to derive each special case from equation (8) one may need to apply L'Hopital's rule. The Appendix provides the exact option pricing formulas respectively for the SV, the SVSI, and the SVJ models. Second, this general model allows for a flexible correlation structure between the stock return and its volatility, as opposed to the perfect correlation assumed in, for instance, Heston (1993). Third, when compared to the model in Scott (1997), the formula in equation (8) is parsimonious in the number of parameters; especially since it is given only as a function of identifiable variables such that all parameters can be estimated.

The pricing formula in equation (8) applies to European equity options. But in reality most option contracts are American in nature. While it is beyond the scope of the present article to derive a model for American options, it is

nevertheless possible to capture the first-order effect of early exercise in the following manner. For options with early exercise potential, compute the Barone-Adesi and Whaley (1987) early-exercise premium, treating it as if the stock volatility and the yield-curve were time-invariant. Adding this early-exercise adjustment component to the European option price in equation (8) should result in a reasonable approximation of the corresponding American option price (e.g., Bates (1996a)). Alternatively, one can follow such a nonparametric approach as in Ait-Sahalia and Lo (1996) and Broadie, Detemple, Ghysels, and Torres (1996) to price American options.

The closed-form option pricing formula in equation (8) makes it possible to derive comparative statics and hedge ratios analytically. In the present context, there are three sources of stochastic variations over time, price risk $S(t)$, volatility risk $V(t)$ and interest rate risk $R(t)$. Consequently, there are three deltas:

$$\Delta_S(t, \tau; K) \equiv \frac{\partial C(t, \tau)}{\partial S} = \Pi_1 \geq 0 \quad (10)$$

$$\Delta_V(t, \tau; K) \equiv \frac{\partial C(t, \tau)}{\partial V} = S(t) \frac{\partial \Pi_1}{\partial V} - KB(t, \tau) \frac{\partial \Pi_2}{\partial V} \quad (11)$$

$$\Delta_R(t, \tau; K) \equiv \frac{\partial C(t, \tau)}{\partial R} = S(t) \frac{\partial \Pi_1}{\partial R} - KB(t, \tau) \left\{ \frac{\partial \Pi_2}{\partial R} - \rho(\tau) \Pi_2 \right\}, \quad (12)$$

where, for $g = V, R$ and $j = 1, 2$,

$$\frac{\partial \Pi_j}{\partial g} = \frac{1}{\pi} \int_0^\infty Re \left[(i\phi)^{-1} e^{-i\phi \ln[K]} \frac{\partial f_j}{\partial g} \right] d\phi. \quad (13)$$

The second-order partial derivatives with respect to these variables are provided in the Appendix. These analytical expressions for the deltas form a convenient basis for constructing hedges such as the ones to be analyzed shortly.

II. Data Description

Based on the following considerations, we use S&P 500 call option prices for our empirical work. First, options written on this index are the most actively traded European-style contracts. Second, the daily dividend distributions are available for the index (from the S&P 500 Information Bulletin). Furthermore, S&P 500 options and options on S&P 500 futures have been the focus of many existing investigations including, among others, Bakshi, Cao, and Chen (1997), Bates (1996c), Dumas, Fleming, and Whaley (1995), Madan and Chang (1996), Nandi (1996), and Rubinstein (1994). Finally, we also use S&P 500 puts to estimate the pricing and hedging errors of all the models and find the results to be qualitatively similar. To save space, we only report the results based on the calls.

The sample period extends from June 1, 1988 through May 31, 1991. The intraday bid-ask quotes for S&P 500 options are obtained from the Berkeley Option Database.⁵ To ease computational burden, for each day in the sample, only the last reported bid-ask quote (prior to 3:00 PM Central Standard Time) of each option contract is employed in the empirical tests. Note that the recorded S&P 500 index values are not the daily closing index levels. Rather, they are the corresponding index levels at the moment when the option bid-ask quote is recorded. Thus, there is no nonsynchronous price issue here, except that the S&P 500 index level itself may contain stale component stock prices at each point in time.

The data on the daily Treasury-bill bid and ask discounts with maturities up to one year are hand-collected from the *Wall Street Journal* and provided to us by Hyuk Choe and Steve Freund. By convention, the average of the bid and ask Treasury bill discounts is used and converted to an annualized interest rate. Since Treasury bills mature on Thursdays while index options expire on the third Friday of the month, we utilize the two Treasury-bill rates straddling an option's expiration date to obtain the interest rate corresponding to the option's maturity. This is done for each contract and each day in the sample. The 30-day Treasury bill rate is the surrogate for the short rate in equation (5).

For European options, the spot stock price must be adjusted for discrete dividends. For each option contract with τ periods to expiration from time t , we first obtain the present value of the daily dividends $D(t)$ by computing

$$\bar{D}(t, \tau) \equiv \sum_{s=1}^{\tau-t} e^{-R(t,s)s} D(t+s), \quad (14)$$

where $R(t, s)$ is the s -period yield-to-maturity. In the next step, we subtract the present value of future dividends from the time- t index level, in order to obtain the dividend-exclusive S&P 500 spot index series that is later used as input into the option models. This procedure is repeated for all option maturities and for each day in our sample.

Several exclusion filters are applied to construct the option bid-ask price data. First, option price quotes that are time-stamped later than 3:00 PM Central Standard Time are eliminated. This ensures that the spot price is recorded synchronously with its option counterpart. Second, as options with less than six days to expiration may induce liquidity-related biases, they are excluded from the sample. Third, to mitigate the impact of price discreteness on option valuation, price quotes lower than $\frac{3}{8}$ are not included. Finally, quotes not satisfying the arbitrage restriction

$$C(t, \tau) \geq \max(0, S(t) - K, S(t) - \bar{D}(t, \tau) - KB(t, \tau)) \quad (15)$$

⁵ Early in the project we used only option transaction price data for the empirical work, but, that data set is much smaller, especially for the hedging exercise. Nonetheless, the results based on the transaction prices are similar to those based on mid-point bid-ask quotes.

Table I
Sample Properties of S&P 500 Index Options

The reported numbers are respectively the average quoted bid-ask mid-point price, the average effective bid-ask spread (ask price minus the bid-ask mid-point) which are shown in parentheses, and the total number of observations (in braces), for each moneyness-maturity category. The sample period extends from June 1, 1988 through May 31, 1991 for a total of 38,749 calls. Daily information from the last quote (prior to 3:00 p.m. CST) of each option contract is used to obtain the summary statistics. S denotes the spot S&P 500 index level and K is the exercise price. OTM, ATM, and ITM denote out-of-the money, at-the-money, and in-the-money options, respectively.

	Moneyness S/K	Days-to-Expiration			Subtotal
		<60	60–180	≥ 180	
OTM	<0.94	\$1.68 (0.06) {542}	\$4.38 (0.16) {2330}	\$8.58 (0.26) {1847}	{4719}
	0.94–0.97	\$2.35 (0.09) {1943}	\$8.02 (0.23) {1966}	\$15.12 (0.35) {965}	
ATM	0.97–1.00	\$4.83 (0.15) {2703}	\$12.79 (0.29) {1919}	\$20.17 (0.40) {1003}	{5625}
	1.00–1.03	\$10.42 (0.23) {2543}	\$18.72 (0.35) {1793}	\$26.44 (0.45) {931}	
ITM	1.03–1.06	\$17.77 (0.30) {2255}	\$25.52 (0.41) {1566}	\$33.00 (0.50) {738}	{4559}
	≥ 1.06	\$39.40 (0.37) {5137}	\$48.06 (0.46) {5269}	\$58.12 (0.50) {3299}	
Subtotal		{15123}	{14843}	{8783}	{38749}

are taken out of the sample. Based on this criterion, 624 observations (approximately 1.3 percent of the original sample) are eliminated and those calls are all deep in-the-money.

We divide the option data into several categories according to either moneyness or term to expiration. Define $S(t) - K$ as the time- t intrinsic value of a call. A call option is then said to be *at-the-money* (ATM) if its $S/K \in (0.97, 1.03)$; *out-of-the-money* (OTM) if $S/K \leq 0.97$; and *in-the-money* (ITM) if $S/K \geq 1.03$. A finer partition resulted in six moneyness categories. By the term to expiration, an option contract can be classified as (i) short-term (<60 days); (ii) medium-term (60–180 days); and (iii) long-term (>180 days). The proposed moneyness and maturity classifications produce 18 categories for which the empirical results will be reported.

Table I describes certain sample properties of the S&P 500 call prices used in the study. Summary statistics are reported for the average bid-ask mid-point price, the average effective bid-ask spread (i.e., the ask price minus the bid-ask midpoint), and the total number of observations, for each moneyness-

maturity category. Note that there are a total of 38,749 call option observations, with ITM and ATM options respectively taking up 47 percent and 28 percent of the total sample, and that the average call price ranges from \$1.68 for short-term, deep OTM options to \$58.12 for long-term, deep ITM calls. The effective bid-ask spread varies from \$0.06 (for short-term deep OTM options) to \$0.50 (for long-term deep ITM options).

III. Structural Parameter Estimation and In-Sample Performance

For the empirical work to follow, we concentrate on the four models: the BS, the SV, the SVSI, and the SVJ.⁶ As stated before, the analysis is intended to present a complete picture of what each generalization of the benchmark BS model can really buy in terms of performance improvement and whether each generalization produces a worthy tradeoff between benefits and costs.

To get a sense of what we should look for in any desirable alternative to the BS model, let us use the described data set to examine the extent and the direction of biases associated with the BS. To do this, we back out a BS implied volatility from each option price in the sample. Then, we equally weigh the implied volatilities of all call options in a given moneyness-maturity category, to produce an average implied volatility. The calculations are similarly done for put options. Table II reports the average BS implied-volatility values across six moneyness and three maturity categories, for both calls and puts as well as for both the entire sample period and different subperiods. Clearly, regardless of sample (sub)period and term to expiration, the BS implied volatility exhibits a strong U-shaped pattern (smile) as the call option goes from deep ITM to ATM and then to deep OTM or as the put option goes from deep OTM to ATM and then to deep ITM, with the deepest ITM call-implied and the deepest OTM put-implied volatilities taking the highest values. Furthermore, the volatility smiles are the strongest for short-term options (both calls and puts), indicating that short-term options are the most severely mispriced by the BS model and present perhaps the greatest challenge to any alternative option pricing model. For a given sample (sub)period and moneyness range, the implied volatility is downward-sloping in most cases and exhibits a slight U-shape in some cases, as the term to expiration increases. This is again true for both calls and puts. These findings of clear moneyness-related and maturity-related biases associated with the BS are consistent with those in the existing literature (e.g., Bates (1996b)). Therefore, any acceptable alternative to the BS model must show an ability to properly price non-ATM options, especially short-term OTM calls and puts. As the smile evidence is indicative of negatively-skewed implicit return distributions with excess kurtosis, a better model must be based on a distributional assumption that allows for negative skewness and excess kurtosis.

⁶ In an earlier version of the article we also report the performance results for the SI model. Since incorporating stochastic interest rates does not help improve performance much, we omit the SI model from the discussions to follow. For the same reason, we do not report the SVSI-J model's results.

Table II
Implied Volatility from the Black-Scholes Model

The implied volatility is obtained by inverting the Black-Scholes model separately for each call (put) option contract. The implied volatilities of individual calls (puts) are then averaged within each moneyness-maturity category and across the days in the sample. Moneyness is determined by S/K , where S denotes the spot S&P 500 index level and K is the exercise price.

Sample Period	S/K	Call Options Days-to-Expiration			Put Options Days-to-Expiration		
		<60	60–180	≥180	<60	60–180	≥180
June 1988– May 1991	<0.94	18.27	17.25	16.58	24.64	20.05	18.93
	0.94–0.97	16.64	16.89	17.30	17.91	17.57	18.11
	0.97–1.00	16.95	17.76	17.72	16.95	18.00	18.54
	1.00–1.03	18.80	18.95	18.83	18.68	19.25	19.63
	1.03–1.06	21.40	20.04	19.91	21.29	20.37	20.80
	≥1.06	28.72	23.14	21.35	26.77	23.72	23.38
June 1988– May 1989	<0.94	17.27	16.55	16.09	23.15	19.80	20.10
	0.94–0.97	16.21	16.42	16.95	17.66	17.62	19.11
	0.97–1.00	16.33	16.89	17.03	16.11	17.51	18.84
	1.00–1.03	17.70	17.58	18.04	17.42	18.19	19.81
	1.03–1.06	19.63	17.56	18.44	19.04	18.24	20.29
	≥1.06	27.03	20.07	18.76	21.84	20.54	22.34
June 1989– May 1990	<0.94	16.16	15.64	15.96	23.20	17.80	17.61
	0.94–0.97	15.10	15.89	17.02	16.58	16.29	17.55
	0.97–1.00	15.83	16.97	17.53	15.95	16.98	17.98
	1.00–1.03	17.93	18.31	18.53	17.81	18.39	19.19
	1.03–1.06	20.74	19.45	19.91	20.65	19.72	20.62
	≥1.06	28.45	23.15	21.40	25.70	23.24	22.95
June 1990– May 1991	<0.94	19.70	18.81	17.55	25.64	20.73	18.87
	0.94–0.97	18.23	18.24	17.70	18.83	18.63	18.09
	0.97–1.00	18.65	19.25	18.37	18.70	19.43	18.88
	1.00–1.03	20.57	20.64	19.55	20.55	20.87	19.92
	1.03–1.06	23.37	22.02	20.58	23.34	22.27	21.20
	≥1.06	30.34	24.94	23.24	29.31	25.57	24.61

Note that in Table II the implied volatility of *calls* in a given ITM (OTM) category is quite similar to the implied volatility of *puts* in the opposing OTM (ITM) category, which is generally true regardless of sample period or term to expiration. Especially, for a fixed term to expiration, calls and puts imply the same U-shaped volatility patterns across strike prices. Such similarities in pricing structure exist between calls and puts mainly due to the working of the put-call parity. It is this link that makes puts and calls of the same strike price and the same expiration exhibit similar levels of mispricing, whenever one side of the put-call parity is mispriced by an option pricing model. For this reason, basing the discussions to follow solely on results obtained from the S&P 500 calls should not present a biased picture of the candidate models (either qualitatively or even quantitatively).

A. Estimation Procedure

In applying option pricing models, one always encounters the difficulty that the spot volatility and the structural parameters are unobservable. Take the SVJ model, for instance. Suppose that a call option is to be priced or hedged. Then, the strike price and the term to expiration are specified in the contract, while the spot stock price, the spot interest rate, and the matching τ -period bond price can be taken from published market data. But, the spot volatility (conditional on no jump), its related structural parameters $(\kappa_v, \theta_v, \sigma_v, \rho)$, and the jump-related parameters $(\mu_J, \sigma_J, \lambda)$ need to be estimated. In principle, one can apply econometric tools (such as maximum likelihood or the generalized methods of moments) to obtain the required estimates. However, such estimation may not be practical or convenient, because of its stringent requirement on historical data. To circumvent this difficulty, practitioners and academics alike have traditionally opted to use option-implied volatility based on the model. This practice has not only reduced data requirement dramatically but also resulted in significant performance improvement (e.g., Bates (1996a,b,c), Bodurtha and Courtadon (1987), and Melino and Turnbull (1990, 1995)). To follow this tradition, we implement each model by adapting the steps below:

Step 1. Collect N option prices on the same stock and taken from the same point in time (or same day), for any N greater than or equal to one plus the number of parameters to be estimated. For each $n = 1, \dots, N$, let τ_n and K_n be respectively the time-to-expiration and the strike price of the n -th option; Let $\hat{C}_n(t, \tau_n, K_n)$ be its observed price, and $C_n(t, \tau_n, K_n)$ its model price as determined by, for example, formula (8) with $S(t)$ and $R(t)$ taken from the market. The difference between \hat{C}_n and C_n is a function of the values taken by $V(t)$ and by $\Phi \equiv \{\kappa_R, \theta_R, \sigma_R, \kappa_v, \theta_v, \sigma_v, \rho, \lambda, \mu_J, \sigma_J\}$. For each n , define

$$\epsilon_n[V(t), \Phi] \equiv \hat{C}_n(t, \tau_n; K_n) - C_n(t, \tau_n; K_n). \quad (16)$$

Step 2. Find $V(t)$ and parameter vector Φ , to solve

$$SSE(t) \equiv \min_{V(t), \Phi} \sum_{n=1}^N |\epsilon_n[V(t), \Phi]|^2. \quad (17)$$

This step results in an estimate of the implied spot variance and the structural parameter values, for date t . Go back to Step 1 until the two steps have been repeated for each day in the sample.

The objective function in equation (17) is defined as the sum of squared dollar pricing errors, which may force the estimation to assign more weight to relatively expensive options (e.g., ITM options and long-term options) and less weight to short-term and OTM options. An alternative could be to minimize the sum of squared *percentage* pricing errors of all options, but that would lead to a more favorable treatment of cheaper options (e.g., OTM options) at the expense of ITM and long-term options. Based on this and other considerations, we choose to adopt the objective function in equation (17). Among others, Bates

(1991, 1996a,c), Dumas, Fleming and Whaley (1995), Longstaff (1995), Madan and Chang (1996), and Nandi (1996) have applied this technique for similar purposes. Applying such an implied-parameter procedure to implement the candidate models should in some sense give each model an “equal” chance, and it is also consistent with the existing practice of judging a new option pricing model’s performance relative to that of the BS when the latter is implemented using the model’s own (time-varying) implied volatility and time-varying interest rates.

B. Implied Parameters and In-Sample Pricing Fit

In implementing the above procedure, we initially use all call options available on each given day, regardless of maturity and moneyness, as inputs to estimate that day’s spot volatility and relevant structural parameters. This estimation is separately done for each model and for each day in the June 1988 to May 1991 period.⁷ The group in Table III under the heading “All Options” reports the daily average and standard error of each so-estimated parameter/volatility series as well as the daily-averaged sum of squared in-sample pricing errors (SSE), respectively for the BS, the SV, the SVSI, and the SVJ models. These reported statistics are quite informative about the internal working of the models. As such, several observations are in order. First, the implied spot volatility is on average less than 0.50 percent apart among the BS, the SV, and the SVSI models, except that the average implied standard deviation under the SVJ is 1.15 percent higher than under the BS model. For each subperiod the implied volatilities (not reported in Table III) are also close across the models. This closeness in implied volatility is somewhat surprising. It should, however, be recognized that option prices and hedge ratios are generally sensitive to the volatility input (see Figlewski (1989)). Even small differences in volatility can lead to significantly different pricing and hedging results.

Second, the estimated structural parameters for the spot volatility process generally differ across the SV, the SVSI, and the SVJ models (each assuming stochastic volatility). To appreciate these estimates, recall that in the SV model the skewness and kurtosis levels of stock returns are respectively controlled, for the most part, by correlation ρ and volatility variation coefficient σ_v ; The SVSI model relies on the same flexibility, with the additional caveat of having stochastic interest rates to ensure more proper discounting of future payoffs; In addition to inheriting all features of the SV, the SVJ model also allows price jumps to occur, which can internalize more negative skewness and higher kurtosis without making other parameters unreasonable. With this in mind, note from Table III that when all calls on a given day are used as input for the estimation, (i) the implied speed-of-volatility-adjustment κ_v is the highest for the SVJ; (ii) the implied long-run mean volatility is 18.65 percent,

⁷ For every model the daily parameter and spot volatility estimates are reasonably stable from subperiod to subperiod. Histogram-based inferences indicate that the majority of the estimated values are centered around the mean. To save space, the subsample results are not reported and are available upon request.

Table III
Implied Parameters and In-Sample Fit

Each day in the sample, the structural parameters of a given model are estimated by minimizing the sum of squared pricing errors between the market price and the model-determined price for each option. The daily average of the estimated parameters is reported first, followed by its standard error in parentheses. The parameters in the groups under “All Options”, “Short-Term Options”, and “At-the-Money Options” are obtained by respectively using all the available options, only short-term options, and only ATM options in the day as input into the estimation. For each model, SSE in a given column group denotes the daily average sum of squared errors for all options after the All-Options-Based, Maturity-Based, or Moneyness-Based treatment. The structural parameters κ_v , θ_v/κ_v , and σ_v (κ_R , θ_R/κ_R , and σ_R) are respectively the speed of adjustment, the long-run mean, and the variation coefficient of the diffusion volatility $V(t)$ (the spot interest rate $R(t)$). The parameter μ_J represents the mean jump size, λ the frequency of the jumps per year, and σ_J the standard deviation of the logarithm of one plus the percentage jump size. V_J is the instantaneous variance of the jump component. BS, SV, SVSI, and SVJ, respectively, stand for the Black-Scholes, the stochastic-volatility model, the stochastic-volatility and stochastic-interest-rate model, and the stochastic-volatility model with random jumps.

Parameters	All Options				Short-Term Options				At-the-Money Options			
	BS	SV	SVSI	SVJ	BS	SV	SVSI	SVJ	BS	SV	SVSI	SVJ
κ_v	1.15 (0.03)	0.98 (0.04)	2.03 (0.06)		1.62 (0.09)	1.47 (0.08)	3.93 (0.08)		0.99 (0.02)	0.71 (0.02)	1.74 (0.04)	
θ_v	0.04 (0.00)	0.04 (0.00)	0.04 (0.00)		0.04 (0.00)	0.04 (0.00)	0.04 (0.00)		0.04 (0.00)	0.04 (0.00)	0.04 (0.00)	
σ_v	0.39 (0.00)	0.42 (0.00)	0.38 (0.00)		0.44 (0.00)	0.45 (0.00)	0.40 (0.00)		0.40 (0.00)	0.43 (0.00)	0.40 (0.00)	
ρ	-0.64 (0.01)	-0.76 (0.01)	-0.57 (0.01)		-0.76 (0.01)	-0.80 (0.01)	-0.52 (0.01)		-0.70 (0.01)	-0.79 (0.01)	-0.58 (0.01)	
λ			0.59 (0.02)				0.61 (0.02)				0.68 (0.02)	
μ_J			-0.05 (0.00)				-0.09 (0.00)				-0.04 (0.00)	
σ_J			0.07 (0.00)				0.14 (0.00)				0.06 (0.00)	
$\sqrt{V_J}$ (%)			6.15 (0.22)				12.30 (0.17)				6.65 (0.21)	
κ_R			0.58 (0.02)				0.40 (0.02)				0.69 (0.02)	
θ_R			0.02 (0.00)				0.02 (0.00)				0.02 (0.00)	
σ_R			0.03 (0.00)				0.03 (0.00)				0.03 (0.00)	
Implied Volatility (%)	18.23 (0.14)	18.66 (0.14)	18.65 (0.15)	19.38 (0.16)	18.15 (0.14)	18.45 (0.14)	18.54 (0.14)	20.65 (0.15)	18.74 (0.14)	18.48 (0.14)	18.36 (0.15)	19.03 (0.16)
SSE	69.60	10.63	10.68	6.46	28.09	5.48	5.16	2.63	25.34	5.98	5.45	5.31

20.20 percent, and 15.32 percent, respectively, for the SV, the SVSI, and the SVJ; and (iii) the variation coefficient σ_v and the magnitude of ρ are the lowest for the SVJ, followed by the SV model. These estimates together present the picture that, to the extent that the pricing structure of the calls can be explained respectively by each model, the SVJ model's demand on the $V(t)$ process is the least stringent as it requires both the lowest σ_v and the lowest ρ (in magnitude), whereas the SVSI requires σ_v and ρ to be respectively as high as 0.42 and -0.76 . The SVJ model attributes part of the implicit negative skewness and excess kurtosis to the possibility of a jump occurring with an average frequency of 0.59 times per year and an average jump size of -5 percent (with the jump size uncertainty estimated at 7 percent).⁸ The finding that the SVSI implied-parameter values seem to be less reasonable than their SV counterparts is surprising, as one would expect the three extra parameters (related to the interest rate process) to make the SVSI model fit the data better. This poor performance by the SVSI will show up in other measures to be examined as well, suggesting that having more parameters in an option pricing model does not necessarily mean better performance. Note that under the SVSI the parameter estimates for the short-rate process are comparable to those reported in Chan *et al.* (1992). We defer further discussion on the reasonableness of the implied parameters until a later point.

Finally, the fact that incorporating stochastic interest rates does not seem to enhance the SV model's fit is further illustrated by each model's sum of squared pricing errors (SSE) across all calls on an average day. From the "All Options" panel of Table III, the SSE is 69.60 for the BS and 6.46 for the SVJ, while it is 10.63 and 10.68 respectively for the SV and the SVSI. Indeed, the SVSI and the SV result in similar in-sample fit. Allowing jumps to occur does, however, improve the SV model's in-sample fit further.

In light of the BS model's moneyness- and maturity-related biases, researchers and especially practitioners have tried to find ways to "live with a smile." One of the proposed ways, while arguably ad hoc, is to estimate and use an "implied-volatility matrix." For example, if the call option being evaluated is ATM and has one month to expiration, use as input to the BS formula the volatility implied by one-month calls of similar moneyness. To see how the candidate models fare against each other under such a matrix treatment, we reestimate and implement the four models, each time using one of six alternative sets of call options traded on a given day: short-term calls, medium-term calls, long-term calls, OTM calls, ATM calls, and ITM calls. Those maturity-based or moneyness-based parameter estimates are then applied to price or hedge options in the corresponding maturity or moneyness category.

⁸ Examining the SV and the SVJ models together, Bates (1996a,c) also finds that the SVJ is less demanding than the SV on the volatility process and its correlation with stock price changes. For the post-1987 crash years, he identifies an infrequent negative price jump implicit in S&P 500 futures options of a magnitude similar to ours. His other parameter estimates for the SV and the SVJ are also similar to ours in magnitude.

Given the space constraint, we present in Table III (i) the parameter sets implied by all short-term calls in the sample and (ii) those implied by all ATM calls. Short-term options have been the most challenging, and ATM options have been much of the focus of empirical option pricing. Let's first look at the average parameter values implied by each day's short-term calls in Table III. For each model, the estimated structural parameters of the volatility process are respectively higher than their counterparts implied by all options of each given day. In particular, the volatility coefficient σ_v is higher for each model than before, meaning that for the short-term options to be priced properly the volatility process needs to be more volatile than for all options of any maturity to be priced. The implied magnitude of ρ is also higher than before, for both the SV and the SVSI models. More strikingly, even though the implied magnitude of ρ is lower now under the SVJ, the implied jump frequency λ , mean jump size μ_J , and jump volatility σ_J are all higher in magnitude than under "All Options." This is to say that for the observed short-term option prices to be consistent with the SVJ model, more frequent and more significant market crashes (on average, 9 percent price drops) would have to be implicit in the underlying stock price process! While the SV and the SVSI attribute the relatively high premiums of short-term options to high volatility variations and significantly negatively correlated volatility shocks with underlying price changes, the SVJ attributes them to the implicit existence of significant and somewhat frequent market crashes.

When only ATM options are used to back out the parameters, the resulting estimates for volatility-related parameters do not significantly differ from their counterparts under "All Options." But to price the observed ATM option prices properly, all the three models with stochastic volatility would require volatility shocks to be more negatively correlated with underlying price changes. More interestingly, under the SVJ model these option prices imply more frequent but lower-magnitude price corrections (on average, 4 percent price drops) than either all options or short-term options do.

As expected, when the parameters and spot volatility are backed out separately using each of the six sets of option prices, the respective in-sample fits of the four models are better than when the same one set of parameters is applied to all options. This is reflected by the significant reduction in each model's daily-averaged SSE. Thus, even though ad hoc, the "implied-parameter matrix" treatment helps improve each model's respective in-sample pricing performance.

The above conclusion has, however, an unfortunate implication as well. That is, if each candidate option pricing model were correctly specified, the six sets of option prices, formed across either moneyness or maturity, should not have resulted in different implied parameter/volatility values nor should the "implied-parameter matrix" treatment have led to any performance improvement. Table III thus indicates that every candidate model is misspecified.

IV. Assessment of Relative Model Misspecification

In the two subsections to follow, we assess each model's misspecification from a different angle.

A. Implied Volatility Graphs

The first diagnostic of relative model misspecification is to compare the implied-volatility patterns of each model across both moneyness and maturity, as is done in Rubinstein (1985). For this exercise, we use the subsample data from July 1990 to December 1990.⁹ The basic procedure for backing out each model's implied-volatility series is as follows. First, substitute the spot index and interest rates of date t as well as the structural parameter values implied by all date $(t - 1)$ option prices, into the option pricing formula, which leaves only the spot volatility undetermined. Next, for each given call option of date t , find a spot volatility value that equates the model-determined price with the observed price of the call. Then, after repeating these steps for all options in the sample, obtain for each moneyness-maturity category an average implied-volatility value. These estimates are grouped into three maturity categories and plotted in Figure 1, respectively for the BS, the SV, the SVSI, and the SVJ models. Due to the difference in sample periods used, the comparable implied-volatility levels may not be exactly the same between Figure 1 and Tables II and III.

In Figure 1, the SVJ model's implied-volatility pattern smiles the least for short-term options, followed in increasing order by the SVSI, the SV, and the BS model. However, all models still show some U-shaped moneyness-related biases, indicating misspecification by all. For medium-term and long-term calls, the implied volatility exhibits a moneyness-related smile only under the BS model, but not so under the SV, the SVSI, and the SVJ models. Overall, the SV's and the SVSI's patterns are quite close, on a maturity-by-maturity basis. Further, relative to other models', the SVJ's implied volatility is persistently higher (by about 1.5 percent on average).

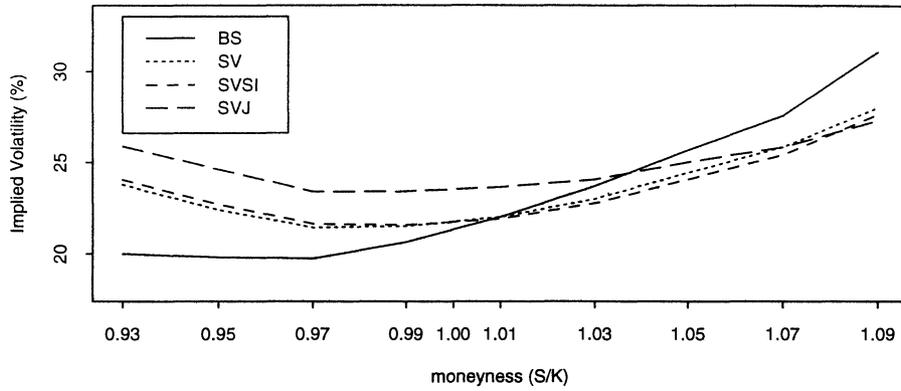
Also in Figure 1, the pricing models (except the SVJ) yield, for each given maturity category, virtually identical implied-volatility values for ATM options. Take as an example the short-term options. The three implied-volatility curves all intersect at about the ATM point. The same is true for the other maturity categories.

B. Internal Consistency of Implied Parameters

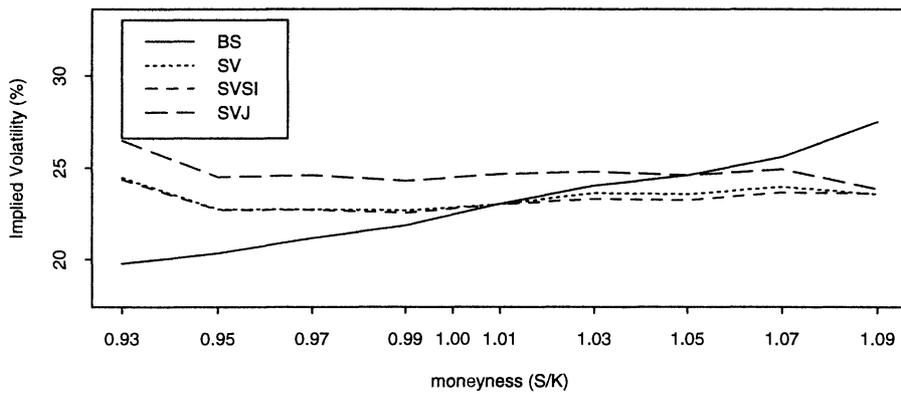
Another way to gauge model misspecification is to follow the approach taken by Bates (1996a,c) and examine whether each model's implied parameters are consistent with those implicit in the time series of (a) the S&P 500 returns, (b) the (implied) volatility, and (c) the spot interest rate. That is, are the daily averages of the implied parameters similar in magnitude to those from the time series counterparts? The closer the implied parameters, the closer the implied time-series path to its observed counterpart for each given variable and hence the less misspecified the model.

⁹ Since volatility changes over time, we focus on the average implied-volatility patterns for a relatively short period of time, rather than for the entire 3-year period. Results from another 6-month subperiod are similar.

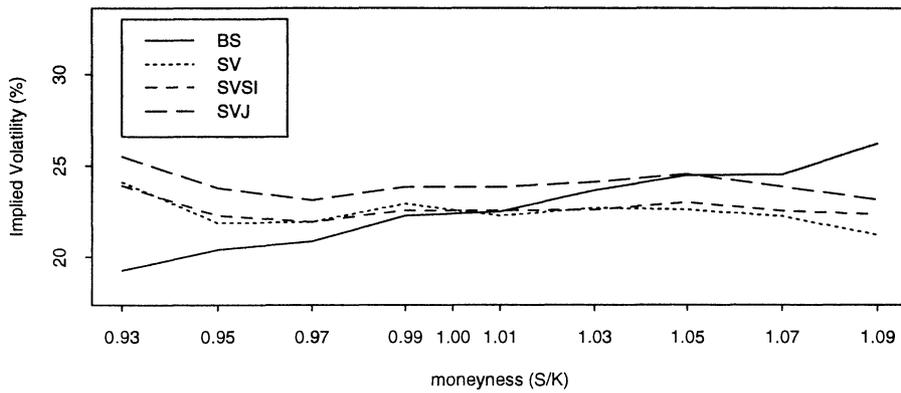
Days-to-expiration < 60



60 < Days-to-expiration < 180



Days-to-expiration > 180



Observe that the option-implied parameters correspond to the risk-neutral distributions while those estimated from observed time-series data are for the true distributions. Thus, before making the desired comparisons, we need to separate out the true distributional parameter values from their risk-neutral counterparts. For this, we rely on the general-equilibrium models of Bakshi and Chen (1997a) and Bates (1996a,c) in which the factor risk premiums are proportional to the respective factors and consequently the processes for $V(t)$, $R(t)$, $q(t)$ and $J(t)$ under the true probability measure share the same stochastic structure as their counterparts under the risk-neutral measure. Specifically, θ_v , σ_v , ρ , θ_R , σ_R , and σ_J are the same under either probability; Only κ_v , κ_R , λ , and μ_J will change when the probability measure changes from the risk-neutral to its true counterpart. Let these parameters under the true probability measure be respectively denoted by $\bar{\kappa}_v$, $\bar{\kappa}_R$, $\bar{\lambda}$, and $\bar{\mu}_J$. According to Bates (1991), when the risk aversion coefficient of the representative agent is bounded within a reasonable range, the parameters of the true distributions will not differ significantly from their risk-neutral counterparts.

For the overall sample period from June 1988 to May 1991, the annualized daily S&P 500 returns have a mean of 12.7 percent, a volatility of 17.47 percent, a skewness of -0.43 , and a kurtosis of 6.58. The historical volatility is indeed lower than its option-implied counterparts (see Table III). The negative skewness and the high kurtosis are in contrast with the skewness (of zero) and kurtosis (of 3) allowed by the log-normal distribution in the BS model. The distributional assumption of the BS is thus overwhelmingly rejected by the data. We only need to focus attention on the relative misspecification of the three models with stochastic volatility. In the rest of this subsection, we treat the volatility implied by *all options* in a given day as a surrogate for the unobservable true spot volatility of that day.

Let us first examine the consistency of the option-implied correlation ρ with the sample correlation between daily returns and volatility changes of the S&P 500 index. If an option model is correctly specified, the average ρ value implied by the option prices must equal its time-series counterpart estimated from the daily price and volatility changes. The row marked "Time-series estimate" in Table IV provides such estimates of ρ at -0.28 , -0.23 , and -0.27 , respectively, under the SV, the SVJ, and the SVSI model. The magnitudes of these estimates are much lower than their option-implied counterparts (-0.64 , -0.57 , and -0.76), suggesting that for each model the correlation level implicit in

Figure 1. The implied volatility graphs are based on the six-month sample period from July 1990 through December 1990. Using as inputs (i) current day's interest rate and S&P 500 index value and (ii) previous day's implied structural parameters, we invert each option formula from the market price of a given option, to obtain the model's implied volatility corresponding to this option contract. For each model, the reported implied volatility in a given moneyness-maturity category is the average of all calls in that moneyness-maturity category and over the entire six-month period. BS, SV, SVSI, and SVJ respectively stand for the Black-Scholes, the stochastic-volatility model, the stochastic-volatility and stochastic-interest-rate model, and the stochastic-volatility model with random jumps.

Table IV
Consistency Tests

The reported parameter values under “Average Implied” correspond to the daily average of the respective implied-parameter values. The parameters κ_v , θ_v/κ_v , and σ_v represent the speed of adjustment, the long-run mean, and the variation coefficient of the risk-neutralized process for diffusion volatility $V(t)$. Similarly, κ_R , θ_R/κ_R , and σ_R represent the speed of adjustment, the long-run mean, and the variation coefficient of the risk-neutralized process for the spot interest rate $R(t)$. The parameter ρ is the correlation between equity shocks and volatility shocks. The corresponding values of these parameters under the true probability measure are reported in the row under “MLE”, and they are obtained by applying the maximum-likelihood estimation method to the volatility (or spot interest rate) time series. The row under “ p -value (H_0 hold)” indicates the probability of the null hypothesis holding that for a given parameter, its daily-averaged option-implied value equals its maximum likelihood estimate. The row under “Time series estimate” reports the sample time-series correlation between daily S&P 500 index returns and daily changes in the implied volatility of a given option model. Standard errors are in parentheses. BS, SV, SVSI, and SVJ, respectively, stand for the Black-Scholes, the stochastic-volatility model, the stochastic-volatility and stochastic-interest-rate model, and the stochastic-volatility model with random jumps.

	SV			SVJ			SVSI			SVSI			
	Volatility Parameters			Volatility Parameters			Volatility Parameters			Interest Rate Parameters			
	$\bar{\kappa}_v$	θ_v	σ_v	$\bar{\kappa}_v$	θ_v	σ_v	$\bar{\kappa}_v$	θ_v	σ_v	ρ	$\bar{\kappa}_R$	θ_R	σ_R
Average implied	1.15	0.04	0.39	2.03	0.04	0.38	0.98	0.04	0.42	-0.76	0.58	0.02	0.03
MLE	1.16	0.04	0.10	1.09	0.03	0.10	1.14	0.04	0.10		3.02	0.14	0.04
	(0.18)	(0.00)	(0.00)	(0.14)	(0.00)	(0.00)	(0.19)	(0.01)	(0.00)		(0.09)	(0.00)	(0.00)
p -value (H_0 hold)	[0.97]	[0.78]	[0.00]	[0.00]	[0.16]	[0.00]	[0.39]	[0.85]	[0.00]		[0.00]	[0.00]	[0.00]
Time series estimate										-0.23			
										(0.03)			

option prices is inconsistent with the time-series relation between stock returns and implied volatility. Each of the three models is hence significantly misspecified. On a relative scale, however, this departure between the average implied and the time-series estimated is the weakest for the SVJ, and the strongest for the SVSI. Based on his estimated general autoregressive conditional heteroskedasticity (EGARCH) specification for equity-return dynamics, Nelson (1991) gives an estimate of -0.12 for the correlation between stock returns and changes in the true volatility, which is closer to our time-series estimates than to the average option-implied values of ρ .

Next, we adopt the maximum-likelihood (ML) method proposed by Bates (1996a) to estimate the structural parameters of $V(t)$ and $R(t)$ (wherever applicable) under the true probability measure. Take the volatility process as an example. Using the implied-volatility time series as inputs, maximize the log-likelihood function

$$\max_{\bar{\kappa}_v, \theta_v, \sigma_v} \sum_{t=1}^T \ln\{P[\ln(V(t+1))|V(t)]\}, \tag{18}$$

where $P[\cdot | \cdot]$ denotes the transition density of the non-central χ^2 distribution given by

$$P[\ln V(t + \Delta t)|V(t)] = \frac{[cV(t + \Delta t)]^{2\theta_v/\sigma_v^2}}{\exp[cV(t + \Delta t) + cV(t)e^{-\bar{\kappa}_v\Delta t}]} \sum_{j=0}^{\infty} \frac{[c^2V(t)V(t + \Delta t)e^{-\bar{\kappa}_v\Delta t}]^j}{G[(2\theta_v/\sigma_v^2) + j]j!}, \tag{19}$$

where $c^{-1} \equiv (1/2\bar{\kappa}_v) \sigma_v^2 (1 - e^{-\bar{\kappa}_v\Delta t})$, and $G(\cdot)$ denotes the (statistical) Gamma function. The ML estimates of the structural parameters are reported in Table IV for the three models. Two observations are in order. First, for each model, the ML estimates of $\bar{\kappa}_v$ and θ_v are statistically indistinguishable from their respective option-implied counterparts (except for the $\bar{\kappa}_v$ estimate of the SVJ model). The p -values for the null hypothesis of equality between the ML and the option-implied estimates are all in excess of 15 percent (except for the SVJ case noted). Second, the implied value of σ_v is, for each model, about four times its ML estimate. The volatility process implicit in option prices is therefore much too volatile, relative to each implied-volatility time series! According to this yardstick, the three models are equally misspecified. This finding is similar to those of Bates (1996a,c) using currency and S&P 500 futures options.

By replacing $V(t)$ in equation (18) with $R(t)$, we also obtain maximum-likelihood estimates for θ_R , $\bar{\kappa}_R$, and σ_R , and report them in Table IV for the SVSI (as it is the sole model assuming stochastic interest rates). Unlike the previous case for the volatility parameters, the ML estimate of σ_R is similar to its option-implied counterpart, but the ML estimates of θ_R and $\bar{\kappa}_R$ are several times as large as their option-implied counterparts. That is, interest rate volatility implicit in option prices is consistent with the interest-rate time series, but the

mean-reverting speed and the long-run mean of the spot rate implicit in option prices are much lower than the spot rate time series suggests. A possible cause for this departure is the existence of a negative interest-rate risk premium, which tends to make the risk-neutral κ_R much lower than the true $\bar{\kappa}_R$.

In summary, the models with stochastic volatility each rely on implausible levels of correlation ρ and volatility variation σ_v to rationalize the observed option prices. While the SV, the SVJ, and the SVSI are clearly misspecified (though to a lesser degree compared to the BS),¹⁰ how will they perform in pricing and hedging options? We answer this question in the sections to follow.

V. Out-of-Sample Pricing Performance

We have shown that the in-sample fit of daily option prices is increasingly better as we extend from the BS to the SV and then to the SVJ model, even though going from the SV to the SVSI does not necessarily improve the fit much further. As one may argue, this increasingly better fit might simply be a consequence of having an increasingly larger number of structural parameters. To lower the impact of this connection on inferences, we turn to examining each model's out-of-sample cross-sectional pricing performance. For out-of-sample pricing, the presence of more parameters may actually cause overfitting and have the model penalized if the extra parameters do not improve its structural fitting.

For this purpose, we rely on *previous day's* option prices to back out the required parameter/volatility values and then use them as input to compute *current day's* model-based option prices. Next, we subtract the model-determined price from its observed counterpart, to compute both the absolute pricing error and the percentage pricing error. This procedure is repeated for every call and each day in the sample, to obtain the average absolute and the average percentage pricing errors and their associated standard errors. These steps are separately followed for the BS, the SV, the SVSI, and the SVJ models. Table V reports the pricing results, where for clarity the standard errors for each estimate are omitted as they are generally very small and close to zero. Three groups of results are presented to reflect differences in the parameter/volatility values used in the model price calculations. Pricing errors reported under the heading "All-Options-Based" are obtained using the parameter/volatility values implied by *all of the previous day's call options*. Those under "Maturity-Based" are obtained using the parameter/volatility values implied by those previous-day calls *whose maturities lie in the same category (short-term, medium-term, or long-term) as the option being priced*. Pricing errors under "Moneyness-Based" are obtained using the parameter/volatility values implied by those previous-day calls *whose moneyness levels lie in the same*

¹⁰ See Bates (1996c) for other types of consistency tests. He also corrects for measurement-error-induced correlations among fitting errors across different contracts. To move on to our pricing and hedging exercise, we provide only the consistency tests just discussed. In addition, we conduct maximum-likelihood estimations using ATM-option-implied volatilities and find the results similar to those reported in Table IV.

category (OTM, ATM, or ITM) as the option being priced. In other words, the pricing errors under “Maturity-Based” and “Moneyness-Based” respectively reflect each model’s results from the “implied-parameter matrix” treatments based first on maturity and then on moneyness.

We begin with the absolute and the percentage pricing errors, respectively given in Panels A and B of Table V, corresponding to “All-Options-Based.” First, both pricing error measures rank the SVJ model first, the SVSI second, the SV next, and the BS last, except that for a few categories either the SV or the SVSI performs slightly better than the others. According to both measures, the SVSI does slightly better than the SVJ in pricing the deepest OTM calls (regardless of maturity) and the long-term deepest ITM calls. The second part of the last statement may not be surprising since one would expect the long-term deep ITM calls to be the most sensitive to interest rates. But, the fact that the SVJ does not surpass the SVSI in pricing deep OTM calls is somewhat a surprise because one would expect the opposite to be true. Second, regardless of option moneyness or maturity, incorporating stochastic volatility produces by far the most important improvement over the BS model, reducing the absolute pricing errors typically by 20 percent to 70 percent. Pricing improvement for both OTM (especially the deepest OTM) and ITM calls is particularly striking. For example, take a typical OTM call with moneyness less than 0.94 and with less than 60 days to expiration. From Table I, the average price for such a call is \$1.68. When the BS is applied to value this call, the resulting absolute pricing error is, on average, \$0.78 as shown in Table V, but when the SV is applied, the average error goes down to \$0.42. As another example, for calls of the deepest moneyness ($S/K \geq 1.06$) and the longest term-to-expiration (greater than 180 days), their average price is \$58.12, the BS gives an average pricing error of \$1.57, and the SV results in an average error of \$0.65. Table V, together with Figure 1, thus suggests that once stochastic volatility is modeled, adding other features will usually lead to second-order pricing improvement. Third, for a given moneyness category and regardless of the pricing model, the absolute pricing errors typically increase from short- to medium- to long-term options. By the percentage pricing error measure, while the BS exhibits clear moneyness- and maturity-related biases, the other three models do not except for short-term options. In fact, except for the deepest OTM calls as well as short-term calls, the percentage pricing errors are all below 1 percent in magnitude for the SV, the SVSI, and the SVJ.

A possible concern about the relatively large mispricing of short-term as well as OTM options is that the objective function in equation (17) is biased in favor of more expensive calls (i.e., long-term and ITM calls). In addition, as shown in Table I, far more sample observations are in the more expensive, ITM categories, which is also to the disadvantage of OTM options. As each estimation tries to minimize the sum of squared dollar pricing errors, these two factors must have exaggerated the extent of poor fit for short-term and OTM options by each candidate pricing model. This possible exaggeration, however, should not affect the overall conclusion regarding the pricing structure of short-term and OTM options relative to others. The reason is that in both Table II and Figure

Table V
Out-of-Sample Pricing Errors

For a given model, we compute the price of each option using the previous day's implied parameters and implied stock volatility. The reported absolute pricing error is the sample average of the absolute difference between the market price and the model price for each call in a given moneyness-maturity category. The reported percentage pricing error is the sample average of the market price minus the model price, divided by the market price. The results under "All-Options-Based" are obtained using the parameters implied by all of the previous day's calls; those under "Maturity-Based" using the parameters implied by the previous day's options of a given maturity (short-, medium-, or long-term) to price the current day's options of the same maturity; those under "Moneyness-Based" using the parameters implied by the previous day's options of a given moneyness (Out-, At-, or In-the-money; OTM, ATM, ITM) to price the current day's options of the same moneyness. The sample period is June 1988–May 1991, with a total of 38,749 call option prices. BS, SV, SVSI, and SVJ, respectively, stand for the Black-Scholes, the stochastic-volatility model, the stochastic-volatility and stochastic-interest-rate model, and the stochastic-volatility model with random jumps.

Moneyness S/K	Model	All-Options-Based Days-to-Expiration			Maturity-Based Days-to-Expiration			Moneyness-Based Days-to-Expiration		
		<60	60–180	≥180	<60	60–180	≥180	<60	60–180	≥180
Panel A: Absolute Pricing Errors										
<0.94	BS	\$0.78	\$1.39	\$1.89	\$1.02	\$1.48	\$1.78	\$0.41	\$0.63	\$0.78
	SV	0.42	0.43	0.61	0.38	0.42	0.58	0.32	0.36	0.53
	SVSI	0.37	0.39	0.57	0.38	0.40	0.52	0.30	0.36	0.53
	SVJ	0.37	0.40	0.59	0.27	0.40	0.58	0.33	0.36	0.54
0.94–0.97	BS	0.76	1.02	1.16	0.73	1.07	1.15	0.45	0.53	0.69
	SV	0.46	0.41	0.54	0.33	0.41	0.54	0.34	0.38	0.53
	SVSI	0.40	0.40	0.55	0.34	0.41	0.52	0.34	0.38	0.52
	SVJ	0.38	0.38	0.53	0.25	0.39	0.53	0.33	0.38	0.51
0.97–1.00	BS	0.61	0.62	0.66	0.51	0.64	0.66	0.70	0.74	0.94
	SV	0.48	0.41	0.53	0.39	0.41	0.52	0.40	0.43	0.60
	SVSI	0.47	0.41	0.54	0.39	0.42	0.51	0.39	0.42	0.60
	SVJ	0.42	0.40	0.52	0.31	0.40	0.51	0.36	0.41	0.63
1.00–1.03	BS	0.52	0.69	0.81	0.45	0.65	0.84	0.47	0.50	0.69
	SV	0.41	0.43	0.53	0.40	0.41	0.51	0.38	0.43	0.54
	SVSI	0.43	0.42	0.53	0.41	0.41	0.49	0.38	0.42	0.52
	SVJ	0.40	0.42	0.51	0.37	0.41	0.50	0.37	0.41	0.51
1.03–1.06	BS	0.76	1.21	1.30	0.77	1.14	1.37	0.51	0.85	1.76
	SV	0.45	0.47	0.55	0.41	0.41	0.51	0.48	0.48	0.67
	SVSI	0.42	0.45	0.54	0.41	0.41	0.50	0.48	0.48	0.66
	SVJ	0.39	0.44	0.53	0.39	0.41	0.51	0.39	0.42	0.53
≥1.06	BS	0.82	1.39	1.57	0.79	1.35	1.64	0.56	0.62	0.72
	SV	0.54	0.49	0.65	0.47	0.40	0.51	0.44	0.41	0.54
	SVSI	0.52	0.51	0.51	0.48	0.42	0.47	0.43	0.41	0.52
	SVJ	0.43	0.43	0.56	0.36	0.39	0.50	0.40	0.42	0.54

Table V—Continued

Moneyness <i>S/K</i>	Model	All-Options-Based Days-to-Expiration			Maturity-Based Days-to-Expiration			Moneyness-Based Days-to-Expiration		
		<60	60–180	≥180	<60	60–180	≥180	<60	60–180	≥180
Panel B: Percentage Pricing Errors										
<0.94	BS	-65.78%	-41.87%	-26.85%	-82.99%	-44.08%	-25.30%	-29.21%	-18.59%	-9.87%
	SV	-26.48	-2.78	-3.50	6.63	1.82	0.32	-15.04	-1.21	0.71
	SVSI	-17.22	-1.38	-1.43	5.26	1.38	-0.41	-10.62	-0.65	-1.71
	SVJ	-19.79	-1.44	2.13	2.88	0.14	-0.43	-12.39	-1.58	0.11
0.94–0.97	BS	-38.23	-15.14	-7.39	-35.52	-15.86	-7.18	-15.38	-2.34	2.73
	SV	-17.29	-1.80	0.28	1.85	-1.42	-0.55	-9.26	0.13	-0.06
	SVSI	-11.59	-0.70	0.12	2.14	-1.08	-0.17	-8.13	0.06	0.08
	SVJ	-11.99	-1.03	0.06	-0.39	-0.81	-0.48	-6.97	0.07	-0.04
0.97–1.00	BS	-13.73	-3.11	-1.25	-11.61	-3.60	-1.18	-17.64	-5.62	-3.63
	SV	-9.87	-0.47	-0.25	-5.08	-1.04	-0.24	-6.17	0.97	0.03
	SVSI	-8.84	-0.57	-0.39	-5.01	-1.10	-0.29	-5.04	0.89	-0.02
	SVJ	-6.87	-0.33	-0.32	-0.81	-0.63	-0.17	-3.11	0.25	-0.74
1.00–1.03	BS	1.09	2.61	2.00	1.58	2.28	2.41	-0.55	0.91	0.12
	SV	-1.16	0.54	-0.32	-1.09	-0.30	0.16	-0.68	0.75	-0.36
	SVSI	-1.73	0.05	-0.50	-1.16	-0.46	-0.09	-0.79	0.64	-0.31
	SVJ	-1.11	0.31	-0.26	0.05	-0.20	0.20	-0.35	0.47	-0.36
1.03–1.06	BS	3.90	4.57	3.70	4.02	4.29	4.01	-0.51	-2.76	-5.05
	SV	1.47	0.92	-0.32	1.05	0.09	0.25	-1.01	-0.58	0.33
	SVSI	1.00	0.50	-0.42	1.04	-0.01	-0.05	-1.07	-0.62	0.44
	SVJ	0.69	0.55	-0.20	0.22	0.03	0.24	0.26	0.34	-0.12
≥1.06	BS	2.49	3.27	2.85	2.41	3.16	3.01	1.45	0.89	-0.30
	SV	1.46	0.79	-0.66	1.18	0.32	-0.02	0.80	0.25	-0.23
	SVSI	1.36	0.74	-0.28	1.21	0.40	0.03	0.72	0.22	-0.16
	SVJ	0.86	0.50	-0.41	0.09	0.17	-0.04	0.65	0.36	-0.36

1, even when the BS implied volatility is estimated for each option individually (so that no weighting across options is involved), the volatility smile is clearly the sharpest for short-term options.

Observe that all four models produce negative percentage pricing errors for options with moneyness $S/K \leq 1.00$, and positive percentage pricing errors for options with $S/K \geq 1.03$, subject to their time-to-expiration not exceeding 180 days. This means that the models systematically overprice OTM calls while they underprice ITM calls. But the magnitude of such mispricing varies dramatically across the models, with the BS producing the highest and the SVJ the lowest errors.

The “Maturity-Based” results in Table V are obtained following the rule that short-term-options-implied parameter/volatility estimates are used to price the next day’s short-term options, medium-term-options-implied estimates to price the next day’s medium-term options, and so on. Given that short-term options are the most mispriced by every model (in terms of percentage pricing errors), this maturity-based treatment should work in favor of short-term options

in particular. While the maturity-based results in Table V do not affect the relative ranking of the four models, they do show differential ability by the models to benefit from this treatment. First, according to the absolute pricing errors, the BS model's performance is better under this treatment than under "All Options-Based" parameters only for some moneyness categories of a given maturity, and it is actually worse for other moneyness categories of the same maturity. For the other three models with stochastic volatility, the absolute pricing errors under this treatment are lower than their respective values under "All-Options-Based," with the improvement for short-term and long-term calls particularly noticeable. Among the four models, the SVJ shows the best ability in improving the pricing of short-term options (over what can be achieved under "All-Options-Based" parameters), while the SVSI is ahead of the others in further improving the pricing of long-term options. The same conclusions can be reached regarding the models even according to the percentage pricing errors.

Results from the moneyness-based treatment, in which OTM-options-based parameters are used to price OTM options and so on, also do not affect the relative ranking of the models. These results in Table V demonstrate, however, that each model can benefit differently from this moneyness-based treatment: the BS model benefits the most while the SVJ benefits the least. This finding may not be surprising given that in Figure 1, the BS shows the strongest moneyness-related biases whereas the SVJ shows the weakest such biases.

To further understand the structure of remaining pricing errors, we appeal to a regression analysis to study the association between the errors and factors that are either contract-specific or market condition-dependent. We first fix an option pricing model, and let $\epsilon_n(t)$ denote the n -th call option's percentage pricing error on day t . Then, we run the regression below for the entire sample:

$$\begin{aligned} \epsilon_n(t) = & \beta_0 + \beta_1 \frac{S(t)}{K_n} + \beta_2 \tau_n + \beta_3 SPREAD_n(t) + \beta_4 SLOPE(t) \\ & + \beta_5 LAGVOL(t - 1) + \eta_n(t), \end{aligned} \quad (20)$$

where K_n is the strike price of the call, τ_n the remaining time to expiration, and $SPREAD_n(t)$ the percentage bid-ask spread at date t of the call (i.e., $(Ask - Bid)/[0.5(Ask + Bid)]$), all of which are contract-specific variables. The variable, $LAGVOL(t - 1)$, is the (annualized) standard deviation of the previous day's intraday S&P 500 returns computed over 5-minute intervals, and it is included in the regression to see whether the previous day's volatility of the underlying may cause systematic pricing biases. The variable, $SLOPE(t)$, represents the yield differential between one-year and 30-day Treasury bills, and it provides information on whether the single-factor Cox-Ingersoll-Ross (1985) term structure model assumed in the present paper (for the SVSI) is sufficient to make the resulting option formula capture all term structure-related effects on the S&P 500 index options. In some sense, the contract-specific variables help detect the existence of cross-sectional pricing biases, whereas $LAGVOL(t - 1)$ and $SLOPE(t)$ serve to indicate whether the pricing errors over time are related to the dynamically changing market conditions.

Table VI
Regression Analysis of Pricing Errors

The regression results below are based on the equation:

$$\epsilon_n(t) = \beta_0 + \beta_1 \frac{S(t)}{K_n} + \beta_2 \tau_n + \beta_3 \text{SPREAD}_n(t) + \beta_4 \text{SLOPE}(t) + \beta_5 \text{LAGVOL}(t - 1) + \eta_n(t),$$

where $\epsilon_n(t)$ is the percentage pricing error of the n th call on date- t ; S/K_n and τ_n respectively represent the moneyness and the term-to-expiration of the option contract; the variable $\text{SPREAD}_n(t)$ is the percentage bid-ask spread; $\text{SLOPE}(t)$ the yield differential between the 1-year and the 30-day Treasury bill rates; and $\text{LAGVOL}(t - 1)$ the previous day's (annualized) standard deviation of S&P 500 index returns computed from 5-minute intradaily returns. The standard errors, reported in parentheses, are White's (1980) heteroskedasticity consistent estimator. The percentage pricing errors under the group "All-Options-Based" are obtained using the parameters implied by all of the previous day's calls. Those under "Maturity-Based" are obtained using the parameters implied by the previous day's options of a given maturity (short-, medium-, or long-term) to price current day's options of the same maturity. Those under "Moneyness-Based" are obtained using the parameters implied by the previous day's options of a given moneyness (Out-, At-, or In-the-money; OTM, ATM, or ITM) to price current day's options of the same moneyness. The sample period is June 1988–May 1991 for a total of 38,749 observations. BS, SV, SVSI, and SVJ, respectively, stand for the Black-Scholes, the stochastic-volatility model, the stochastic-volatility and stochastic-interest-rate model, and the stochastic-volatility model with random jumps.

Coefficient	All-Options-Based				Maturity-Based				Moneyness-Based			
	BS	SV	SVSI	SVJ	BS	SV	SVSI	SVJ	BS	SV	SVSI	SVJ
Constant	-0.01 (0.02)	0.20 (0.01)	0.07 (0.01)	0.13 (0.01)	-0.04 (0.02)	0.05 (0.01)	0.01 (0.01)	0.02 (0.01)	0.03 (0.02)	0.11 (0.01)	0.04 (0.01)	0.07 (0.01)
S/K	0.13 (0.02)	-0.16 (0.01)	-0.05 (0.01)	-0.10 (0.01)	0.17 (0.02)	-0.04 (0.01)	-0.00 (0.01)	-0.02 (0.01)	0.00 (0.01)	-0.09 (0.01)	-0.03 (0.01)	-0.06 (0.01)
τ	-0.02 (0.00)	0.05 (0.00)	0.02 (0.00)	0.04 (0.00)	-0.01 (0.00)	0.01 (0.00)	0.01 (0.00)	0.00 (0.00)	0.02 (0.00)	0.03 (0.00)	0.02 (0.00)	0.01 (0.00)
SPREAD	-5.37 (0.10)	-1.86 (0.06)	-1.26 (0.06)	-1.26 (0.06)	-5.18 (0.10)	-0.27 (0.06)	-0.20 (0.06)	-0.12 (0.05)	-2.78 (0.08)	-1.07 (0.05)	-0.81 (0.05)	-0.78 (0.05)
SLOPE	2.70 (0.13)	0.33 (0.07)	0.08 (0.07)	0.15 (0.07)	1.85 (0.13)	-0.13 (0.07)	-0.30 (0.07)	-0.37 (0.06)	1.79 (0.09)	0.12 (0.06)	0.03 (0.07)	0.08 (0.06)
LAGVOL	0.11 (0.01)	0.03 (0.01)	0.03 (0.01)	0.02 (0.01)	-0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.03 (0.01)	0.12 (0.01)	0.04 (0.01)	0.04 (0.01)	0.04 (0.01)
Adj. R^2	0.31	0.09	0.06	0.05	0.31	0.00	0.00	0.00	0.19	0.05	0.03	0.02

Table VI reports the regression results based on the entire sample period, where the standard error for each coefficient estimate is adjusted according to the White (1980) heteroskedasticity-consistent estimator and is given in the parentheses. The three groups of results under headings "All-Options-Based," "Maturity-Based," and "Moneyness-Based" have the same respective interpretations as in Table V. We first examine the results for the group under "All-Options-Based." Regardless of the model, each independent variable has statistically significant explanatory power of the remaining pricing errors.

That is, the pricing errors from each model have some moneyness, maturity, intra-daily volatility, bid-ask spread, and term-structure related biases. The magnitude and sign of each such bias, however, differ among the models. The pricing errors due to the three models with stochastic volatility are always biased in the same direction. To look at some point estimates, the BS percentage pricing errors will on average be 2.70 points higher when the yield spread $SLOPE(t)$ increases by one point, whereas the SV and the SVSI percentage errors will only be, respectively, 0.33 and 0.08 points higher in response. The SVSI is the only model whose pricing errors are statistically insensitive to $SLOPE(t)$, which is expected since it is the only model incorporating a dynamic term structure setup. This points out that modeling stochastic interest rates can lead to pricing improvement, albeit small. Other noticeable patterns include the following. The BS pricing errors are significantly, while the SV, the SVSI, and the SVJ pricing errors are only barely, increasing in the underlying stock's volatility on the previous day, which confirms that modeling stochastic volatility is important. The deeper in-the-money the call or the wider its bid-ask spread, the lower the SV's, the SVSI's, and the SVJ's mispricing. But, for the BS model, its mispricing increases with moneyness and decreases with bid-ask spread.

Even though all four models' pricing errors are, in most cases, statistically significantly related to each independent variable, the collective explanatory power of these variables is quite high only for the BS but not so for the others. The adjusted R^2 is 31 percent for the BS formula's pricing errors, 9 percent for the SV's, 6 percent for the SVSI's, and 5 percent for the SVJ model's.

The Maturity-Based and the Moneyness-Based results together present a rather intriguing picture about the four models. After the Maturity-Based pricing treatment, the SV, the SVSI, and the SVJ models' remaining pricing errors no longer show any biases at all in relation to moneyness, term-to-expiration, or lagged volatility from the previous day, while the BS model's are still as moneyness-related as under All-Options-Based parameters. The adjusted R^2 for the SV, the SVSI, and the SVJ is now zero. In contrast, the BS model's adjusted R^2 stays at 31 percent, close to its previous magnitude. Therefore, applying the maturity-based pricing treatment helps the models with stochastic volatility eliminate all contract-specific pricing biases, but it does not help the BS model improve its performance much.

The Moneyness-Based treatment produces just the opposite result: it helps improve the BS model's performance by a wide margin, but does not help the other three. This conclusion is supported by comparing the four model's respective R^2 values and coefficient estimates to those either under All-Options-Based or under Maturity-Based parameters. Strikingly, the Moneyness-Based treatment is supposed to neutralize any moneyness-related pricing biases by any model, but the resulting moneyness-related biases for the SV, the SVSI, and the SVJ are actually stronger than those obtained under the Maturity-Based treatment. It does however clear the BS of any remaining moneyness-related bias.

To see why the Moneyness-Based treatment favors the BS the most while the Maturity-Based favors the other three, look again at the results in Panel B of Table V. In the group under All-Options-Based, the dramatic pricing errors for the BS come mostly from OTM calls (in the first two rows), whereas the large errors for the other three models are all associated with short-term calls (in the first column). Outside of the OTM categories in the first two rows, the BS percentage pricing errors show no particular relation to maturity or moneyness. For medium-term and long-term calls, the pricing errors due to the SV, the SVSI, or the SVJ are quite random across strike prices (also see the relatively flat implied-volatility graphs in Figure 1, for the three models and corresponding to medium-term and long-term options). Therefore, using Moneyness-Based implied volatilities for the BS and Maturity-Based parameter/volatility estimates for the other three models serves to correct for their respective weaknesses.

VI. Dynamic Hedging Performance

For all the hedging exercises conducted in this section, the spot S&P 500 index, rather than an S&P 500 futures contract, is used in place of the “spot asset” in each hedge.¹¹ We divide our discussion into two parts: (i) single-instrument hedges and (ii) delta-neutral hedges.

A. Single-Instrument Hedges

We first examine hedges in which only a single instrument (i.e., the underlying stock) can be employed. Under this constraint, dimensions of uncertainty that move a target option value but are uncorrelated with the underlying stock price cannot be hedged by any position in the stock and will necessarily be uncontrolled for. But, as discussed before, such factors as model misspecification and transaction costs may render this type of hedge more practical to adopt.

To make the point precise, imagine a situation in which a financial institution intends to hedge a short position in a call option with τ periods to expiration and strike price K . As before, we use the SVSI-J model as the point of discussion. Let $X_S(t)$ be the number of shares of the stock to be purchased and $X_0(t)$ be the residual cash position, so that the time- t value of a replicating

¹¹ This is done out of two considerations. First, the spot S&P 500 and the immediate-expiration-month S&P 500 futures price generally have a correlation coefficient close to one. This means that whether the spot index or the futures price is used in a hedge, the conclusions are most likely the same. Second, if a futures contract is used in constructing a hedge, a futures pricing formula has to be adopted. That will introduce another dimension of model misspecification (due to stochastic interest rates), which will, in turn, produce a compounded effect on the hedging results. For these reasons, using the spot index may lead to a cleaner comparison among the four option models.

portfolio is $X_0(t) + X_S(t)S(t)$. Solving this standard minimum-variance hedging problem under the SVSI-J model, we obtain

$$X_S(t) = \frac{\text{Cov}_t[dS(t), dC(t, \tau)]}{\text{Var}[dS(t)]} = \frac{V}{V + V_J} \Delta_S(t, \tau) + \rho\sigma_v \Delta_V(t, \tau) \frac{V}{S(V + V_J)} + \frac{\lambda}{S(V + V_J)} [\Lambda_1(t) - \Lambda_2(t) - \mu_J C(t, \tau)], \quad (21)$$

where $\Lambda_1(t)$ and $\Lambda_2(t)$ are respectively given in equations (A19) and (A20) of the Appendix, and the resulting cash position for the hedge is

$$X_0(t) = C(t, \tau) - X_S(t)S(t). \quad (22)$$

This solution is quite intuitive. First, if there is no jump risk (i.e., $\lambda = 0$) and stock volatility is deterministic (i.e., $\sigma_v = 0$) (or stock returns are not correlated with volatility changes, i.e., $\rho = 0$), then one only needs to be long $\Delta_S(t)$ shares of the stock. However, if volatility is stochastic and correlated with stock returns, the position to be taken in the stock must control not only for the direct impact of underlying stock price changes on the target option, but also for the indirect impact of that part of volatility changes which is correlated with stock price fluctuations. This is reflected in the second term on the right-hand side of (21), which shows that the additional number of shares needed besides Δ_S is increasing in ρ (assuming $\sigma_v > 0$). Furthermore, if jump risk is present as well, the position to be taken in the underlying stock must also hedge the impact of jump risk on the target option, which is reflected in the last term of equation (21). This term is increasing in λ and μ_J , meaning that the larger the random-jump risk, the more adjustment need be made in the hedging position. Therefore, by considering an option model with jumps, one makes the resulting hedging strategy also immunized against jump risk.

In theory the constructed partial hedge requires continuous rebalancing to reflect the changing market conditions. In practice, only discrete rebalancing is possible. To derive a hedging effectiveness measure, suppose that portfolio rebalancing takes place at intervals of length Δt . As described above, at time t short the call option, go long in $X_S(t)$ shares of the stock and invest the residual, $X_0(t)$, in an instantaneously maturing riskfree bond. The combined position is a self-financed portfolio. Next, at time $t + \Delta t$ calculate the hedging error as follows:

$$H(t + \Delta t) = X_S(t)S(t + \Delta t) + X_0(t)e^{R(t)\Delta t} - C(t + \Delta t, \tau - \Delta t). \quad (23)$$

At the same time, reconstruct the self-financed portfolio, repeat the hedging error calculation at time $t + 2\Delta t$, and so on. Record the hedging errors $H(t + l\Delta t)$, for $l = 1, \dots, M \equiv (\tau - t)/\Delta t$. Finally, compute the average absolute hedging error as a function of rebalancing frequency Δt : $H(\Delta t) = (1/M)$

$\frac{\sum_{l=1}^M |H(t + l\Delta t)|}{\sum_{l=1}^M H(t + l\Delta t)}$, and the average dollar-value hedging error: $\bar{H}(\Delta t) = (1/M)$

Single-instrument hedging errors under the BS, the SV, the SVSI, and the SVJ models are similarly determined accounting for their modeling differences. In the case of the SVJ model, the same three terms as in equation (21) still determine the single stock position, except that the characteristic functions used in the calculations should be adjusted to reflect the constant-interest-rate assumption. For the SV and the SVSI models, the jump risk-related term (the last term) does not appear and the other two terms remain. For the BS model, only the first term in equation (21) is used to determine the minimum-variance hedge.

To obtain the hedging results presented in Table VII, we follow the three steps below. First, estimate the set of parameter/volatility values implied by all call options of day $t - 1$. Next, on day t , use these parameter/volatility estimates and the current day's spot index and interest rates, to construct the desired hedge as given in equation (21) or its model-specific version. Finally, calculate the hedging error as of day $t + 1$ if the hedge is rebalanced daily or as of day $t + 5$ if the rebalancing takes place every five days. These steps are repeated for each option and every trading day in the sample. The average absolute and the average dollar hedging errors for each moneyness-maturity category are then reported for each model in Table VII. Note that hedging results obtained respectively from the Maturity-Based and the Moneyness-Based treatments are almost the same as these in Table VII and hence not reported.

Based on the absolute hedging errors in Table VII, the SV model is the best overall performer, followed by the SVJ model, and then by the SVSI. But, according to the dollar hedging errors, the SVSI performs the best among all four in hedging both OTM calls (irrespective of maturity) and long-term ITM calls. It is also clear from both Panels A and B that, regardless of hedge rebalancing frequency, the real significant improvement by the stochastic-volatility models over the BS occurs only when OTM calls are being hedged. When other categories of calls are the hedging target, the performance is in most cases virtually indistinguishable among the four models. The hedging-based ranking of the models is thus in sharp contrast with that obtained earlier based either on out-of-sample pricing or on internal consistency of a model's estimated structural parameters.

The finding that the SVJ does not improve over the SV's hedging performance seems somewhat surprising, especially given the SVJ's better out-of-sample pricing performance (Table V). As discussed by Bates (1996a) in a different context, a possible explanation is as follows. In Table III, the average implied jump-intensity parameter λ (under "All Options") is 0.59 times per year, which means it takes, on average, about a year and a half for a jump of the average magnitude to occur. In Table VII, the results are obtained when each hedge is either rebalanced daily or once every five days. Clearly, during a one-day or five-day interval the chance for a significant price jump (or fall) to occur is very small. Thus, once stochastic volatility is modeled, hedging per-

Table VII
Single-Instrument Hedging Errors

In this table, all hedges of calls use only the underlying asset as the hedging instrument. Parameters and spot volatility implied by all options of the previous day are used to establish the current day's hedges, which are then liquidated the following day or five days later. For each target call option, its hedging error is, as of the liquidation day, the difference between its market price and the replicating portfolio value. The average absolute hedging error and the average dollar hedging error are reported for each model and for each moneyness-maturity category. The sample period is June 1988–May 1991. In calculating the hedging errors generated with daily (or 5-day) hedge rebalancing, 15,041 (or 11,704) observations are used. BS, SV, SVSI, and SVJ respectively stand for the Black-Scholes, the stochastic-volatility model, the stochastic-volatility and stochastic-interest-rate model, and the stochastic-volatility model with random jumps.

Moneyness S/K	Model	1-Day Revision Days-to-Expiration			5-Day Revision Days-to-Expiration		
		<60	60–180	≥180	<60	60–180	≥180
Panel A: Absolute Hedging Errors							
<0.94	BS	NA	\$0.35	\$0.43	NA	\$0.95	\$0.83
	SV		0.25	0.37		0.57	0.71
	SVSI		0.27	0.39		0.62	0.78
	SVJ		0.27	0.39		0.59	0.74
0.94–0.97	BS	0.24	0.33	0.42	0.70	0.77	0.80
	SV	0.23	0.32	0.43	0.59	0.66	0.72
	SVSI	0.23	0.34	0.46	0.60	0.72	0.80
	SVJ	0.23	0.33	0.44	0.60	0.69	0.78
0.97–1.00	BS	0.36	0.36	0.45	0.74	0.74	0.81
	SV	0.30	0.35	0.42	0.75	0.71	0.73
	SVSI	0.31	0.38	0.43	0.77	0.77	0.79
	SVJ	0.30	0.36	0.43	0.77	0.73	0.76
1.00–1.03	BS	0.37	0.38	0.44	0.78	0.73	0.75
	SV	0.38	0.37	0.43	0.80	0.68	0.68
	SVSI	0.39	0.39	0.45	0.83	0.75	0.75
	SVJ	0.37	0.37	0.44	0.79	0.70	0.71
1.03–1.06	BS	0.41	0.37	0.45	0.69	0.65	0.77
	SV	0.39	0.38	0.45	0.70	0.65	0.71
	SVSI	0.40	0.40	0.48	0.73	0.73	0.76
	SVJ	0.39	0.38	0.46	0.68	0.66	0.75
≥1.06	BS	0.37	0.38	0.46	0.51	0.55	0.67
	SV	0.36	0.39	0.45	0.50	0.55	0.62
	SVSI	0.36	0.41	0.46	0.51	0.60	0.66
	SVJ	0.35	0.39	0.46	0.50	0.55	0.62

formance may not be improved any further by incorporating jumps into the option pricing framework (at least when the hedge is rebalanced frequently).

B. Delta-Neutral Hedges

Now, suppose that one can use whatever instruments it takes to create a perfect delta-neutral hedge. The need for a perfect hedge can arise in situations where not only is the underlying price risk present, but also are volatility,

Table VII—Continued

Moneyness S/K	Model	1-Day Revision Days-to-Expiration			5-Day Revision Days-to-Expiration		
		<60	60–180	≥ 180	<60	60–180	≥ 180
Panel B: Dollar Hedging Errors							
<0.94	BS	NA	\$-0.03	\$-0.02	NA	\$-0.33	\$-0.19
	SV		-0.03	-0.01		-0.12	-0.10
	SVSI		-0.02	-0.00		-0.08	-0.09
	SVJ		-0.04	-0.02		-0.16	-0.11
0.94–0.97	BS	-0.13	-0.07	-0.00	-0.49	-0.20	-0.12
	SV	-0.03	-0.02	-0.02	-0.19	-0.12	-0.10
	SVSI	-0.02	-0.01	-0.01	-0.16	-0.12	-0.09
	SVJ	-0.04	-0.03	-0.02	-0.22	-0.13	-0.11
0.97–1.00	BS	-0.05	-0.03	0.00	-0.34	-0.09	-0.06
	SV	-0.04	-0.02	0.00	-0.28	-0.13	-0.09
	SVSI	-0.04	-0.02	0.00	-0.28	-0.15	-0.12
	SVJ	-0.04	-0.03	0.00	-0.30	-0.14	-0.11
1.00–1.03	BS	-0.08	-0.01	0.00	-0.37	-0.09	-0.11
	SV	-0.04	-0.02	0.00	-0.32	-0.13	-0.12
	SVSI	-0.04	-0.01	0.00	-0.30	-0.14	-0.11
	SVJ	-0.05	-0.02	-0.01	-0.33	-0.12	-0.12
1.03–1.06	BS	-0.04	-0.02	-0.03	-0.31	-0.08	-0.18
	SV	-0.05	-0.02	-0.02	-0.27	-0.12	-0.14
	SVSI	-0.05	-0.01	-0.01	-0.26	-0.11	-0.09
	SVJ	-0.05	-0.02	-0.03	-0.28	-0.10	-0.16
≥ 1.06	BS	-0.05	-0.03	-0.02	-0.18	-0.09	-0.09
	SV	-0.05	-0.03	-0.01	-0.16	-0.11	-0.08
	SVSI	-0.05	-0.02	-0.01	-0.16	-0.12	-0.07
	SVJ	-0.05	-0.03	-0.02	-0.16	-0.11	-0.07

interest rate and jump risks. In conducting this exercise, however, we should first recognize that a *perfect* hedge may not be practically feasible in the presence of stochastic jump sizes (e.g., for the SVJ and the SVSI-J models). This difficulty is seen from the existing work by Bates (1996a), Cox and Ross (1976), and Merton (1976). For this reason, whenever jump risk is present, we follow Merton (1976) and only aim for a partial hedge in which diffusion risks are completely neutralized but jump risk is left uncontrolled for. We do this with the understanding that the overall impact on hedging effectiveness of not controlling for jump risk can be small or large, depending on whether the hedge is frequently rebalanced or not.

Suppose again that the target is a short position in a call option with τ periods to expiration and strike price K . Taking the SVSI-J model as the point of discussion, the hedger will need a position in (i) some $X_S(t)$ shares of the underlying stock (to control for price risk), (ii) some $X_B(t)$ units of a τ -period discount bond (to control for $R(t)$ risk), and (iii) some $X_C(t)$ units of another call option with the same maturity but a different strike price \bar{K} (or any option on

the stock with a different maturity) to control for volatility risk $V(t)$. The time- t value of this replicating portfolio is then $X_0(t) + X_S(t)S(t) + X_B(t)B(t, \tau) + X_C(t)C(t, \tau; \bar{K})$, where $X_0(t)$ denotes the residual cash position. Deriving the dynamics for the replicating portfolio and comparing them with those of $C(t, \tau, K)$, we find the following solution:

$$X_C(t) = \frac{\Delta_V(t, \tau; K)}{\Delta_V(t, \tau; \bar{K})} \quad (24)$$

$$X_S(t) = \Delta_S(t, \tau; K) - \Delta_S(t, \tau; \bar{K})X_C(t) \quad (25)$$

$$X_B(t) = \frac{1}{B(t, \tau)\rho(\tau)} \{\Delta_R(t, \tau; \bar{K})X_C(t) - \Delta_R(t, \tau; K)\} \quad (26)$$

$$X_0(t) = C(t, \tau; K) - X_S(t)S(t) - X_C(t, \tau; \bar{K}) - X_B(t)B(t, \tau), \quad (27)$$

where all the primitive deltas, Δ_S , Δ_R and Δ_V , are as determined in equations (10)–(12).

To examine the hedging effectiveness, at time t short the call option and establish the hedge as just described. After the next interval, compute the hedging error according to

$$H(t + \Delta t) = X_0 e^{R(t)\Delta t} + X_S(t)S(t + \Delta t) + X_B(t)B(t + \Delta t, \tau - \Delta t) + X_C(t)C(t + \Delta t, \tau - \Delta t; \bar{K}) - C(t + \Delta t, \tau - \Delta t; K). \quad (28)$$

Like in the previous case, repeat this calculation for each date t and every target call in the sample to obtain a collection of hedging errors, which is then used to compute the average absolute and the average dollar-value hedging errors, both as functions of rebalancing frequency Δt .

For the BS model, the delta-neutral hedge is the same as the previous single-instrument hedge and its hedging error measures are similarly calculated as in (28), except that $\lambda = X_B(t) = X_C(t) = 0$ and $X_S(t)$ is the BS delta. Thus, the BS delta-neutral hedge involves no other instrument than the underlying stock. In the case of the SV model, set $\lambda = X_B(t) = 0$ and let Δ_S and Δ_V be as determined in the SV model. Its delta-neutral hedge hence consists of a position in both the stock and the second option contract. For the SVJ model, set $X_B(t) = 0$ and let Δ_S and Δ_V be as determined in the SVJ model. Clearly, when $\lambda = 0$, the hedge created in equations (24)–(27) becomes the one corresponding to the SVSI model.

In the cases of the SV, the SVSI, and the SVJ models, the hedge in equation (28) requires (i) the availability of prices for four time-matched target and hedging-instrumental options: $C(t, \tau, K)$, $C(t, \tau, \bar{K})$, $C(t + \Delta t, \tau - \Delta t; K)$, $C(t + \Delta t, \tau - \Delta t; \bar{K})$ and (ii) the computation of Δ_S , Δ_V , and Δ_R for both the target and the instrumental option. Due to this requirement, we use as hedging instruments only options whose prices on both the hedge-construction day and the following liquidation day were quoted no more than 15 seconds apart from the

times when the respective prices for the target option were quoted. This constraint guarantees that the deltas for the target and instrumental options on the same day are computed based on the same spot price. The remaining sample for both this delta-neutral hedging exercise and the previous single-instrument hedging contains 15,041 matched pairs when hedging revision occurs daily, and 11,704 matched pairs when rebalancing takes place at five-day intervals.

As before, we use the current day's spot index and interest rates, but parameter/volatility values implied by all of the previous-day's options, to determine the current day's hedging positions for each target call. Table VIII presents the average absolute and the average dollar hedging errors across the 18 moneyness-maturity categories and for each of the four models. A striking pattern emerging from this table is that, irrespective of moneyness-maturity category, the three models with stochastic volatility have virtually identical delta-neutral hedging errors. Therefore, consistent with the results of the previous subsection, adding jumps or stochastic interest rates to the SV model does not improve its hedging performance, at least with respect to our sample data.

When the hedges are revised daily, the BS delta-neutral hedging errors are usually two to three times as high as the corresponding hedging errors for the other three models. Improvement by the stochastic-volatility models is even more evident when the hedge revision frequency changes from daily to once every five days: the BS hedging errors increase dramatically while the other models' do not increase by much. This seems to suggest that the other three models perform much better than the BS.

The last observation perhaps raises more questions than answers. Is the hedging improvement by the three models with stochastic volatility a consequence of better model specification, or is it mostly due to the inclusion of a second option in their delta-neutral hedges? Is the fact that hedge revision frequency does not affect the hedging effectiveness of the three models by as much as it affects the performance of the BS a consequence of better model specification, or is it due to the indirect effect of the second call option on the position gamma measure?

To answer the first question, we implement the so-called delta-plus-vega-neutral hedge for the BS model, in which the underlying stock and a second call option are used respectively to neutralize the sensitivity of the hedge to underlying price risk and volatility risk. This type of strategy is clearly inconsistent with the BS setup, but such a treatment may in some sense give the BS a fairer chance. In particular, if the BS delta-plus-vega-neutral hedge results in hedging errors comparable to those from the delta-neutral hedges of the other three models, it will simply suggest that model misspecification may only have a secondary effect on hedging. We report the average hedging errors of this BS delta-plus-vega-neutral strategy under the abbreviation "BSDV" in Table VIII. Except for the ITM categories, hedging performance is indistinguishable between the BS delta-plus-vega-neutral strategy and the delta-neutral strategies for the other three models. For the two ITM call option groups (with $S/K > 1.03$), however, incorporating stochastic volatility does improve upon the BS delta-plus-vega-neutral hedging performance. Thus, for hedging these ITM calls, more appropriate model specifi-

Table VIII
Delta-Neutral Hedging Errors

In this table, all delta-neutral hedges of calls use as many hedging instruments as there are sources of risk (except the jump risk) assumed in a given option model. The only exception is the BS delta-plus-vega-neutral strategy, denoted by BSDV, which uses the underlying asset and a second call option to neutralize both the delta and vega risks of the target call, based on the Black-Scholes model. Parameters and spot volatility implied by all options of the previous day are used to establish the current day's hedges, which are then liquidated the following day or five days later. For each target call option, its hedging error is, as of the liquidation day, the difference between its market price and the replicating portfolio value. The average absolute hedging error and the average dollar hedging error are reported for each model and for each moneyness-maturity category. The sample period is June 1988–May 1991. In calculating the hedging errors generated with daily (once every five days) hedge rebalancing, 15,041 (11,704) observations are used. BS, SV, SVSI, and SVJ, respectively, stand for the Black-Scholes, the stochastic-volatility model, the stochastic-volatility and stochastic-interest-rate model, and the stochastic-volatility model with random jumps.

Moneyness S/K	Model	1-Day Revision Days-to-Expiration			5-Day Revision Days-to-Expiration		
		<60	60–180	≥180	<60	60–180	≥180
Panel A: Absolute Hedging Errors							
<0.94	BS	NA	\$0.35	\$0.43	NA	\$0.95	\$0.83
	BSDV		0.15	0.18		0.18	0.28
	SV		0.14	0.20		0.16	0.31
	SVSI		0.14	0.22		0.16	0.33
	SVJ		0.14	0.21		0.16	0.32
0.94–0.97	BS	0.24	0.33	0.42	0.70	0.77	0.80
	BSDV	0.08	0.13	0.19	0.11	0.15	0.24
	SV	0.08	0.13	0.20	0.09	0.14	0.25
	SVSI	0.08	0.13	0.20	0.09	0.14	0.25
	SVJ	0.08	0.13	0.20	0.10	0.14	0.27
0.97–1.00	BS	0.36	0.36	0.45	0.74	0.74	0.81
	BSDV	0.13	0.13	0.16	0.16	0.16	0.22
	SV	0.13	0.13	0.16	0.15	0.15	0.22
	SVSI	0.13	0.13	0.16	0.15	0.15	0.22
	SVJ	0.13	0.13	0.16	0.15	0.15	0.22
1.00–1.03	BS	0.37	0.38	0.44	0.78	0.73	0.75
	BSDV	0.13	0.14	0.16	0.16	0.16	0.22
	SV	0.13	0.14	0.16	0.16	0.16	0.22
	SVSI	0.13	0.14	0.16	0.16	0.16	0.22
	SVJ	0.14	0.14	0.17	0.17	0.17	0.23
1.03–1.06	BS	0.41	0.37	0.45	0.69	0.65	0.77
	BSDV	0.16	0.13	0.16	0.18	0.16	0.21
	SV	0.15	0.13	0.16	0.17	0.16	0.20
	SVSI	0.15	0.13	0.16	0.16	0.16	0.21
	SVJ	0.16	0.13	0.16	0.17	0.15	0.20
≥1.06	BS	0.37	0.38	0.46	0.51	0.55	0.67
	BSDV	0.19	0.16	0.23	0.22	0.21	0.33
	SV	0.15	0.14	0.19	0.17	0.17	0.25
	SVSI	0.15	0.14	0.19	0.16	0.17	0.24
	SVJ	0.15	0.14	0.20	0.17	0.17	0.26

Table VIII—Continued

Moneyness <i>S/K</i>	Model	1-Day Revision Days-to-Expiration			5-Day Revision Days-to-Expiration		
		<60	60–180	≥180	<60	60–180	≥180
Panel B: Dollar Hedging Errors							
<0.94	BS	NA	\$-0.03	\$-0.02	NA	\$-0.33	\$-0.19
	BSDV		0.01	-0.01		0.02	0.00
	SV		0.01	-0.00		0.02	0.01
	SVSI		0.01	-0.00		0.02	0.00
	SVJ		0.01	-0.00		0.02	0.01
0.94–0.97	BS	-0.13	-0.07	-0.00	-0.49	-0.20	-0.12
	BSDV	-0.04	-0.01	0.00	-0.03	0.00	-0.01
	SV	0.01	-0.01	0.00	0.00	-0.01	-0.01
	SVSI	0.01	-0.01	0.00	0.00	-0.01	0.00
	SVJ	0.01	-0.01	0.00	0.02	-0.01	-0.03
0.97–1.00	BS	-0.05	-0.03	0.00	-0.34	-0.09	-0.06
	BSDV	0.01	0.01	0.01	0.06	0.01	0.03
	SV	0.01	0.00	0.01	0.03	0.00	0.02
	SVSI	0.01	0.00	0.01	0.02	0.00	0.01
	SVJ	0.01	0.00	0.01	0.00	0.00	0.02
1.00–1.03	BS	-0.08	-0.01	0.00	-0.37	-0.09	-0.11
	BSDV	-0.01	-0.00	0.01	-0.01	-0.00	0.03
	SV	-0.01	-0.00	0.01	-0.01	0.00	0.02
	SVSI	-0.01	-0.00	0.01	-0.00	0.00	0.03
	SVJ	-0.01	0.00	0.01	-0.02	0.00	0.02
1.03–1.06	BS	-0.04	-0.02	-0.03	-0.31	-0.08	-0.18
	BSDV	0.01	0.00	-0.01	-0.05	-0.01	-0.03
	SV	0.01	0.00	-0.00	-0.02	-0.02	-0.03
	SVSI	0.01	0.00	-0.00	-0.01	-0.01	-0.02
	SVJ	0.01	0.00	-0.01	-0.03	-0.01	-0.03
≥1.06	BS	-0.05	-0.03	-0.02	-0.18	-0.09	-0.09
	BSDV	-0.02	-0.01	-0.01	-0.05	-0.02	-0.02
	SV	-0.01	-0.01	-0.00	-0.03	-0.02	-0.01
	SVSI	-0.01	0.00	0.00	-0.02	-0.00	0.01
	SVJ	-0.01	-0.01	0.00	-0.03	-0.02	-0.01

cation matters, whereas for hedging other calls, even ad hoc hedging strategies based on the BS may do just fine.

Given the performance of the BSDV in Table VIII, it is apparent that the relative insensitivity of the other three models' hedging errors to revision frequency must be mostly due to the use of the instrumental call option. It is the instrumental option position that not only neutralizes the volatility risk but also dramatically reduces the remaining gamma risk in the hedge. To see this point, take the SV delta-neutral hedge as an example. Denote the gamma (with respect to the spot price) of the target call by $\Gamma_S(t, \tau, K)$ and that of the instrumental call by $\Gamma_S(t, \tau, \bar{K})$, a detailed expression of which is provided in the Appendix. Since the position in the instrumental call is $X_C(t) = [\Delta_V(t, \tau, K)]/[\Delta_V(t,$

\bar{K}], the remaining gamma value of the SV delta-neutral hedge is given by $\Gamma_S(t, \tau; K) - X_C(t)\Delta_S(t, \tau; \bar{K})$. For a typical delta-neutral hedge under stochastic volatility, the remaining gamma value is close to zero. The following are some examples based on June 3, 1988, when the spot S&P 500 was at 265.42:

- For hedging the ATM call with strike price 265, the position taken in the chosen instrumental option is $X_C = 0.97$ and the remaining SV gamma value in the hedge is $0.020 - 0.97 \times 0.022 = -0.001$;
- For hedging the ITM call with strike price 250, $X_C = 0.81$ and the remaining SV gamma of the hedge is $0.011 - 0.81 \times 0.014 = -0.001$;
- For hedging the OTM call with strike price 275, $X_C = 0.97$ and the remaining SV gamma of the hedge is $0.022 - 0.97 \times 0.022 = 0.001$.

For a typical BSDV hedge, the remaining BS gamma is also close to zero, which explains why the BSDV hedging errors are relatively insensitive to revision frequency as well.

Another pattern to note from Table VIII is that the BS model's dollar hedging errors are always negative, indicating that the model overhedges each target option, whereas the dollar hedging errors of the other models are more random and can take either sign. Therefore, the BS formula exhibits a systematic hedging bias, while the others do not.

Comparing Tables VII and VIII, one can see that for a given option model, the conventional delta-neutral hedge (using as many instruments as there are sources of uncertainty) performs far better than its single-instrument counterpart, for every moneyness-maturity category. This may not be surprising as the former type of hedge involves more instruments (except under the BS model).¹²

VII. Concluding Remarks

We have developed a parsimonious option pricing model that admits stochastic volatility, stochastic interest rates, and random jumps. It is shown that this closed-form pricing formula is practically implementable, leads to useful analytical hedge ratios, and contains many known option formulas as special cases. This last feature has made it relatively straightforward to study the relative empirical performance of several models of distinct interest.

Our empirical evidence indicates that regardless of performance yardstick, taking stochastic volatility into account is of the first-order importance in

¹² According to Rubinstein (1985), the volatility smile pattern and the nature of pricing biases are time-dependent. To see whether our conclusion may be reversed, we separately examine the pricing and the hedging performance of the models in three subperiods: June 1988–May 1989, June 1989–1990, and June 1990–May 1991. Each subperiod contains about 10,000 call option observations. The subperiod results are qualitatively the same as those, respectively, in Tables V and VII. Separately, we examine the pricing and hedging error measures of each model when the structural parameters are not updated daily. Rather, we retain the structural parameter values estimated from the options of the first day of each month and then, for the remainder of the month, use them as input to compute the corresponding model-based prices for each traded option, except that the implied spot volatility is updated each day based on the previous day's option prices. The obtained absolute pricing and hedging errors for the subperiod June 1990–May 1991 indicate that the performance ranking of the four models also remains the same.

improving upon the BS formula. In terms of internal consistency, the SV, the SVJ, and the SVSI are still significantly misspecified. In particular, to rationalize the negative skewness and excess kurtosis implicit in option prices, each model with stochastic volatility requires highly implausible levels of volatility-return correlation and volatility variation. But, such structural misspecifications do not necessarily preclude these models from performing better otherwise. According to the out-of-sample pricing measures, adding the random-jump feature to the SV model can further improve its performance, especially in pricing short-term options; whereas modeling stochastic interest rates can enhance the fit of long-term options. With both the SVSI and the SVJ, the remaining pricing errors show the least contract-specific or market-conditions-related biases. For hedging purposes, however, incorporating either the jump or the SI feature does not seem to improve the SV model's performance further. The SV achieves the best hedging results among all the models studied, and its remaining hedging errors are generally quite small. Therefore, the three performance yardsticks employed in this article can rank a given set of models differently as they capture and reveal distinct aspects of a pricing model. Overall, our results support the claim that a model with stochastic volatility and random jumps is a better alternative to the BS formula, because the former not only performs far better but also is practically implementable.

The empirical issues and questions addressed in this article can also be reexamined using data from individual stock options, American-style index options, options on futures, currency and commodity options, and so on. Eventually, the acceptability of option pricing models with added features will be judged not only by its implementability, its internal consistency, and its pricing and hedging performance as demonstrated in this paper, but also by its success or failure in pricing and hedging other types of options. These extensions are left for future research.

APPENDIX

Proof of the Option Pricing Formula in Equation (8). The valuation partial differential equation (PDE) in equation (7) can be rewritten as:

$$\begin{aligned} & \frac{1}{2} V \frac{\partial^2 C}{\partial L^2} + \left(R - \lambda \mu_J - \frac{1}{2} V \right) \frac{\partial C}{\partial L} + \rho \sigma_v V \frac{\partial^2 C}{\partial L \partial V} + \frac{1}{2} \sigma_v^2 V \frac{\partial^2 C}{\partial V^2} \\ & + [\theta_v - \kappa_v V] \frac{\partial C}{\partial V} + \frac{1}{2} \sigma_R^2 R \frac{\partial^2 C}{\partial R^2} + [\theta_R - \kappa_R R] \frac{\partial C}{\partial R} - \frac{\partial C}{\partial \tau} - RC \\ & + \lambda E_Q \{ C(t, \tau; L + \ln[1 + J], R, V) - C(t, \tau; L, R, V) \} = 0, \quad (A1) \end{aligned}$$

where we have applied the transformation $L(t) \equiv \ln[S(t)]$. Inserting the conjectured solution in equation (8) into (A1) produces the PDEs for the risk-

neutralized probabilities, Π_j for $j = 1, 2$:

$$\begin{aligned} & \frac{1}{2} V \frac{\partial^2 \Pi_1}{\partial L^2} + \left(R - \lambda \mu_J + \frac{1}{2} V \right) \frac{\partial \Pi_1}{\partial L} + \rho \sigma_v V \frac{\partial^2 \Pi_1}{\partial L \partial V} + \frac{1}{2} \sigma_v^2 V \frac{\partial^2 \Pi_1}{\partial V^2} \\ & + [\theta_v - (\kappa_v - \rho \sigma_v) V] \frac{\partial \Pi_1}{\partial V} + \frac{1}{2} \sigma_R^2 R \frac{\partial^2 \Pi_1}{\partial R^2} + [\theta_R - \kappa_R R] \frac{\partial \Pi_1}{\partial R} - \frac{\partial \Pi_1}{\partial \tau} \\ & - \lambda \mu_J \Pi_1 + \lambda E_Q \{ (1 + \ln[1 + J]) \Pi_1(t, \tau; L + \ln[1 + J], R, V) \\ & \quad - \Pi_1(t, \tau; L, R, V) \} = 0, \quad (\text{A2}) \end{aligned}$$

and

$$\begin{aligned} & \frac{1}{2} V \frac{\partial^2 \Pi_2}{\partial L^2} + \left(R - \lambda \mu_J - \frac{1}{2} V \right) \frac{\partial \Pi_2}{\partial L} + \rho \sigma_v V \frac{\partial^2 \Pi_2}{\partial L \partial V} + \frac{1}{2} \sigma_v^2 V \frac{\partial^2 \Pi_2}{\partial V^2} \\ & + [\theta_v - \kappa_v V] \frac{\partial \Pi_2}{\partial V} + \frac{1}{2} \sigma_R^2 R \frac{\partial^2 \Pi_2}{\partial R^2} + \left[\theta_R - \left(\kappa_R - \frac{\sigma_R^2}{B(t, \tau)} \frac{\partial B(t, \tau)}{\partial R} \right) R \right] \frac{\partial \Pi_2}{\partial R} - \frac{\partial \Pi_2}{\partial \tau} \\ & + \lambda E_Q \{ \Pi_2(t, \tau; L + \ln[1 + J], R, V) - \Pi_2(t, \tau; L, R, V) \} = 0. \quad (\text{A3}) \end{aligned}$$

Observe that equations (A2) and (A3) are the Fokker-Planck forward equations for probability functions. This implies that Π_1 and Π_2 must indeed be valid probability functions, with values bounded between 0 and 1. These PDEs must be solved separately subject to the terminal condition:

$$\Pi_j(t + \tau, 0) = 1_{L(t+\tau) \geq K} \quad j = 1, 2. \quad (\text{A4})$$

The corresponding characteristic functions for Π_1 and Π_2 will also satisfy similar PDEs:

$$\begin{aligned} & \frac{1}{2} V \frac{\partial^2 f_1}{\partial L^2} + \left(R - \lambda \mu_J + \frac{1}{2} V \right) \frac{\partial f_1}{\partial L} + \rho \sigma_v V \frac{\partial^2 f_1}{\partial L \partial V} + \frac{1}{2} \sigma_v^2 V \frac{\partial^2 f_1}{\partial V^2} \\ & + [\theta_v - (\kappa_v - \rho \sigma_v) V] \frac{\partial f_1}{\partial V} + \frac{1}{2} \sigma_R^2 R \frac{\partial^2 f_1}{\partial R^2} + [\theta_R - \kappa_R R] \frac{\partial f_1}{\partial R} - \frac{\partial f_1}{\partial \tau} \\ & - \lambda \mu_J f_1 + \lambda E_Q \{ (1 + \ln[1 + J]) f_1(t, \tau; L + \ln[1 + J], R, V) \\ & \quad - f_1(t, \tau; L, R, V) \} = 0, \quad (\text{A5}) \end{aligned}$$

and

$$\begin{aligned} & \frac{1}{2} V \frac{\partial^2 f_2}{\partial L^2} + \left(R - \lambda \mu_J - \frac{1}{2} V \right) \frac{\partial f_2}{\partial L} + \rho \sigma_v V \frac{\partial^2 f_2}{\partial L \partial V} + \frac{1}{2} \sigma_v^2 V \frac{\partial^2 f_2}{\partial V^2} + [\theta_v - \kappa_v V] \\ & \cdot \frac{\partial f_2}{\partial V} + \frac{1}{2} \sigma_R^2 R \frac{\partial^2 f_2}{\partial R^2} + \left[\theta_R - \left(\kappa_R - \frac{\sigma_R^2}{B(t, \tau)} \frac{\partial B(t, \tau)}{\partial R} \right) R \right] \frac{\partial f_2}{\partial R} - \frac{\partial f_2}{\partial \tau} \\ & + \lambda E_Q \{ f_2(t, \tau; L + \ln[1 + J], R, V) - f_2(t, \tau; L, R, V) \} = 0. \quad (\text{A6}) \end{aligned}$$

with the boundary condition:

$$f_j(t + \tau, 0; \phi) = e^{i\phi L(t+\tau)} \quad j = 1, 2. \tag{A7}$$

Conjecture that the solution to the PDEs (A5) and (A6) is respectively given by

$$f_1(t, \tau, S(t), R(t), V(t); \phi) = \exp\{u(\tau) + x_r(\tau)R(t) + x_v(\tau)V(t) + i\phi \ln[S(t)]\} \tag{A8}$$

$$f_2(t, \tau, S(t), R(t), V(t); \phi) = \exp\{z(\tau) + y_r(\tau)R(t) + y_v(\tau)V(t) + i\phi \ln[S(t)] - \ln[B(t, \tau)]\} \tag{A9}$$

with $u(0) = x_r(0) = x_v(0) = 0$ and $z(0) = y_r(0) = y_v(0) = 0$. Solving the resulting systems of differential equations and noting that $B(t + \tau, 0) = 1$ respectively produce the following desired characteristic functions:

$$f_1(t, \tau) = \exp\left\{ -\frac{\theta_R}{\sigma_R^2} \left[2 \ln \left(1 - \frac{[\xi_R - \kappa_R](1 - e^{-\xi_R \tau})}{2\xi_R} \right) + [\xi_R - \kappa_R]\tau \right] - \frac{\theta_v}{\sigma_v^2} \left[2 \ln \left(1 - \frac{[\xi_v - \kappa_v + (1 + i\phi)\rho\sigma_v](1 - e^{-\xi_v \tau})}{2\xi_v} \right) \right] - \frac{\theta_v}{\sigma_v^2} [\xi_v - \kappa_v + (1 + i\phi)\rho\sigma_v]\tau + i\phi \ln[S(t)] + \frac{2i\phi(1 - e^{-\xi_R \tau})}{2\xi_R - [\xi_R - \kappa_R](1 - e^{-\xi_R \tau})} R(t) + \lambda(1 + \mu_J)\tau[(1 + \mu_J)^{i\phi} e^{(i\phi/2)(1+i\phi)\sigma_J^2} - 1] - \lambda i\phi\mu_J\tau + \frac{i\phi(i\phi + 1)(1 - e^{-\xi_v \tau})}{2\xi_v - [\xi_v - \kappa_v + (1 + i\phi)\rho\sigma_v](1 - e^{-\xi_v \tau})} V(t) \right\}, \tag{A10}$$

and,

$$f_2(t, \tau) = \exp\left\{ -\frac{\theta_R}{\sigma_R^2} \left[2 \ln \left(1 - \frac{[\xi_R^* - \kappa_R](1 - e^{-\xi_R^* \tau})}{2\xi_R^*} \right) + [\xi_R^* - \kappa_R]\tau \right] - \frac{\theta_v}{\sigma_v^2} \left[2 \ln \left(1 - \frac{[\xi_v^* - \kappa_v + i\phi\rho\sigma_v](1 - e^{-\xi_v^* \tau})}{2\xi_v^*} \right) + [\xi_v^* - \kappa_v + i\phi\rho\sigma_v]\tau \right] + i\phi \ln[S(t)] - \ln[B(t, \tau)] + \frac{2(i\phi - 1)(1 - e^{-\xi_R^* \tau})}{2\xi_R^* - [\xi_R^* - \kappa_R](1 - e^{-\xi_R^* \tau})} R(t) + \lambda\tau[(1 + \mu_J)^{i\phi} e^{(i\phi/2)(i\phi-1)\sigma_J^2} - 1] - \lambda i\phi\mu_J\tau + \frac{i\phi(i\phi - 1)(1 - e^{-\xi_v^* \tau})}{2\xi_v^* - [\xi_v^* - \kappa_v + i\phi\rho\sigma_v](1 - e^{-\xi_v^* \tau})} V(t) \right\}, \tag{A11}$$

where

$$\xi_R = \sqrt{\kappa_R^2 - 2\sigma_R^2 i\phi}, \quad \xi_v = \sqrt{[\kappa_v - (1 + i\phi)\rho\sigma_v]^2 - i\phi(i\phi + 1)\sigma_v^2},$$

$$\xi_R^* = \sqrt{\kappa_R^2 - 2\sigma_R^2(i\phi - 1)}, \quad \text{and} \quad \xi_v^* = \sqrt{[\kappa_v - i\phi\rho\sigma_v]^2 - i\phi(i\phi - 1)\sigma_v^2}.$$

The SI, the SV, the SVSI, the SVJ models are all nested within the general formula in equation (8). In the SVJ case, for instance, the partial derivatives with respect to R vanishes in equation (A1). The general solution in equations (A8)–(A9) will still apply except that now $R(t) = R$ (a constant) and $B(t, \tau) = e^{-R\tau}$. The final characteristic functions \hat{f}_j for the SVJ model are respectively given by

$$\begin{aligned} \hat{f}_1 = \exp \left\{ -i\phi \ln[B(t, \tau)] - \frac{\theta_v}{\sigma_v^2} \left[2 \ln \left(1 - \frac{[\xi_v - \kappa_v + (1 + i\phi)\rho\sigma_v](1 - e^{-\xi_v\tau})}{2\xi_v} \right) \right] \right. \\ \left. - \frac{\theta_v}{\sigma_v^2} [\xi_v - \kappa_v + (1 + i\phi)\rho\sigma_v]\tau + i\phi \ln[S(t)] \right. \\ \left. + \lambda(1 + \mu_J)\tau[(1 + \mu_J)^{i\phi} e^{(i\phi/2)(1+i\phi)\sigma_J^2} - 1] - \lambda i\phi\mu_J\tau \right. \\ \left. + \frac{i\phi(i\phi + 1)(1 - e^{-\xi_v\tau})}{2\xi_v - [\xi_v - \kappa_v + (1 + i\phi)\rho\sigma_v](1 - e^{-\xi_v\tau})} V(t) \right\}, \end{aligned} \quad (\text{A12})$$

and

$$\begin{aligned} \hat{f}_2 = \exp \left\{ -i\phi \ln[B(t, \tau)] - \frac{\theta_v}{\sigma_v^2} \left[2 \ln \left(1 - \frac{[\xi_v^* - \kappa_v + i\phi\rho\sigma_v](1 - e^{-\xi_v^*\tau})}{2\xi_v^*} \right) \right] \right. \\ \left. - \frac{\theta_v}{\sigma_v^2} [\xi_v^* - \kappa_v + i\phi\rho\sigma_v]\tau + i\phi \ln[S(t)] \right. \\ \left. + \lambda\tau[(1 + \mu_J)^{i\phi} e^{(i\phi/2)(i\phi-1)\sigma_J^2} - 1] - \lambda i\phi\mu_J\tau \right. \\ \left. + \frac{i\phi(i\phi - 1)(1 - e^{-\xi_v^*\tau})}{2\xi_v^* - [\xi_v^* - \kappa_v + i\phi\rho\sigma_v](1 - e^{-\xi_v^*\tau})} V(t) \right\}, \end{aligned} \quad (\text{A13})$$

The characteristic functions for the SV and the SVSI models can be obtained by respectively setting $\lambda = 0$ in (A12)–(A13) and in (A10)–(A11). Q.E.D.

Expressions for the gamma measures. The various second-order partial derivatives of the call price in equation (8), which are commonly referred to as gamma measures, are given below for the SVSI-J model:

$$\Gamma_S \equiv \frac{\partial^2 C(t, \tau)}{\partial S^2} = \frac{\partial \Pi_1}{\partial S} = \frac{1}{\pi} \int_0^\infty \text{Re} \left[(i\phi)^{-1} e^{-i\phi \ln[K]} f_1 \frac{i\phi}{S} \right] d\phi > 0. \quad (\text{A14})$$

$$\Gamma_V \equiv \frac{\partial^2 C(t, \tau)}{\partial V^2} = S(t) \frac{\partial^2 \Pi_1}{\partial V^2} - KB(t, \tau) \frac{\partial^2 \Pi_2}{\partial V^2} \quad (\text{A15})$$

$$\Gamma_R \equiv \frac{\partial^2 C(t, \tau)}{\partial R^2} = S(t) \frac{\partial^2 \Pi_1}{\partial R^2} - KB(t, \tau) \left\{ \frac{\partial^2 \Pi_2}{\partial R^2} - 2\varrho(\tau) \frac{\partial \Pi_2}{\partial R} + \varrho^2(\tau) \Pi_2 \right\}. \tag{A16}$$

$$\Gamma_{S,V} \equiv \frac{\partial^2 C(t, \tau)}{\partial S \partial V} = \frac{\partial \Pi_1}{\partial V} = \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[(i\phi)^{-1} e^{-i\phi \ln[K]} \frac{\partial f_1}{\partial V} \right] d\phi. \tag{A17}$$

where for $g = V, R$ and $j = 1, 2$:

$$\frac{\partial^2 \Pi_j}{\partial g^2} = \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[(i\phi)^{-1} e^{-i\phi \ln[K]} \frac{\partial^2 f_j}{\partial g^2} \right] d\phi. \tag{A18}$$

Q.E.D.

Proof of the Minimum Variance Hedge in equation (21). To derive a closed-form expression for $\operatorname{Cov}_t[dS(t), dC(t, \tau)]$, one needs to evaluate $E_t[JC(S(1 + J), R, V)]$. With the aid of equations (3) and (8), derive this conditional expectation directly, which upon simplification results in equation (21) with

$$\begin{aligned} \Lambda_1(t) &= \frac{S(t)}{2} [\mu_J + \mu_J^2 + (e^{\sigma_J^2} - 1)(1 + \mu_J)^2] \\ &\quad + \frac{S(t)}{\pi} \int_0^\infty \operatorname{Re} \left[e^{-i\phi \ln[K]} \frac{f_1(t, \tau) \bar{m}_1}{i\phi} \right] d\phi \end{aligned} \tag{A19}$$

$$\Lambda_2(t) = \frac{KB(t, \tau) \mu_J}{2} + \frac{KB(t, \tau)}{\pi} \int_0^\infty \operatorname{Re} \left[e^{-i\phi \ln[K]} \frac{f_2(t, \tau) \bar{m}_2}{i\phi} \right] d\phi \tag{A20}$$

where

$$\begin{aligned} \bar{m}_1 &\equiv \exp \left[(2 + i\phi) \left(\ln[1 + \mu_J] - \frac{1}{2} \sigma_J^2 \right) + \frac{1}{2} (2 + i\phi)^2 \sigma_J^2 \right] \\ &\quad - \exp \left[(1 + i\phi) \left(\ln[1 + \mu_J] - \frac{1}{2} \sigma_J^2 \right) + \frac{1}{2} (1 + i\phi)^2 \sigma_J^2 \right] \\ \bar{m}_2 &\equiv \exp \left[(1 + i\phi) \left(\ln[1 + \mu_J] - \frac{1}{2} \sigma_J^2 \right) + \frac{1}{2} (1 + i\phi)^2 \sigma_J^2 \right] \\ &\quad - \exp \left[i\phi \left(\ln[1 + \mu_J] - \frac{1}{2} \sigma_J^2 \right) - \frac{1}{2} \phi^2 \sigma_J^2 \right] \end{aligned}$$

Q.E.D.

REFERENCES

- Ait-Sahalia, Yacine, and Andrew Lo, 1996, Nonparametric estimation of state-price densities implicit in financial asset prices, Working paper, University of Chicago and MIT.
- Amin, Kaushik, and Robert Jarrow, 1992, Pricing options on risky assets in a stochastic interest rate economy, *Mathematical Finance* 2, 217–237.
- Amin, Kaushik, and Victor Ng, 1993, Option valuation with systematic stochastic volatility, *Journal of Finance* 48, 881–910.
- Bailey, Warren, and René Stulz, 1989, The pricing of stock index options in a general equilibrium model, *Journal of Financial and Quantitative Analysis* 24, 1–12.
- Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Pricing and hedging long-term options, Working paper, University of Maryland, Pennsylvania State University, and The Ohio State University.
- Bakshi, Gurdip, and Zhiwu Chen, 1997a, An alternative valuation model for contingent claims, *Journal of Financial Economics* 44, 123–165.
- Bakshi, Gurdip, and Zhiwu Chen, 1997b, Equilibrium valuation of foreign exchange claims, *Journal of Finance* 52, 799–826.
- Barone-Adesi, Giovanni, and Robert Whaley, 1987, Efficient analytic approximation of American option values, *Journal of Finance* 42, 301–320.
- Bates, David, 1991, The crash of 87: Was it expected? The evidence from options markets, *Journal of Finance* 46, 1009–1044.
- Bates, David, 1996a, Jumps and stochastic volatility: Exchange rate processes implicit in Deutschemark options, *Review of Financial Studies* 9, 69–108.
- Bates, David, 1996b, Testing option pricing models, in G. S. Maddala and C. R. Rao, eds: *Handbook of Statistics, Vol. 15: Statistical Methods in Finance* (North Holland, Amsterdam), pp. 567–611.
- Bates, David, 1996c, Post-87 crash fears in S&P 500 futures options, Working paper, University of Iowa.
- Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–659.
- Bodurtha, James, and George Courtadon, 1987, Tests of the American option pricing model on foreign currency-options markets, *Journal of Financial and Quantitative Analysis* 22, 153–168.
- Broadie, Mark, Jerome Detemple, Eric Ghysels, and Oliver Torres, 1996, American options with stochastic volatility, Working paper, Pennsylvania State University.
- Cao, Charles, 1993, Pricing foreign currency options with stochastic volatility, Ph.D. dissertation, University of Chicago.
- Chan, K. C., Andrew Karolyi, Francis Longstaff, and Anthony Sanders, 1992, An empirical comparison of alternative models of the short-term interest rate, *Journal of Finance* 47, 1209–1227.
- Cox, John, Jonathan Ingersoll, and Stephen Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–408.
- Cox, John, and Stephen Ross, 1976, The valuation of options for alternative stochastic processes, *Journal of Financial Economics* 3, 145–166.
- Dumas, Bernard, Jeff Fleming, and Robert Whaley, 1995, Implied volatility smiles: Empirical tests, Working paper, Duke University and Rice University.
- Figlewski, Stephen, 1989, Option arbitrage in imperfect markets, *Journal of Finance* 44, 1289–1311.
- Galai, Dan, 1983, The components of the return from hedging options against stocks, *Journal of Business* 56, 45–54.
- Harvey, Campbell, and Robert Whaley, 1992, Market volatility and the efficiency of the S&P 100 index option market, *Journal of Financial Economics* 31, 43–73.
- Heston, Steven, 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6, 327–343.

- Hull, John, and Alan White, 1987a, The pricing of options with stochastic volatilities, *Journal of Finance* 42, 281–300.
- Hull, John, and Alan White, 1987b, Hedging the risks from writing foreign currency options, *Journal of International Money and Finance* 6, 131–152.
- Longstaff, Francis, 1995, Option pricing and the martingale restriction, *Review of Financial Studies* 8, 1091–1124.
- Madan, Dilip, and Eric Chang, 1996, The VG option pricing model, Working paper, University of Maryland and Georgia Institute of Technology.
- Melino, Angelo, and Stuart Turnbull, 1990, Pricing foreign currency options with stochastic volatility, *Journal of Econometrics* 45, 239–265.
- Melino, Angelo, and Stuart Turnbull, 1995, Misspecification and the pricing and hedging of long-term foreign currency options, *Journal of International Money and Finance* 14, 373–393.
- Merton, Robert, 1973, Theory of rational option pricing, *Bell Journal of Economics* 4, 141–183.
- Merton, Robert, 1976, Option pricing when the underlying stock returns are discontinuous, *Journal of Financial Economics* 4, 125–144.
- Nandi, Saikat, 1996, Pricing and hedging index options under stochastic volatility, Working paper, Federal Reserve Bank of Atlanta.
- Nelson, Daniel, 1991, Conditional heteroskedasticity in asset returns, *Econometrica* 59, 347–370.
- Ross, Stephen, 1995, Hedging long-run commitments: Exercises in incomplete market pricing, Working paper, Yale School of Management.
- Rubinstein, Mark, 1985, Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE options classes from August 23, 1976 through August 31, 1978, *Journal of Finance* 40, 455–480.
- Rubinstein, Mark, 1994, Implied binomial trees, *Journal of Finance* 49, 771–818.
- Scott, Louis, 1987, Option pricing when the variance changes randomly: Theory, estimators, and applications, *Journal of Financial and Quantitative Analysis* 22, 419–438.
- Scott, Louis, 1997, Pricing stock options in a jump-diffusion model with stochastic volatility and interest rates: Application of Fourier inversion methods, *Mathematical Finance*, Forthcoming.
- Stein, Elias, and Jeremy Stein, 1991, Stock price distributions with stochastic volatility, *Review of Financial Studies* 4, 727–752.
- White, Halbert, 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 48, 817–838.
- Wiggins, James, 1987, Option values under stochastic volatilities, *Journal of Financial Economics* 19, 351–372.