Information Content of Cross-Sectional Option Prices: A Comparison of Alternative Currency Option Pricing Models on the Japanese Yen

BRICE DUPOYET

This article implements a currency option pricing model for the general case of stochastic volatility, stochastic interest rates, and jumps in an attempt to reconcile levels of risk-neutral skewness and kurtosis with observed option prices on the Japanese yen and to analyze the information content of the cross section of option prices by investigating the hedging and pricing performance of various currency option pricing models. The study makes use of both a method of moments and a more traditional generalized-least-squares (GLS) estimation technique, taking advantage of the fact that methods of moments do not specifically require the use of cross-sectional option prices, whereas GLS does. Results centered around the Asia economic crisis of 1997 and 1998 indicate that the cross section of option prices surprisingly does not appear to contain superior information as the two estimation techniques yield relatively similar results once idiosyncratic differences between them are acknowledged. Extensions of

For correspondence, Brice Dupoyet, Assistant Professor of Finance, Florida International University, School of Business, 11200 SW 8th Street, RB 209A, Miami, Florida 33199; e-mail: dupoyetb@fiu.edu

Received June 2004; Accepted March 2005

■ *Brice Dupoyet is Assistant Professor of Finance at Florida International University in Miami, Florida.*

The Journal of Futures Markets, Vol. 26, No. 1, 33–59 (2006)

Published online in Wiley InterScience (www.interscience.wiley.com).
DOI: 10.1002/fut.20177

the G. Bakshi, C. Cao, and Z. Chen (1997) results to currencies are also provided. Jrl Fut Mark 26:33–59, 2006

INTRODUCTION

Varying degrees of skewness and kurtosis present in a security's distribution of returns are known to affect the price of an option written on that security, as modeled by Bakshi, Cao, and Chen (1997); Bates (2000); Duffie, Pan, and Singleton (2000); or Pan (2002) and formalized more recently by Bakshi, Kapadia, and Madan (2003). By deriving and implementing in this article a general option-pricing model allowing for stochastic volatility, jumps, and stochastic interest rates in the case of a currency, pricing implications on Japanese yen options as well as the value of the information content of the cross section of option prices are investigated for purposes of parameter estimation, pricing, and hedging. Although it is accepted that incorporating stochastic volatility and return jumps in option pricing models reduces hedging and pricing errors—as demonstrated by Bakshi et al. (1997), Bates (2000), and Pan (2002)—and that incorporating volatility jumps provides a better fit yet—as shown by Eraker, Johannes, and Polson (2003) and Bakshi and Cao (2004)—it is, however, less established what parameter-estimation techniques and option data sets are the most relevant for the calibration of complex options models and the subsequent pricing and hedging with these models. This study is an attempt to determine the value of the information content of the cross section of option prices by implementing two estimation techniques that make use of radically different information sets. By using options on the Japanese yen, the study also extends Bakshi, Cao, and Chen (1997) results on equity options to the case of a currency.

When dealing with the estimation of parameters associated with option-pricing models, two important classes of estimation methods are found in the literature: moment-based methods and least-squares minimization techniques. These two classes vary considerably in nature. Moment-based methods usually do not make use of the whole cross-sectional data set of options available, but estimate parameters in a time-series fashion by relying on incorporating the dynamics of the specified process in the estimation itself. Least-squares minimization procedures instead estimate parameters by using the full cross section of option prices at a given point in time, thereby ensuring a good fit at that date but somewhat preventing forecasting more than one period ahead as the dynamics of the process are not part of the estimation. For purposes of this study, one estimation technique from each class is selected in an attempt to determine whether one class fares better in terms of parameter estimation and subsequent option pricing and hedging abilities. The moment-based method selected is the Pan (2002) implied-state generalized method of moments (IS-GMM), to be compared against a more traditional generalized-least-squares (GLS) estimation as seen in Bates (1996); Bakshi et al. (1997); and Dumas, Fleming, and Whaley (1998) to only name a few.

The generalized-least-squares estimation methodology involving minimizing a sum of squared errors over an option data set cross sectionally is relatively simple to implement and is therefore used by academics and practitioners alike, as seen in Bates (1996); Bakshi et al. (1997); or Dumas, Fleming, and Whaley (1998). The GLS estimation is generally performed daily with parameters obtained by averaging the results over the sample period. This methodology explores the information in the cross section of option prices. The second estimation technique implemented, the implied-state generalized method of moments of Pan (2002), allows for a latent variable such as the volatility to be inferred from the data given the value of the parameters being estimated. IS-GMM investigates the time series properties of the data holding the parameters fixed over time. The estimation is also only making use of at-the-money options data. It thus uses a different data set than the generalized-least-squares procedure, because the GLS estimation makes use of the full cross sectional data set. It is also important to note that the implied-state GMM technique is computer intensive. Although GLS relies on a nonlinear least-squares minimization approach, IS-GMM derives an implied-volatility time series at each loop of the estimation, a time-consuming exercise. The mathematical derivation of the various moments can be tedious as well. One may ask whether the benefit of IS-GMM truly warrants the added computational costs.

This study also tests whether results by Bakshi et al. (1997) hold for a currency. An important difference between a currency and an equity security is the fact that although both currencies and stocks depend on a myriad of financial and economic factors, a currency's "natural" path can be altered through the use of government intervention. Central banks are known to employ foreign-exchange reserves and fiscal and monetary policies as a means of influencing the currency's price. Hence even a so-called free-floating currency often goes through managed regimes where the currency is not only the product of various macroeconomic factors but also the result of a government's actions. This begs the question of whether a currency is thus fundamentally different from a stock or an equity index and whether its inner structure differs significantly from that of a stock, particularly in times of crisis when the government is likely to intervene. The answer to this question has important implications regarding the hedging of foreign-exchange risk by firms as well as the pricing of derivatives allowing this hedging.

Results by Duffie et al. (2000) on affine processes can be used to produce a closed-form currency option-pricing model featuring stochastic volatility, stochastic interest rates, and jumps with all features embedded in one general model. This solution allows for convenient comparison of several submodels nested within one another and thus for the implementation of a study on currency options in the spirit of Bakshi et al. (1997) equity option study. Because currencies are governed by a somewhat different set of dynamics than equities, establishing whether results for options on an equity index also hold for options on a currency is important. The currency selected is the Japanese yen, and the period studied is from 1996 to 1999, a window of time surrounding the Asian economic crisis of 1997 along with known central bank interventions.

The models considered in this study are the Black-Scholes model (Black & Scholes, 1973), the stochastic-volatility model, the stochastic-volatility and stochastic interest-rates model, and finally the stochastic volatility with jumps model. Henceforth these models will be referred to as BS, SV, SVSI, and SVJ, respectively. General findings in this study are first that the Bakshi et al. (1997) results for equity options do hold overall for a currency. The stochastic volatility feature provides the largest incremental pricing and hedging improvement over the Black-Scholes benchmark. Including jumps into the model does improve the fit in terms of pricing or implied volatility smile, but contributes little in terms of hedging. Finally, the stochastic interest-rates feature improves the pricing only in the case of in-the-money long-term options, whereas its effect is insignificant in the hedging exercise. Perhaps more surprisingly, the IS-GMM and GLS estimation methods seem to yield pricing and hedging results of comparable levels. The IS-GMM estimation actually produces slightly poorer fits and hedging errors, and although it might thus be tempting to conclude that the information contained in the cross section of option prices—used by GLS and not by IS-GMM—is valuable, this apparent slight superiority is actually marginal and most likely the result of the recalibration of parameters granted to GLS whereas the IS-GMM parameters remain unchanged. The somewhat surprising conclusion is thus that the information contained in the cross section of option prices does not seem to be of the highest importance for purposes of pricing and hedging.

The rest of the article proceeds as follows. Next the model and the solution to the pricing problem are given. Then the options data on the Japanese yen are described. In the following sections the two estimation techniques implemented in the article are presented, results from the estimation of the parameters and pricing exercises are described, and the hedging abilities of the models are presented. Finally, a conclusion is given.

THE MODEL

The spot value for the currency is denoted by $S(t)$ and is assumed to follow a mixed jump-diffusion process with stochastic volatility. The domestic and foreign instantaneous risk-free rates are denoted by $R(t)$ and $R_f(t)$, and are assumed to follow a Cox, Ingersoll, and Ross (1985) process. The source of Brownian risk associated with the currency is denoted by $d\omega_S(t)$, a standard Brownian motion. The instantaneous volatility of the currency process is denoted by $V(t)$ and is also assumed to follow a Cox et al. (1985) type of motion, with its Brownian source of risk correlated with the currency returns' diffusion component. The assumptions of Bakshi et al. (1997) are adapted to the case of currency options and the data-generating processes under the risk-neutral probability measure Q are thus assumed to follow:

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

$$dV(t) = [\theta_V - \kappa_V V(t)]\,dt + \sigma_V\sqrt{V(t)}\,d\omega_V(t) \quad (2)$$

$$dR(t) = [\theta_R - \kappa_R R(t)]\,dt + \sigma_R\sqrt{R(t)}\,d\omega_R(t) \quad (3)$$

$$dR_f(t) = [\theta_{R_f} - \kappa_{R_f} R_f(t)]\,dt + \sigma_{R_f}\sqrt{R_f(t)}\,d\omega_{R_f}(t) \quad (4)$$

with $Cov_t(d\omega_S(t), d\omega_V(t)) = \rho\,dt$ and where

$R(t)$ and $R_f(t)$ are the time-t instantaneous domestic and foreign interest rates.

$V(t)$ is the diffusion component of return variance, conditional on no jump.

λ is the frequency of jumps per year and $J(t)$ is the percentage jump size conditional on a jump occurring assumed to be lognormally, identically, and independently distributed over time, with unconditional mean μ_J. Conditional on a jump, the spot value of the currency instantly goes from S to $S[1 + J(t)]$.

The binary variable $q(t)$ is a Poisson jump counter with intensity λ with

$$\Pr\{dq(t) = 1\} = \lambda\, dt \quad \text{and} \quad \Pr\{dq(t) = 0\} = 1 - \lambda\, dt$$

The parameters 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 $V(t)$. The variables $q(t)$ and $J(t)$ are uncorrelated with each other or with the Brownian motions. The solution shown below stems from the work of Duffie et al. (2000) on affine jump transforms and is stated for the most general case: stochastic volatility, stochastic interest rates, and jumps, as found in Pan (2002). The various nested models subsequently compared are Black-Scholes (BS), stochastic volatility (SV), stochastic volatility and stochastic interest rates (SVSI), and stochastic volatility with jumps (SVJ).

The call option price can be derived as

$$C(S_t, K, T) = P_1 - KP_2 \tag{5}$$

where

$$\begin{aligned} P_1 &= \frac{\psi(1, X_t, 0, T)}{2} - \frac{1}{\pi}\int_0^\infty \frac{\operatorname{Im}[\psi(1 - iv, X_t, 0, T)e^{iv \ln K}]}{v}\, dv \\ P_2 &= \frac{\psi(0, X_t, 0, T)}{2} - \frac{1}{\pi}\int_0^\infty \frac{\operatorname{Im}[\psi(-iv, X_t, 0, T)e^{iv \ln K}]}{v}\, dv \end{aligned} \tag{6}$$

$$\psi(u, x, t, T) = E_t^Q\left(\exp\left[-\int_t^T R(X_s)\, ds\right] \exp(u \cdot X_T)\right) = \exp[\alpha(t) + \beta(t) \cdot x] \tag{7}$$

where details can be found in Pan (2002).

DATA DESCRIPTION

This study uses daily prices for European and American currency call options written on the Japanese yen from March 29, 1996 to December 31, 1999 traded on the Philadelphia exchange. An important advantage of the Philadelphia exchange is that the currency spot and option price are simultaneously recorded, eliminating any synchronicity issues that may arise with OTC data. One drawback, however, is that options traded on the exchange are not as liquid as OTC records. In order to minimize this issue, both European options and American options are used in the sample.

When the foreign risk-free rate is significantly lower than the domestic risk-free rate, an American call option on the currency has approximately the same value as a European call option on that same currency. In the case of the Japanese yen, the interest-rate differential is high. This allows the inclusion of both American and European call options, increasing the sample size and the reliability of the results. Note that this approximation would not hold in the case of put options, as put options may be exercised at low moneyness levels; thus put options are not included in this study. American options constitute 75% of the sample against 25% for European options. It is therefore obvious that the inclusion of American options in the study is paramount to the reliability of the results. To obtain an estimate of the size of the errors associated with including American call options in the sample, the theoretical percentage price differences between American and European options are computed. The variables S, R, R_f, and V are assigned their sample average, and K and τ are given a series of values yielding a grid of various moneyness and time-to-maturity categories. The theoretical errors associated with including American options in the sample are summarized by moneyness and maturity in Table I.

Table I demonstrates that given the high interest-rate differential, the pricing error due to the inclusion of American options is never more than 0.011%. The results in this article should therefore not be biased by the approximation.

The data set is then placed through a series of filters. Options with maturity less than 12 days are discarded to avoid near-maturity pricing

TABLE I

Summary Statistics on Theoretical Approximation Errors

	Days to expiration	
S/K	*<60*	*60–180*
<0.94	.0113%	.0037%
0.94–0.97	.0112%	.0068%
0.97–1.00	.0098%	.0080%
1.00–1.03	.0023%	.0020%
>1.03	.0010 %	.0007%

Note. Theoretical differences in prices between American and European options of the same strike and maturity are computed and expressed as a percentage of the American option price. The spot, domestic rate, foreign rate, and volatility variables S, R, R_f, and V are assigned their sample average, and the exercise price and time-to-maturity K and τ are given a series of values yielding a grid of various moneyness and time-to-maturity categories. Errors associated with including American options in the sample are summarized by moneyness and maturity.

anomalies. Options violating the European boundary condition are discarded, as well as options with prices lower than 0.0000002 cents. These criteria removed 63 options, leaving a total of 11,109 call options for the April 1996–December 1999 period. Finally, options whose maturity exceeds 180 days are removed for lack of liquidity in this category. The final total number of options remaining is 9,806.

Figure 1 describes the objective distribution of returns on the Japanese yen, indicating a degree of positive skewness and large levels of kurtosis. The distribution thus displays fat tails, but is fairly symmetric, as is often the case with objective distributions.

A brief description of risk-neutral distributional properties is warranted as well. To obtain the levels of risk-neutral skewness and kurtosis, the methodology described in Bakshi et al. (2003) with the third and fourth moments spanned by using out-of-the-money calls and puts is adopted. The levels of risk-neutral skewness and kurtosis can be significantly different from their objective counterparts. Bakshi et al. (2003)

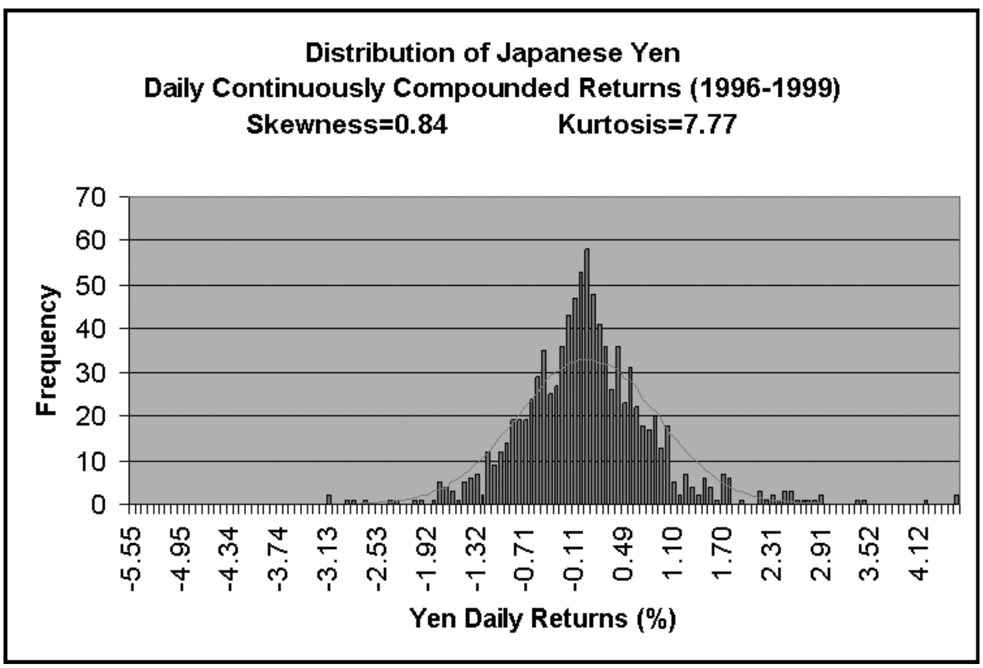

FIGURE 1
Distribution of daily rates of return on the Japanese yen. Actual daily continuously compounded rates of return on the Japanese yen are plotted in a histogram, whereas a normal distribution exhibiting identical first and second moments is plotted as a smooth line in order to demonstrate the contrast between the two. The distribution of the actual rates of return exhibits both skewness and kurtosis.

show that risk-aversion causes the risk-neutral density to have a negative skew provided that the kurtosis of the physical distribution is more than 3, that a negative skew in the physical distribution will also create a left skew in the risk-neutral distribution, and finally that the higher the physical measure level of kurtosis, the deeper the left skew in the risk-neutral distribution. Attention is restricted to short-term options, and risk-neutral levels of skewness and kurtosis of -0.39 and 6.54, respectively, are obtained. These risk-neutral moments are indeed different from their physical counterparts, and it is important to note that despite the positive physical skewness of the objective distribution, negative levels of risk-neutral skewness are observed, most likely the result of high physical kurtosis and risk aversion.

The option-pricing model on the FX market will attempt to incorporate these risk-neutral levels of skewness and kurtosis through features such as stochastic volatility and jumps. In a stochastic-volatility model, the volatility variation controls the kurtosis of the spot return, and the correlation between the volatility and spot returns produces skewness. Jumps can also model skewness and fat tails, although one should note that stochastic volatility and jump-diffusion components incorporate skewness and kurtosis in the underlying security distribution differently: the effects of jump diffusion are short term (3 months or less), whereas the effects of stochastic volatility are long term.

Before describing the estimation procedures, a description of how much the Black-Scholes model misprices the options is in order. The way this is shown is by backing out the Black-Scholes implied volatility for each option price in the sample and equally weighting the implied volatilities of all call options in each given moneyness-maturity class, yielding a matrix of average implied volatilities with values reported in Table II along with the number of options in each category. Table III reports average call prices in cents along with average trade size in each category.

Table II is graphically translated in Figure 2, where it can be seen that the Black-Scholes implied volatility exhibits a very strong U-shape or volatility smile as the call option goes from being deep in-the-money to deep out-of-the money. The smile is most pronounced for short-term options, that is, options with a maturity of less than 60 days. The short-term options thus seem to be the most mispriced, as also found by Bakshi et al. (1997) in the case of S&P 500 index options.

The main challenge faced by option pricing models is therefore to explain the behavior of short-term deep out-of-the-money and short-term deep in-the-money call options, as they appear to be the most mispriced

TABLE II
Implied Volatility and Number of Options by Moneyness and Maturity

	Call options April 1996–December 1999 days to expiration	
S/K	*<60*	*60–180*
<0.94	19.03 {204}	15.12 {895}
0.94–0.97	14.04 {958}	13.69 {1436}
0.97–1.00	11.64 {2926}	12.45 {1184}
1.00–1.03	11.92 {908}	13.14 {549}
>1.03	20.53 {374}	16.47 {372}

Note. The implied volatility is obtained by inverting the Black-Scholes model separately for each call option contract on the Japanese yen for the period from April 1996 to December 1999, for a total of 9806 records. The implied volatilities are then averaged within each moneyness-maturity class. Moneyness is represented by S/K, the ratio of the spot exchange rate of the Japanese yen currency over the exercise price. Two different maturities categories are plotted: less than 60 days to maturity and between 60 and 180 days. The number of options by moneyness/maturity class is shown below each implied volatility value.

TABLE III
Average Price and Average Trade Size by Moneyness and Maturity

	Call options April 1996–December 1999 days to expiration	
S/K	*<60*	*60–180*
<0.94	0.0027 (44)	0.0086 (38)
0.94–0.97	0.0048 (23)	0.0122 (29)
0.97–1.00	0.0087 (25)	0.0194 (20)
1.00–1.03	0.0193 (11)	0.0341 (11)
>1.03	0.0669 (16)	0.0769 (13)

Note. Average call prices are given in cents, with average trade size below, for two different maturities: less than 60 days to maturity and between 60 and 180 days.

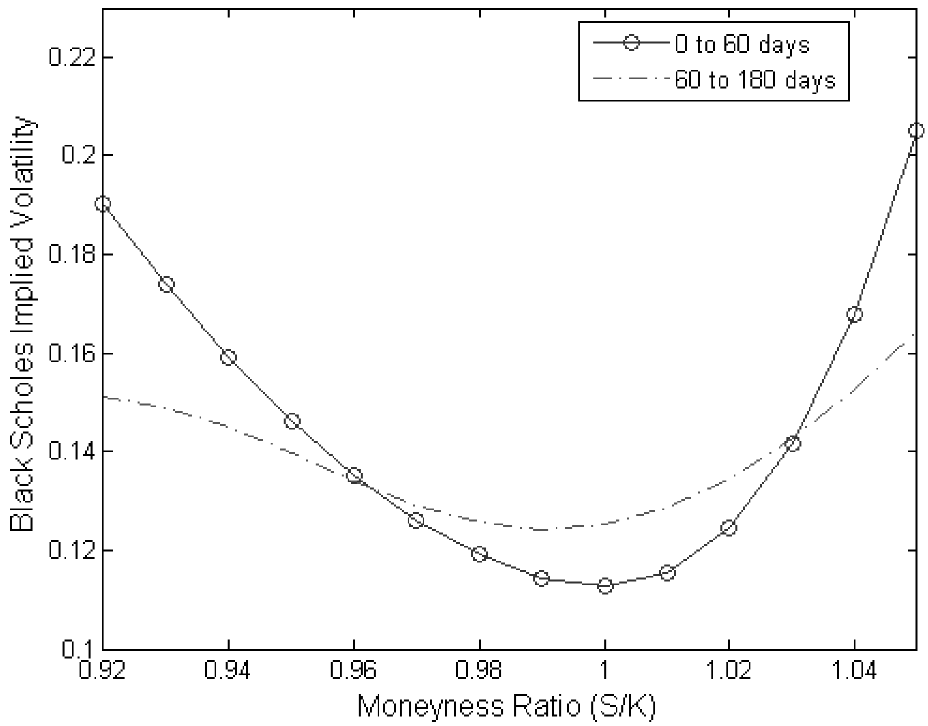

FIGURE 2
Black-Scholes implied volatility smiles. The implied volatility is obtained by inverting the Black-Scholes model separately for each call option contract on the Japanese yen for the period April 1996 to December 1999. The implied volatilities are then averaged within each moneyness-maturity class. Moneyness is represented by S/K, the ratio of the spot exchange rate of the Japanese yen currency over the exercise price. Two different maturities categories are plotted: less than 60 days to maturity and between 60 and 180 days.

category. For options with longer term to expiration, that is, options with more than 60 days to maturity, the smile is not as pronounced; it is only for deep in-the-money call options that the extreme values are found.

Next comes the issue of what short-term rate to use to proxy for the instantaneous risk-free rate. In order to replicate the instantaneous interest rate, an overnight rate should theoretically be used. However, such a short period can give rise to unwanted issues such as the second Wednesday settlement effect in the Federal Funds market. To avoid this problem Ait-Sahalia (1996) selects the 7-day eurodollar rate. The issue of what rate to use is tackled by Chapman, Long, and Pearson (1999), who show that in an affine bond price model, accurate estimates of the drift and diffusion of the short rate process can be obtained, even using proxies with maturities as long as 3 months. Constant-maturity 3-month Treasury bill rates are therefore used along with 3-month Japanese rates collected from DataStream as proxies for the domestic and foreign short rates.

Following Pan (2002), the parameters associated with the risk-free rates are estimated separately and later used as the true parameters when implementing and estimating the option models. The interest-rate parameters are obtained by maximum-likelihood estimation with the use of the known transition density function.

ESTIMATION METHODOLOGY

Estimations are performed with the use of both the implied-state generalized method-of-moments method and the generalized-least-squares technique. In the implied-state GMM setting, the term *implied state* stands for the fact that at each step of the GMM procedure, implied volatilities for each day in the sample are inferred with the use of the model being estimated, observed data, and the current set of parameters. The volatility used at a given iteration is therefore not the true volatility, but an estimate based on the values taken by the parameters at that specific stage of the estimation. As demonstrated by Pan (2002), as the parameters eventually converge to their true (risk-neutral) values the implied volatility time series, in turns, reaches its "true" value.

The moments needed to implement the IS-GMM estimation are derived from the data-generating process for the Japanese yen, differing from its Equation (1–2) risk-neutral counterpart and given by

$$dS(t) = S(t)[R(t) - R_f(t) - \lambda\mu_J^* + \eta^S V(t)]\,dt + S(t)\sqrt{V(t)}\,d\omega_S^*(t) + S(t)J^*(t)\,dq^*(t) \tag{8}$$

$$dV(t) = [\theta_V - \kappa_V^* V(t)]\,dt + \sigma_V\sqrt{V(t)}\,d\omega_V^*(t) \tag{9}$$

where $d\omega_S^*$ and $d\omega_V^*$ are standard Brownian motions under the objective measure with correlation level ρ, where $\kappa_V^* = \kappa_V + \eta^V$ with η^V reflecting the additional volatility risk premium, and where η^S reflects the risk premium for Brownian returns.

The moment-generating function used is the joint conditional moment between the continuously compounded return and the volatility. With $y_t = \Delta\ln(S_t) - \int_{t-1}^{t}(R_u - R_{fu})\,du$, the time-$t$ conditional joint moment-generating function between y_t and V_t between period t and period $t + 1$ is

$$M_{y,V}(u_y, u_V) = \exp[\alpha(u_y, u_V) + \beta(u_y, u_V)V_t] \tag{10}$$

where details for $\alpha(\cdot)$ and $\beta(\cdot)$ can be found in Pan (2002). The moments subsequently used in the GMM estimation are derived by taking derivatives of the moment-generating function and are obtained as

$$E_t(y_{t+1}^i V_{t+1}^j) = \left.\frac{\partial^{i+j} M_{y,V}(u_y, u_V)}{\partial^i u_y\, \partial^j u_V}\right|_{u_y=0,\, u_V=0} \qquad i, j \in \{0, 1, \ldots\} \tag{11}$$

The computation of the derivatives is tedious work, albeit relatively straightforward with derivations performed in Matlab with the symbolic toolbox. Following Pan (2002), the implied volatility is estimated with the use of short-term ATM options, for two reasons: Bakshi et al. (1997) graphically show that most models' volatility smiles seem to intersect near-the-money. One may thus conjecture that ATM options reflect the true volatility. ATM options also typically tend to be the most liquid options, which helps guarantee that the implied volatility can always be backed out for each day in the sample.

In order to implement IS-GMM, the moment conditions are given by $E_t(g_{t+1}) = 0$ with the vector of errors defined as

$$g_{t+1} = [g_{t+1}^{y1}, g_{t+1}^{y2}, g_{t+1}^{y3}, g_{t+1}^{y4}, g_{t+1}^{V1}, g_{t+1}^{V2}, g_{t+1}^{yV}]^T \tag{12}$$

and where

$$\begin{aligned} g_{t+1}^{y1} &= y_{t+1} - E_t(y_{t+1}), & g_{t+1}^{y2} &= y_{t+1}^2 - E_t(y_{t+1}^2) \\ g_{t+1}^{y3} &= y_{t+1}^3 - E_t(y_{t+1}^3), & g_{t+1}^{y4} &= y_{t+1}^4 - E_t(y_{t+1}^4) \\ g_{t+1}^{V1} &= V_{t+1} - E_t(V_{t+1}), & g_{t+1}^{V2} &= V_{t+1}^2 - E_t(V_{t+1}^2) \\ g_{t+1}^{yV} &= y_{t+1}V_{t+1} - E_t(y_{t+1}V_{t+1}) \end{aligned} \tag{13}$$

Armed with these moment conditions, the estimation of the parameters follows the Hansen (1982) generalized method of moments. The Newey-West correction is applied to prevent the possibility of correlation and autocorrelation, with the lag set to 20. Bartlett weights are used in order to ensure that the resulting matrix is positive semidefinite.

The generalized-least-squares procedure is implemented as follows. For a given day a cross-section of option prices of a given maturity range is recorded. Let $C_n(t, S, \tau_n, R, R_f, K_n)$ be the observed option price on the nth option in the sample on day t, let $\hat{C}(t, S, \tau_n, R, R_f, K_n)$ be the model-driven option price given a set of parameters Θ and a volatility level V, let K_n and τ_n be the exercise price and time-to-maturity of the nth option, let R and

R_f be, respectively, the domestic and foreign risk-free rates on day t and let S be the value of the spot index on that same day. The generalized-least-squares procedure is performed for each day t by minimizing

$$\underset{V,\Theta}{\mathrm{Min}} \sum_{n=1}^{N} [C_n(t, S, \tau_n, R, R_f, K_n) - \hat{C}_n(t, S, \tau_n, R, R_f, K_n)]^2 \qquad (14)$$

The volatility V is estimated along with the rest of the model parameters Θ and is therefore treated as such. Following Bakshi et al. (1997), the sets of parameters obtained for each day in the sample are finally averaged over time in order to yield a unique set of model-implied estimated parameters.

PARAMETERS ESTIMATION RESULTS AND OUT-OF-SAMPLE FIT

The methodologies discussed above are implemented on both short-term options (less than 60 days to maturity) and medium-to-long-term options (between 60 and 180 days to maturity), with IS-GMM and GLS estimates reported in Tables IV and V. It is important to keep in mind

TABLE IV
Implied-State GMM Implied Parameters

	Implied-state GMM							
	Short-term options				*Mid/Long-term options*			
Parameters	*BS*	*SV*	*SVSI*	*SVJ*	*BS*	*SV*	*SVSI*	*SVJ*
κ_V		3.86	3.85	2.75		6.19	6.17	3.02
		(1.99)	(1.98)	(1.52)		(3.65)	(3.59)	(1.70)
θ_V		0.05	0.05	0.03		0.06	0.06	0.03
		(4.85)	(4.71)	(3.55)		(4.78)	(4.75)	(2.72)
σ_V		0.30	0.28	0.24		0.21	0.21	0.23
		(3.86)	(3.88)	(2.35)		(2.30)	(2.20)	(2.70)
ρ		−0.14	−0.14	−0.09		−0.12	−0.13	−0.08
		(2.47)	(2.39)	(2.09)		(2.39)	(2.38)	(2.15)
λ				0.24				0.15
				(1.85)				(1.87)
μ_J				−0.28				−0.06
				(2.01)				(1.99)
σ_J				0.16				0.17
				(3.27)				(3.05)
Implied Velocity (%)	12.92	13.86	13.84	13.02	13.78	15.59	15.60	13.51

Note. The IS-GMM methodology is applied to the period from April 1996 to December 1999 on the set of at-the-money options. The numbers in parentheses are the *t* statistics. The final parameter values yield a true implied volatility time series (calculated from the ATM short-term options time series) as a by-product. Its average is reported at the bottom. The models are Black-Scholes (BS), stochastic volatility (SV), stochastic-volatility stochastic interest rates (SVSI), and stochastic volatility with jumps (SVJ).

TABLE V
Generalized-Least-Squares (GLS) Implied Parameters

	Generalized Least Squares							
	Short-term options				*Mid/Long-term options*			
Parameters	*BS*	*SV*	*SVSI*	*SVJ*	*BS*	*SV*	*SVSI*	*SVJ*
κ_V		5.18	5.16	4.36		5.02	5.01	5.22
		(2.15)	(2.09)	(2.26)		(2.24)	(2.23)	(2.36)
θ_V		0.11	0.10	0.10		0.12	0.11	0.09
		(1.76)	(1.79)	(1.87)		(2.12)	(2.13)	(2.00)
σ_V		0.61	0.64	0.33		0.54	0.51	0.40
		(2.08)	(1.99)	(2.37)		(2.51)	(2.46)	(2.17)
ρ		0.04	0.03	−0.02		−0.02	−0.02	0.01
		(2.34)	(2.01)	(−2.04)		(−2.05)	(−2.09)	(2.73)
λ				0.18				0.15
				(−1.88)				(−1.91)
μ_J				−0.49				−0.12
				(−2.72)				(−2.77)
σ_J				0.18				0.18
				(2.75)				(2.11)
Implied Volatility (%)	12.92	13.11	13.15	13.21	13.78	12.53	12.52	11.90

Note. Each day in the sample, model parameters are estimated by minimizing the sum of squared pricing errors between the observed market prices and the model-driven prices. In each cell the parameters' daily averages are reported above the *t* statistic. The models are: Black-Scholes (BS), stochastic volatility (SV), stochastic-volatility stochastic interest rates (SVSI), and stochastic volatility with jumps (SVJ).

that the two estimation techniques differ in nature. Because the GLS procedure estimates parameters over the cross-section of option prices by using the option pricing formula and observed option prices at a given time t, the estimation yields risk-neutral parameters, whereas IS-GMM yields both risk-neutral and objective parameters by including the time series of the process in the estimation, in a panel data fashion.

For purposes of clarity and ease of comparison with GLS results, Table IV thus only reports estimated risk-neutral parameters. However, risk-premium parameters are also obtained as part of the estimation. The Brownian risk premium η^S is equal to 4.2, 2.0, and 4.0 for the SV, SVJ, and SVSI models, respectively, when estimated from short-term options. The short-term options volatility risk premium η^V is equal to 4.8, 2.8, and 2.7, respectively, and the objective mean jump size μ_J^*—specific to the SVJ model—is estimated at 0.02 from short-term options. The Brownian risk premium η^S is equal to 4.5, 2.4, and 4.4 for the SV, SVJ, and SVSI models when estimated from long-term options. The volatility risk premium η^V is equal to 5.3, 3.1, and 5.4, respectively, and the objective mean jump size μ_J^*—specific to the SVJ model—is estimated at 0.03 from long-term options. The risk-neutral and objective mean jump sizes exhibit significant differences, indicating that the fear of a jump impacts the value of the risk-neutral mean jump size in a negative fashion, sending the risk-neutral parameter below its objective counterpart. The volatility risk premium η^V is positive and significantly different from 0 for both short-term and long-term options, with values higher for long-term options, indicating that volatility impacts options of longer maturities more significantly. A noticeable feature is that implied volatility levels are consistent across models. IS-GMM and GLS estimations both lead to implied volatility levels that are never more than 2% apart. Implied volatility levels are also more consistent across models in the case of short-term options than in the case of longer-term options. Within IS-GMM or GLS, the SV and SVSI parameters are close in value, whereas SVJ results tend to lie further apart. Across estimation techniques, however, the parameters are not always consistent. In the case of short-term options, for instance, the correlation coefficient r estimated from the SV model is around 0.04 for GLS, and around -0.14 for IS-GMM. However, because a USD/JPY call option can also be interpreted as a JPY/USD put option, it is not obvious which correlation coefficient is the most realistic.[1] Ultimately, as Bates (1996) points out, different parameters can sometimes produce similar results when it comes to performance, and the real test of quality of these estimated parameters comes in the form of pricing and hedging exercises.

Flattening the Implied Volatility Smile

A presumably misspecified option-pricing model exhibits biases. When implied volatility levels are inferred from observed option prices and plotted against moneyness (S/K) levels, as with Black-Scholes, these biases graphically translate into an implied volatility plot having the shape of a smile or smirk across moneyness levels. A more flexible option model should theoretically display fewer biases and thus a flattened implied volatility smile. This section investigates by how much the smile is reduced when different models and estimation methodologies are implemented. In the interest of space only graphs on short-term options are reported. Short-term options are the most difficult to price, as revealed by the deeper curvature of the Black-Scholes implied volatility smile.

[1]The author thanks an anonymous referee for this remark.

For parameters obtained through the GLS estimation, out-of-sample implied volatility plots are obtained by replicating the procedure described in Bakshi et al. (1997): For each model parameters are estimated at time $t - 1$ and used as inputs along with time-t observed variables (K, S, τ, R, R_f) in the computation of implied volatility levels at time t. Results are then grouped by moneyness/maturity categories and implied volatility values plotted against the moneyness, with results reported in Figure 3.

For parameters obtained through IS-GMM out-of-sample implied volatility plots are obtained in the following manner: Because IS-GMM does not allow parameters to change over time, IS-GMM parameters are estimated on the first half of the sample and used as inputs to the models

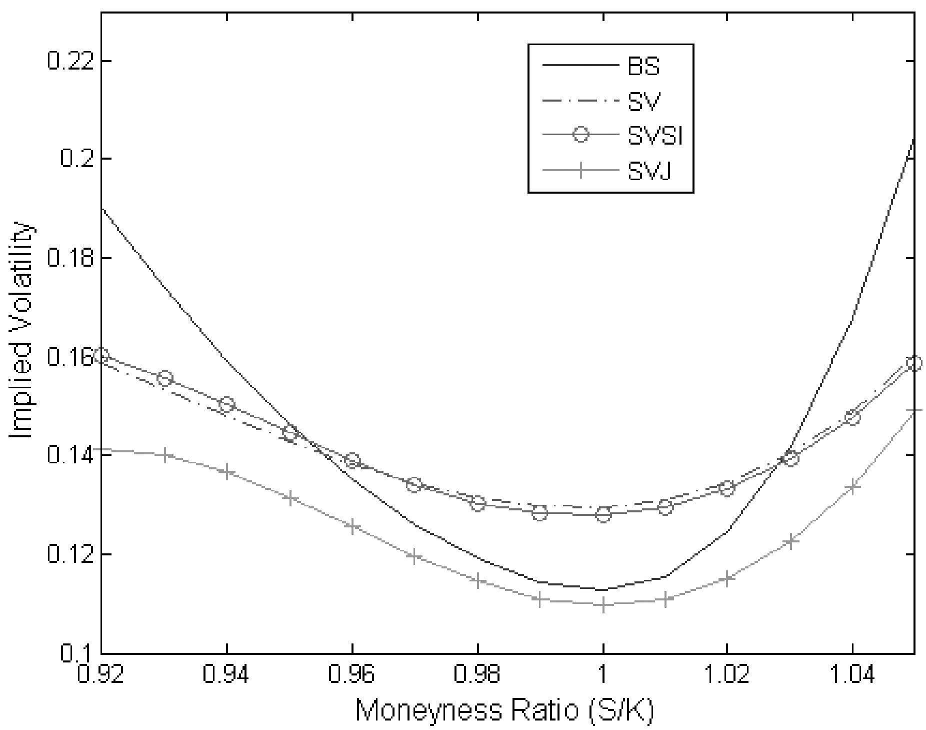

FIGURE 3

Volatility smiles for options of maturity between 0 and 60 days: GLS. The implied volatility is obtained by inverting the models separately for each call option contract on the Japanese yen for the period April 1996 to December 1999 when the parameters are estimated with GLS. The implied volatilities are then averaged within each moneyness-maturity class. Moneyness is represented by S/K, the ratio of the spot exchange rate of the Japanese yen currency over the exercise price. The models are Black-Scholes (BS), stochastic volatility (SV), stochastic-volatility stochastic interest rates (SVSI), and stochastic volatility with jumps (SVJ).

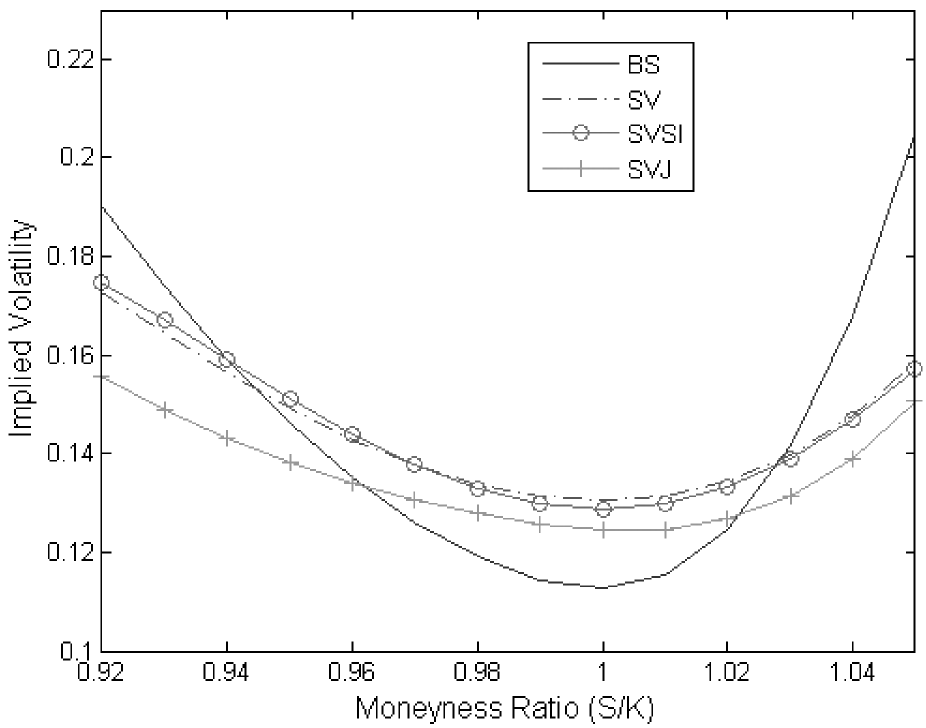

FIGURE 4

Volatility smiles for options of maturity between 0 and 60 days: IS-GMM. The implied volatility is obtained by inverting the models separately for each call option contract on the Japanese yen for the second half of the sample when the parameters are estimated with IS-GMM. The implied volatilities are then averaged within each moneyness-maturity class. Moneyness is represented by S/K, the ratio of the spot exchange rate of the Japanese yen currency over the exercise price. The models are Black-Scholes (BS), stochastic volatility (SV), stochastic-volatility stochastic interest rates (SVSI), and stochastic volatility with jumps (SVJ).

for every option in the second half of the sample. Within each moneyness/maturity class, implied volatility levels are computed and averaged, with results reported in Figure 4.

For ease of comparison, Figures 3 and 4 report implied volatility levels for all models (BS, SV, SVSI, and SVJ). The Black-Scholes implied volatility plot displays the deepest smile, and the other models eliminate some of the pricing biases and correspondingly produce flatter curves. Overall, model performance does not seem strongly dependent on the estimation technique used, as comparing the implied volatility plots produced by IS-GMM and GLS reveals very few differences. The only noticeable idiosyncrasy is the fact that the SVJ plot appears a little lower in the GLS case than in the IS-GMM case. It otherwise appears that when either GLS or IS-GMM is used as an estimation method, the SVJ model performs best, and the SV model provides the largest incremental improvement, thus flattening the curve. Finally, the SVSI model is almost indistinguishable from the SV model.

Out-of-Sample Percentage Pricing Errors

Another approach to studying model performance is to analyze pricing errors by moneyness and time to expiration. For parameters obtained through the GLS estimation, out-of-sample pricing errors are obtained by replicating the procedure described in Bakshi et al. (1997): For each model, parameters are estimated at time $t - 1$ and used as inputs along with time-t observed variables (K, S, τ, R, R_f) in the computation of pricing errors at time t. Pricing errors are defined as the difference between model-predicted prices and observed option prices. Results are then grouped by moneyness/maturity categories and reported in Table VI.

TABLE VI
Out-of-Sample Percentage Pricing Errors

		IS-GMM		*GLS*	
		Percentage pricing errors		*Percentage pricing errors*	
Moneyness S/K	*Model*	*<60 days*	*60–180 days*	*<60 days*	*60–180 days*
<0.94	BS	35.77	9.26	35.77	9.26
	SV	6.75	8.32	1.83	0.13
	SVSI	6.39	8.31	1.62	0.12
	SVJ	3.50	2.10	−0.41	0.08
0.94–0.97	BS	7.36	−0.60	7.36	−0.60
	SV	1.55	5.01	0.39	0.48
	SVSI	1.46	4.88	0.38	0.45
	SVJ	−1.11	1.59	4.00	0.33
0.97–1.00	BS	−1.32	−2.35	−1.32	−2.35
	SV	−0.28	1.66	0.05	0.06
	SVSI	−0.30	1.01	0.04	0.05
	SVJ	−1.63	0.69	0.22	−0.01
1.00–1.03	BS	−1.08	−1.47	−1.08	−1.47
	SV	−0.55	0.98	−0.24	0.28
	SVSI	−0.51	0.86	−0.16	0.11
	SVJ	−0.79	0.88	−0.37	0.19
>1.03	BS	−1.15	0.57	−0.15	0.57
	SV	−0.12	0.47	0.02	0.12
	SVSI	−0.10	0.28	−0.01	0.05
	SVJ	−0.10	0.13	−0.33	−0.48

Note. The table reports the pricing errors (model-predicted vs. market prices) as a percentage of the option price for each call option contract on the Japanese yen for the 1996–1999 period in the GLS case and for the second half of that period only in the IS-GMM case. The pricing errors are averaged within each moneyness-maturity class. The results are shown for each model: Black-Scholes (BS), stochastic volatility (SV), stochastic-volatility stochastic interest rates (SVSI), and stochastic volatility with jumps (SVJ). The performance of IS-GMM is reported on the left-hand side, and the performance of the GLS is reported on the right.

For parameters obtained through IS-GMM out-of-sample, pricing errors are obtained as follows: Because IS-GMM does not allow parameters to change over time, the IS-GMM parameters estimated in the first half of the sample are then used as inputs to the models for every option in the second half of the sample, in the same manner the implied volatility smile exercise is conducted. Within each moneyness/maturity class, pricing errors are computed and averaged with results reported in Table VI as well. Pricing errors are reported as percentages because option prices reflect one unit of the Japanese yen as the underlying asset and errors are thus very small in scale.

Table VI shows that the SV model provides the largest incremental improvement in pricing over Black-Scholes in most maturity/moneyness categories. The stochastic interest-rates feature only seems to improve pricing abilities further in the case of options with longer maturities that are at least in-the-money, as found by Bakshi et al. (1997). In other categories, incorporating stochastic interest rates does not provide any noticeable pricing improvement, regardless of the estimation technique employed. Introducing jumps to the model generally does lead to further pricing improvement with both estimation techniques, although it is possible for the jump feature to hurt the performance of the model in some categories with no particularly discernable pattern.

The main noticeable difference between IS-GMM and GLS is the fact that the out-of-sample fit is actually better overall when using the GLS methodology. This is most likely explained by the implementation of the exercise. Errors are computed in the GLS case by recalibrating the parameters on the day prior to the date when the pricing error is computed, whereas in the IS-GMM case the parameters are estimated once and kept unchanged when computing pricing errors. When comparing the fit of the models with respect to each other, the patterns described in the previous paragraph are very similar regardless of whether IS-GMM or GLS is employed.

DYNAMIC HEDGING PERFORMANCE

The tests conducted in the previous sections provide an indication of how the models perform in a static pricing environment. It may also be of interest to practitioners to learn how the models perform in a portfolio hedging exercise. The Black-Scholes, SV, SVSI, and SVJ models are thus put to the test and compared in a dynamic hedging setting. Comparisons between the GLS and IS-GMM estimation techniques are also of interest, because the outcome determines the way a practitioner should estimate the parameters.

In the Black-Scholes setting, only the underlying asset is needed to hedge a position in an option on that asset because the asset is the only source of risk. In more realistic settings such as the SVJ model, more sources of risk come into play. Random volatility, jumps and interest rates are additional sources of risk assumed away in the Black-Scholes setup. The delta-neutral hedge thus uses the underlying currency but also other assets in order to hedge the various sources of randomness. Note, however, that jump risk cannot be controlled for, as shown by Merton (1976), implying that frequent rebalancing is the main tool to hedge against the chance of a jump. Also note that in order for the strategy to make sense, it is assumed that transaction costs are not excessive. If the transactions costs are very large, then a more imperfect (but less expensive) single-instrument hedge is probably in order.

Assume that one has a short position in a call option $C(t, \tau)$ on the Japanese yen. The goal is to construct a hedging portfolio whose change in value replicates any change in value of the call option on the currency. In the most general setting (SVSI-J), this implies that one must take a position in $X_S(t)$ units of the Japanese yen currency in order to control for currency price risk, $X_B(t)$ units of a τ-period domestic discount bond to account for domestic interest-rate risk, $X_F(t)$ units of a τ-period Japanese discount bond to account for foreign interest-rate risk, and $X_C(t)$ units of another call option on the Japanese yen with a least one different characteristic. An option with a different exercise price or different maturity is thus all that is needed, even if the difference is not large; the hedge ratios will adjust accordingly. Finally, $X_0(t)$ is the residual cash position invested or borrowed at the risk-free rate.

By matching and thus eliminating the various types of risk between the option and the individual securities, it can be shown that the units invested in the various securities described above must have the following solutions, briefly displayed here and in the Appendix, as stochastic interest rates in a currency setting produce hedging ratios different enough from their equity counterparts to justify it:

$$X_C(t) = \frac{\Delta_V(t, \tau, K)}{\Delta_V(t, \tau', K')} \tag{15}$$

$$X_B(t) = \frac{1}{B(t, \tau)\psi(\tau)}\{\Delta_R(t, \tau', K')X_C(t) - \Delta_R(t, \tau, K)\} \tag{16}$$

$$X_F(t) = \frac{X_C(t)\Delta_{R_f}(t, \tau', K') - \Delta_{R_f}(t, \tau, K)}{S(t)\psi_F(\tau)F(t, \tau)} \tag{17}$$

$$X_S(t) = \Delta_S(t, \tau, K) - \Delta_S(t, \tau', K')X_C(t) - X_F(t)F(t, \tau) \tag{18}$$

$$X_0(t) = C(t, \tau, K) - X_S(t)S(t) - X_B(t)B(t, \tau) \\ - X_F(t)S(t)F(t, \tau) - X_C(t)C(t, \tau', K') \tag{19}$$

where

$$\Delta_j(t, \tau, K) = \frac{\partial C(t, \tau, K)}{\partial j} \quad \text{and}$$

$$\Delta_j(t, \tau', K') = \frac{\partial C(t, \tau', K')}{\partial j}, \quad \text{for } j = S, R, R_f, V$$

and where $B(t, \tau)$ is the dollar-denominated price of a domestic (U.S.) Treasury bill of maturity τ, and where $F(t, \tau)$ is the yen-denominated price of a foreign (Japan) Treasury bill of maturity τ. The details for $\Delta_S(t, T, K)$, $\Delta_V(t, T, K)$, and $\Delta_R(t, T, K)$ are given in the Appendix. The hedging effectiveness of the portfolio is measured by the hedging error term expressed as percentages, since straight hedging errors on one unit of the Japanese yen are too small to be visually meaningful in a table. At each rebalancing period an error term is computed, and the hedging portfolio is rebalanced. This time, the error term is given by

$$H(t + \Delta t) = X_S(t)S(t + \Delta t) + X_0(t)e^{R(t)\Delta t} - C(t + \Delta t, \tau - \Delta t) \\ + X_B(t)B(t + \Delta t, \tau - \Delta t) + X_F(t)S(t + \Delta t)F(t + \Delta t, \tau - \Delta t) \\ + X_C(t)C(t + \Delta t, \tau' - \Delta t, K') \tag{20}$$

Parameters associated with the IS-GMM methodology are obtained from the first half of the sample and held fixed for the hedging exercises. On any given date t from the second half of the sample, the parameter estimates are used along with the state variables ($S, \tau \ldots$) date-t values to construct the hedges corresponding to the various models for each option, and for each option, on the next closest available day—when the option reappears—hedging errors are recorded. The procedure is repeated for every option every day in the sample. The hedging exercise associated with GLS is conducted in the following manner. Information at time $t - 1$ is used to compute parameter estimates using GLS, and on the next day, on date t, these parameters are used to construct the specific hedges desired. Finally, on the next closest available day, hedging errors are recorded with the procedure repeated for every option every day in the sample. Note that Bakshi et al. (1997) are able to vary the rebalancing period (1 and 5 days) thanks to high liquidity levels present in index options, whereas the compromise here is to rebalance when next possible, yielding an average

TABLE VII
Delta-Neutral Hedging Average Pricing Errors

		IS-GMM		*GLS*	
		Percentage hedging errors		*Percentage hedging errors*	
Moneyness S/K	*Model*	*<60 days*	*60–180 days*	*<60 days*	*60–180 days*
<0.94	BS	−4.74	−5.98	−4.74	−5.98
	SV	−1.12	1.64	−0.38	1.57
	SVSI	−1.10	1.66	−0.40	1.58
	SVJ	0.25	1.66	−1.64	2.71
0.94–0.97	BS	−5.78	−8.14	−5.78	−8.14
	SV	−1.59	2.71	−0.87	2.89
	SVSI	−1.58	2.73	−0.87	2.92
	SVJ	−1.56	1.20	−0.96	3.00
0.97–1.00	BS	−7.60	−15.14	−7.60	−15.14
	SV	−0.51	1.31	−0.08	1.07
	SVSI	−0.49	1.24	−0.07	1.06
	SVJ	−1.03	−0.88	−0.06	2.85
1.00–1.03	BS	−8.43	−4.24	−8.43	−4.24
	SV	−2.47	0.39	−1.87	0.33
	SVSI	−2.45	0.40	−1.85	0.31
	SVJ	−2.53	−0.25	−1.44	0.42
>1.03	BS	−1.94	−4.02	−1.94	−4.02
	SV	−1.01	−0.78	−0.93	−0.60
	SVSI	−1.00	−0.67	−0.91	−0.55
	SVJ	−0.98	−1.29	−0.90	−0.68

Note. This table reports the percentage hedging errors associated with a hedge constructed with the use of a delta-neutral portfolio. For each call option contract on the Japanese yen, an average hedging error is computed over the life of the option. The hedging errors are averaged within each moneyness-maturity class. The results are shown for each model: Black-Scholes (BS), stochastic volatility (SV), stochastic-volatility stochastic interest rates (SVSI), and stochastic volatility with jumps (SVJ). The performance of IS-GMM is reported on the left-hand side, and the performance of the GLS is reported on the right.

rebalancing frequency of about 3 trading days. All hedging errors are finally grouped by moneyness and maturity in Table VII.

As in the pricing exercise, for ease of visual comparison percentage errors only are reported due to the errors' small scale. Note that straight hedging errors are reported instead of absolute hedging errors to reflect the point of view of a market maker whose gains and losses would be expected to somewhat offset each other over time, assuming the market maker maintains the position for more than one rebalancing period.

Table VII reveals that the largest improvement in hedging error is provided by the stochastic volatility feature. Incremental changes produced by the stochastic rates are negligible, whereas additional improvement as a result of adding a jump feature to the model is somewhat debatable. Some categories reveal a slight improvement, but many others do not. Most importantly, no clear pattern emerges from adding the jump feature. The introduction of a second option in order to hedge the stochastic volatility eliminates most of the gamma exposure, and further model modifications do not add much in terms of hedging ability. The GLS appears to yield slightly smaller hedging errors than IS-GMM overall, similarly to what was found in the pricing exercise. However, considering that the GLS parameters are recalibrated often and the IS-GMM parameters are not, the IS-GMM method can be deemed to perform as well as GLS. The information contained in the cross section of option prices thus does not seem highly critical for pricing and hedging purposes. There is otherwise no major difference to report regarding hedging error patterns among the different models, as the same conclusions can drawn regardless of whether the GLS or the IS-GMM table is read.

CONCLUSION

A currency option pricing model is implemented that features stochastic volatility, stochastic interest rates, and jumps, allowing for comparison of nested models in the spirit of the equity options study by Bakshi et al. (1997) and allowing for the modeling of skewness and kurtosis present in the risk-neutral pricing distribution of the Japanese yen. Two competing estimation methodologies are then conducted and tested: GLS (generalized least squares) and IS-GMM (implied-state generalized method of moments). Parameters are estimated under both techniques and used as inputs in various types of pricing and hedging comparisons with two goals in mind. One goal is to determine the importance of the information embedded in the cross section of option prices for hedging and pricing purposes, using the fact that GLS uses this cross-sectional information but IS-GMM does not. Another is to establish whether results obtained by Bakshi et al. (1997) in an index options market remain true in the case of a currency.

One finding is that most results obtained by Bakshi et al. (1997) in the equity case still hold well in the case of a currency, despite the currency's unique structural characteristics resulting from government interventions. The stochastic volatility model generally displays the largest incremental improvement over Black-Scholes; the jump feature provides lower but still meaningful pricing improvement, and the stochasticity of interest rates is useful mainly in cases of in-the-money longer-term options.

A second and major finding is that IS-GMM and GLS display relatively comparable levels of performance. Although GLS parameters actually perform marginally better in pricing and hedging tests, it is important to realize that the two estimation methods are fundamentally different and that conducting fair comparison tests is challenging. Because of the differences in their natures, GLS is often granted frequent readjusting of the parameters, but IS-GMM is not, a fact seriously undermining the apparent marginal superiority of GLS. In other words, because IS-GMM performs nearly as well as GLS, it appears that the cross-section of option prices does not contain information significantly valuable for pricing and hedging purposes.

This surprising conclusion thus seems to indicate that moment-based methods are somewhat superior to GLS because they are also internally more consistent than least-squares estimation techniques in the way the dynamics of the process are handled in the estimation. However, despite its shortcomings, GLS provides a feature that IS-GMM cannot claim: simplicity. GLS may require more frequent updating of the parameters, forecast can only be made one step ahead, and the information contained in the cross section of option prices may seem to be low, but the time and effort required to implement it are significantly less than what is needed in most moment-based estimations. Therefore, even though the information content of cross-sectional option prices does not seem to be of major importance, estimation techniques that make use of this information set will most likely remain a popular tool so long as they are straightforward to implement.

APPENDIX

Expressions for the Delta-Neutral Hedge (SVSI-J, Most General Case)

The number of units of the various assets needed for the hedge is

- $$\Delta_S(t,T,K) = \frac{\psi(1,X_t,0,T)}{2S} - \frac{1}{\pi}\int_0^{\infty}\frac{\mathrm{Im}[(1-iv)\psi(1-iv,X_t,0,T)e^{iv\ln K}]}{vS}\,dv - \frac{K}{\pi}\int_0^{\infty}\frac{\mathrm{Im}[i\psi(-iv,X_t,0,T)e^{iv\ln K}]}{S}\,dv$$
- $$\Delta_V(t,T,K) = \frac{1-e^{-\gamma_{V,1}\tau}}{2\gamma_{V,1}-(\gamma_{V,1}+\rho\sigma_V-\kappa_V)(1-e^{-\gamma_{V,1}\tau})}\frac{\psi(1,X_t,0,T)}{2} - \frac{1}{\pi}\int_0^{\infty}\mathrm{Im}\left[\frac{(1-e^{-\gamma_{V,2}\tau})(1-iv)^2\psi(1-iv,X_t,0,T)}{2\gamma_{V,2}-(\gamma_{V,2}+(1-iv)\rho\sigma_V-\kappa_V)(1-e^{-\gamma_{V,2}\tau})}\frac{e^{iv\ln K}}{v}\right]dv + \frac{K}{\pi}\int_0^{\infty}\mathrm{Im}\left[\frac{(1-e^{-\gamma_{V,3}\tau})(-iv)^2\psi(-iv,X_t,0,T)}{2\gamma_{V,3}-(\gamma_{V,3}-iv\rho\sigma_V-\kappa_V)(1-e^{-\gamma_{V,3}\tau})}\frac{e^{iv\ln K}}{v}\right]dv$$

where

$$\gamma_{V,1} = \sqrt{(\rho\sigma_V - \kappa_V)^2 - \sigma_V^2}$$
$$\gamma_{V,2} = \sqrt{((1 - iv)\rho\sigma_V - \kappa_V)^2 - (1 - iv)^2\sigma_V^2}$$
$$\gamma_{V,3} = \sqrt{(-iv\rho\sigma_V - \kappa_V)^2 - (-iv)^2\sigma_V^2}$$

- $$\Delta_R(t, T, K) = \frac{1}{\pi}\int_0^\infty \mathrm{Im}\left[\frac{2iv(1 - e^{-\gamma_{R,1}\tau})\psi(1 - iv, X_t, 0, T)}{2\gamma_{R,1} - (\gamma_{R,1} - \kappa_R)(1 - e^{-\gamma_{R,1}\tau})}\frac{e^{iv\ln K}}{v}\right]dv + K\frac{(1 - e^{-\gamma_{R,2}\tau})\psi(0, X_t, 0, T)}{2\gamma_{R,2} - (\gamma_{R,2} - \kappa_R)(1 - e^{-\gamma_{R,2}\tau})} - \frac{K}{\pi}\int_0^\infty \mathrm{Im}\left[\frac{2(1 + iv)(1 - e^{-\gamma_{R,3}\tau})\psi(-iv, X_t, 0, T)}{2\gamma_{R,3} - (\gamma_{R,3} - \kappa_R)(1 - e^{-\gamma_{R,3}\tau})}\frac{e^{iv\ln K}}{v}\right]dv$$

where

$$\gamma_{R,1} = \sqrt{\kappa_R^2 + 2iv\sigma_R^2}$$
$$\gamma_{R,2} = \sqrt{\kappa_R^2 + 2\sigma_R^2}$$
$$\gamma_{R,3} = \sqrt{\kappa_R^2 + 2(i + iv)\sigma_R^2}$$

- $$\Delta_{R_f}(t, T, K) = \frac{-(1 - e^{-\gamma_{R_f,1}\tau})\psi(1, X_t, 0, T)}{2\gamma_{R_f,1} - (\gamma_{R_f,1} - \kappa_{R_f})(1 - e^{-\gamma_{R_f,1}\tau})} + \frac{1}{\pi}\int_0^\infty \mathrm{Im}\left[\frac{2(1 - iv)(1 - e^{-\gamma_{R_f,2}\tau})\psi(1 - iv, X_t, 0, T)}{2\gamma_{R_f,2} - (\gamma_{R_f,2} - \kappa_{R_f})(1 - e^{-\gamma_{R_f,2}\tau})}\frac{e^{iv\ln K}}{v}\right]dv + \frac{K}{\pi}\int_0^\infty \mathrm{Im}\left[\frac{2iv(1 - e^{-\gamma_{R_f,3}\tau})\psi(-iv, X_t, 0, T)}{2\gamma_{R_f,3} - (\gamma_{R_f,3} - \kappa_{R_f})(1 - e^{-\gamma_{R_f,3}\tau})}\frac{e^{iv\ln K}}{v}\right]dv$$

where

$$\gamma_{R_f,1} = \sqrt{\kappa_{R_f}^2 + 2\sigma_{R_f}^2}, \quad \gamma_{R_f,2} = \sqrt{\kappa_{R_f}^2 + 2(1 - iv)\sigma_{R_f}^2},$$
$$\gamma_{R_f,3} = \sqrt{\kappa_{R_f}^2 - 2iv\sigma_{R_f}^2}$$

BIBLIOGRAPHY

Ait-Sahalia, Y. (1996). Nonparametric pricing of interest rate derivative securities. Econometrica, 64, 527–560.

Bakshi, G., & Cao, C. (2004). Disentangling the contribution of return-jumps and volatility-jumps on individual equity option prices (working paper). Pennsylvania State University and the University of Maryland.

Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical performance of alternative option pricing models. Journal of Finance, 52, 2003–2049.

Bakshi, G., Kapadia, N., & Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. The Review of Financial Studies, 16, 101–143.

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

Bates, D. (2000). Post-'87 crash fears in the S&P 500 futures option market. Journal of Econometrics, 94, 181–238.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. Journal of Political Economy, 81, 637–659.

Chapman, D., Long, J., & Pearson, N. (1999). Using proxies for the short rate: When are three months like an instant? Review of Financial Studies, 12, 763–806.

Cox, J., Ingersoll, J., & Ross, S. (1985). A theory of the term structure of interest rates. Econometrica, 53, 363–384.

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

Dumas, B., Fleming, J., & Whaley, R. (1998). Implied volatility smiles: An empirical investigation. Journal of Finance, 53, 2059–2106.

Eraker, B., Johannes, M., & Polson, N. (2003). The impact of jumps in equity index volatility and returns. Journal of Finance, 58, 1269–1300.

Hansen, L. (1982). Large sample properties of generalized method of moments estimators. Econometrica, 50, 1029–1054.

Merton, R. (1976). Option pricing when underlying stock returns are discontinuous. Journal of Financial Economics, 3, 125–144.

Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated time-series study. Journal of Financial Economics, 63, 3–50.