

The performance of alternative valuation models in the OTC currency options market

Nicolas P.B Bollen ^{a,*}, Emma Rasiel ^b

^a *Owen Graduate School of Management, Vanderbilt University, 401 21st Avenue South, Nashville, TN 37203, USA*

^b *Fuqua School of Business, Duke University, Box 90120, Durham, NC 27708, USA*

Abstract

We compare option valuation models based on regime-switching, GARCH, and jump-diffusion processes to a standard “smile” model, in which Black and Scholes (1973) implied volatilities are allowed to vary across strike prices. The regime-switching, GARCH, and jump-diffusion models provide significant improvement over a fixed smile model in fitting GBP and JPY option prices both in-sample and out-of-sample. The jump-diffusion model achieves the tightest fit. A time-varying smile model, however, provides hedging performance that is comparable to the other models for the GBP options. This result suggests that standard option valuation techniques may provide a reasonable basis for trading and hedging strategies.

© 2003 Elsevier Science Ltd. All rights reserved.

Keywords: Option valuation; Currency options; GARCH; Regime-switching; Jump-diffusion

1. Introduction

The size of currency derivative markets reflects the growing need to manage exchange rate risk in an integrated global economy. According to the Bank for International Settlements (1999), the notional value of all over-the-counter (OTC) currency derivatives outstanding at the end of June 1999 was \$14.9 trillion. The search for optimal valuation and hedging strategies in a market of this size deserves considerable effort. The standard approach is to use the Garman and Kohlhagen (1983) modification to the Black and Scholes (1973, hereafter Black–Scholes) model when valuing currency options. In this paper, we compare the performance of the Garman

* Corresponding author. Tel.: +1-615-343-5029; fax: +1-615-343-7177.

E-mail address: nick.bollen@owen.vanderbilt.edu (N.P.B. Bollen).

Nomenclature

μ	is the expected rate of appreciation in the foreign currency;
λ	is the frequency of jumps;
\bar{k}	$= E(e^\gamma - 1)$ is the mean percentage jump in the exchange rate;
Z	is a standard Weiner process;
γ	is the random jump in the log of the exchange rate conditional on a jump occurring and is distributed $N(\ln(1 + \bar{k}) - \delta^2/2, \delta^2)$, and;
q	is a Poisson counter with frequency λ : $\Pr(dq = 1) = \lambda dt$.

and Kohlhaugen (1983) model to option valuation models based on regime-switching, GARCH, and jump-diffusion processes, using a rich data set of OTC dealer quoted volatilities.

We focus on regime-switching, GARCH, and jump-diffusion processes because all have been used extensively in prior research to model the dynamics of foreign exchange rates. Hamilton (1990); Engel and Hamilton (1990), and Bollen et al. (2000) reject a random walk process in favor of a regime-switching model with constant within-regime mean and volatility. In addition, Bekaert and Hodrick (1993) attempt to explain the ‘forward premium puzzle’ bias in the currency markets using regime switches. West and Cho (1995); Jorion (1995), and Ederington and Guan (1999) all compare the abilities of various GARCH models to predict volatility. Regime-switching and GARCH models are popular because both can capture, at least in part, salient time series properties of foreign exchange rates: time-varying volatility and leptokurtosis or “fat tails” relative to the normal distribution. Neither model, however, permits discontinuities in the evolution of exchange rates. Jorion’s (1988) jump-diffusion model, in contrast, is able to capture shocks to an exchange rate as well as smooth evolution in “normal” times.

The first application of modern option valuation techniques to currency options is generally credited to Grabbe (1983) and Garman and Kohlhaugen (1983), who show how the Black–Scholes model can be modified to value currency options. Although the Black–Scholes assumption of geometric Brownian motion is violated empirically in the foreign exchange market, their model is still widely used due to its simplicity and tractability. Campa and Chang (1995, 1998), for example, argue that the Black–Scholes model generates accurate option values even though it may be misspecified, and use Black–Scholes implied volatilities to forecast exchange rate variance and the correlations between exchange rates.

Recent work focuses on developing efficient techniques for valuing options when underlying asset returns are governed by alternative stochastic processes. These valuation algorithms permit estimation of process parameters directly from option prices. Duan (1995) presents an equilibrium-based strategy for valuing options when a GARCH process governs underlying asset returns. Ritchken and Trevor (1999) describe a lattice algorithm for valuing options within the Duan (1995) model. Naik

(1993); Bollen (1998), and Duan et al. (1999) show how to value options when a regime-switching process governs underlying asset returns instead. Bates (1996a,b) develops a closed-form solution for the problem of valuing options when underlying asset returns are governed by a jump-diffusion model.

This paper contributes to the literature on currency option valuation by comparing the performance of a standard option valuation model to the performance of option valuation models based on regime-switching, GARCH, and jump-diffusion processes. To our knowledge, the analysis in Bates (1996a,1996b) of jump-diffusion option valuation and the Bollen et al. (2000) empirical study of regime-switching option valuation are the only published tests of these models in the foreign exchange market. Our study offers three insights. First, we can determine whether more complex option valuation models offer any advantages over standard models. The result will suggest whether future studies of currency options would benefit from moving beyond the Black–Scholes model. Second, a direct comparison of the models will provide guidance for practitioners regarding the optimal framework for developing trading and hedging strategies. Third, since the alternative time series models studied here are non-nested, standard econometric tests are invalid.¹ The options-based tests will therefore provide an additional diagnostic to determine which model is superior in capturing the time series properties of foreign exchange rates.

In this paper, we analyze a rich set of OTC option quotes on British Pounds Sterling (GBP) and Japanese Yen (JPY). We infer parameter values of the regime-switching, GARCH, and jump-diffusion models from option prices in a numerical procedure that is analogous to inferring a Black–Scholes implied volatility. We then use the implied model parameters to compute fitted values for the options in order to determine which model can match market prices most accurately. We compare the models to a standard “smile” model, in which Black–Scholes implied volatilities are allowed to vary across strike prices. We find that the regime-switching, GARCH, and jump-diffusion option valuation models perform better than a fixed smile model in fitting option prices in-sample and forecasting option prices out-of-sample. The jump-diffusion model achieves the tightest fit. When the smile is allowed to change over time, however, the smile model provides hedging performance that is comparable to the other models for the GBP options. This result suggests that in practice, standard option valuation techniques may provide a reasonable basis for trading and hedging strategies.

The rest of the paper is organized as follows: section 2 describes the data used in the study. Section 3 reviews regime-switching, GARCH, and jump-diffusion mod-

¹ Hamilton and Susmel (1994) present a regime switching model in which each regime is characterized by a different ARCH(q) process. In their model, the null of a single regime would be equivalent to a single ARCH(q) process governing innovations. In this sense their regime switching model nests a standard ARCH(q) process. However, as Hamilton and Susmel point out, the general econometric problem of regime switching models remains: the null of a single regime is not testable in the usual way since parameters of the other regime(s) are not identified. In their paper, Hamilton and Susmel rely on standard LR statistics as a simple measure of relative fit, even though they are strictly speaking invalid. Garcia (1998) and Hansen (1992) have developed solutions for these and related econometric problems.

els, both from a time series and option valuation perspective. Section 4 compares the ability of the alternative option valuation models to fit option prices. Section 5 concludes.

2. Data

2.1. Description

Currency options are traded both on organized exchanges and in OTC markets. The Philadelphia Stock Exchange lists currency options, while the Chicago Mercantile Exchange and the Singapore International Monetary Exchange list currency futures options. The vast majority of currency options, however, are traded in the OTC inter-dealer market. According to the Bank for International Settlements (1999), the notional value of outstanding OTC currency derivatives at the end of June 1999 was approximately 200 times the value of outstanding exchange-traded currency derivatives. Obtaining price data on OTC currency options has proved difficult in the past, since there is no requirement for dealers to report OTC prices for the public record. Fortunately, we were able to obtain a rich set of quotes from the OTC currency options desk of Goldman Sachs.

We have daily dealer quoted volatilities (DQVs) for European-style options on the GBP and the JPY for the period 11/13/97–1/29/99, a total of 304 trading days. According to the Bank for International Settlements (1999), options on these two currencies had the largest outstanding notional value at the end of June 1999, except for options on the Euro. Our data set contains midpoint DQVs for one-month, two-month and three-month options. Since the quotes are at the midpoint, we can ignore microstructure issues related to the bid-ask spread when inferring model parameters from option prices. At each maturity, the DQVs are further categorized by option delta, which is the first derivative of the Black–Scholes option value with respect to the exchange rate. There are 13 deltas reported for each maturity: six put options with deltas of -0.05 , -0.1 , -0.15 , -0.2 , -0.25 , and -0.35 , ranging from out-of-the-money to near-the-money, an at-the-money call option with a delta of 0.5 , and six call options with deltas of 0.35 , 0.25 , 0.2 , 0.15 , 0.1 , and 0.05 , ranging from near-the-money to out-of-the-money. A broad cross-section of option prices enables us to examine the ability of different valuation models to capture the empirical relation between strike price and implied volatility.

In order to convert a DQV to an option price, we must first infer the option's strike price from the DQV and the option's delta. The conversion requires observations of the spot exchange rate and appropriate domestic and foreign interest rates. The spot exchange rates are reported by Goldman Sachs in addition to the quoted volatilities. For the interest rates, we use daily observations of one and three month Eurodeposit rates obtained from *Datastream*. To proceed, let us first review the Black–Scholes model applied to foreign exchange rates. The model assumes that the underlying foreign exchange rate S follows a geometric Brownian motion with instantaneous volatility σ_S , and that the domestic interest rate r_d and the foreign interest rate r_f

are constant continuously compounded rates. Define the forward exchange rate over the option maturity T as:

$$F = Se^{(r_d - r_f)T}. \tag{1}$$

Under these assumptions, the value of a European-style call option with exercise price X and time to maturity T is given by Black's (1976) model of an option on a futures contract:

$$c = Fe^{-r_d T}N(d_1) - Xe^{-r_d T}N(d_2), \tag{2}$$

and the value of a European-style put option is given by

$$p = Xe^{-r_d T}N(-d_2) - Fe^{-r_d T}N(-d_1), \tag{3}$$

where

$$d_1 = \frac{\ln(F/X) + .5\sigma^2 T}{\sigma\sqrt{T}} \quad \text{and} \quad d_2 = d_1 - \sigma\sqrt{T}, \tag{4}$$

and $N(z)$ is the cumulative distribution function for a standard normal random variable with upper integral limit z .

The deltas of the options are defined in this context as the first derivative of option value with respect to the forward exchange rate. The deltas of European-style call and put options are

$$\delta_C = e^{-r_d T}N(d_1), \tag{5}$$

and

$$\delta_P = -e^{-r_d T}N(-d_1). \tag{6}$$

Now, given a particular delta, δ^* , we can solve for d_1 using the inverse of the cumulative normal distribution function. For the call,

$$d_1^* = N^{-1}(e^{r_d T}\delta_C^*), \tag{7}$$

and for the put,

$$d_1^* = -N^{-1}(-e^{r_d T}\delta_P^*). \tag{8}$$

Finally, the strike price of the option is computed by inverting Eq. (4):

$$X = Fe^{.5\sigma^2 T - d_1^* \sigma\sqrt{T}}, \tag{9}$$

in which σ is the DQV. Option prices are then computed from Eqs. (2) and (3) using the inferred strike price.

2.2. Summary statistics

Perhaps the most widely studied empirical feature of option prices is the relation between the strike price of an option and the corresponding implied volatility.

According to the Black–Scholes assumptions, there should be no relation. In currency option markets, however, there generally exists a symmetric relation, referred to as a smile, in which implied volatilities are higher for options that are either in-the-money or out-of-the-money than for at-the-money options. Fig. 1 displays the volatility smiles over our sample of options using the average of the DQVs at each delta.

Note that for both exchange rates, the smiles of the three maturities are quite similar. We therefore focus on the one-month options in our empirical work since much of the analysis is computationally demanding. For both exchange rates the smiles are relatively symmetric, in contrast to the implied volatility “skew” of equity index options that has prevailed since the crash of 1987. Rubinstein (1994) and Bates (2000) suggest that the implied volatility skew typically observed in stock options is caused by investors’ concern for large market down-turns, resulting in higher demand for out-of-the-money puts as insurance. The symmetric smiles in Fig. 1 can be interpreted as evidence that on average, there exists demand for hedges against both appreciation and depreciation in the foreign exchange market.

Although the average smiles of both the GBP and JPY are symmetric, there may

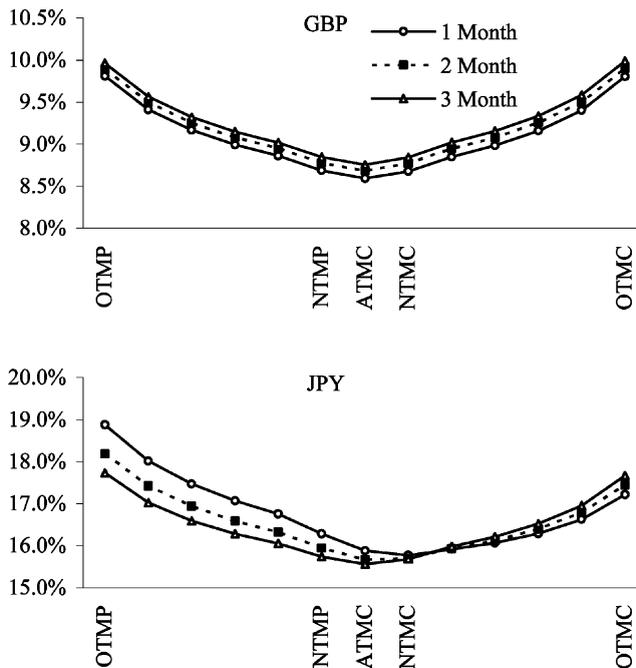


Fig. 1. Dealer quoted volatility smiles. Three volatility smiles corresponding to three different maturities are displayed for both currencies. The sample of options are quoted daily from 11/13/97–1/29/99, a total of 304 trading days. The smiles are constructed by averaging the dealer quoted volatilities at each delta over the sample. Five deltas are particularly important: OTMP and OTMC are out of the money puts and calls with Black–Scholes deltas of -5% and 5% respectively; NTMP and NTMC are near the money puts and calls with deltas of -35% and 35% respectively; and ATMC are at the money calls with deltas of 50% .

be a time variation in the shape of the smile, resulting from either changes in investors' beliefs about the future distribution of exchange rates or from a time variation in hedging demand. To explore this, we compute the slope of the smile every day on either side of the at-the-money option. Each day, we run two linear regressions for each currency of the form:

$$\sigma_i = \alpha + \beta M_i + \varepsilon_i, \quad (10)$$

where σ denotes the DQV and M denotes moneyness, defined as:

$$M_i = X_i/F - 1 \quad (11)$$

where X is the exercise price of the option and F is the forward exchange rate from Eq. (1). One regression (denoted OTMP) uses the six out-of-the-money put options and the at-the-money call option, the other regression (denoted OTMC) uses the at-the-money call option and the six out-of-the-money call options. The two slopes on each day provide a measure of the shape of the smile. Fig. 2 shows the evolution of the slopes on either side of the at-the-money-call over time. For the GBP, the OTMC slope is always positive and the OTMP slope is always negative, indicating that the smile is always convex. The slopes have roughly equal magnitude until September 1998, indicating a symmetric smile, then the smile appears to flatten first on the OTMP side, then the OTMC side, consistent with a distribution characterized by positive, then negative skewness. For the JPY, the OTMC slope is much flatter in the latter half of the sample than the OTMP slope, suggesting persistent negative skewness. Table 1 provides some summary statistics of the slopes. The OTMC slope is on average 0.1745 for the JPY and 0.2660 for the GBP, whereas the minimum OTMC slope is -0.0354 for the JPY and 0.1471 for the GBP, suggesting again that during particular time periods the JPY has substantial negative skewness.

One can view the study of alternative option valuation models as a search for a model that can generate the volatility smile while remaining internally consistent. In practice, one approach to valuing a cross-section of options is to take the previous trading day's smile as a given. To value an option at a particular strike price today, one simply uses the Black–Scholes model, inserting the prior day's volatility at the appropriate strike price and inserting today's underlying asset price, time to maturity, and prevailing interest rates. In fact the data used in this paper are recorded in a format that facilitates this procedure. Our options-based tests in section 4 will use this ad hoc procedure as a competitor in order to benchmark the performance of the regime-switching, GARCH, and jump-diffusion option valuation models.

3. Alternative option valuation models

This section reviews the structure of the regime-switching, GARCH, and jump-diffusion option valuation models, both from a time series and option valuation perspective.

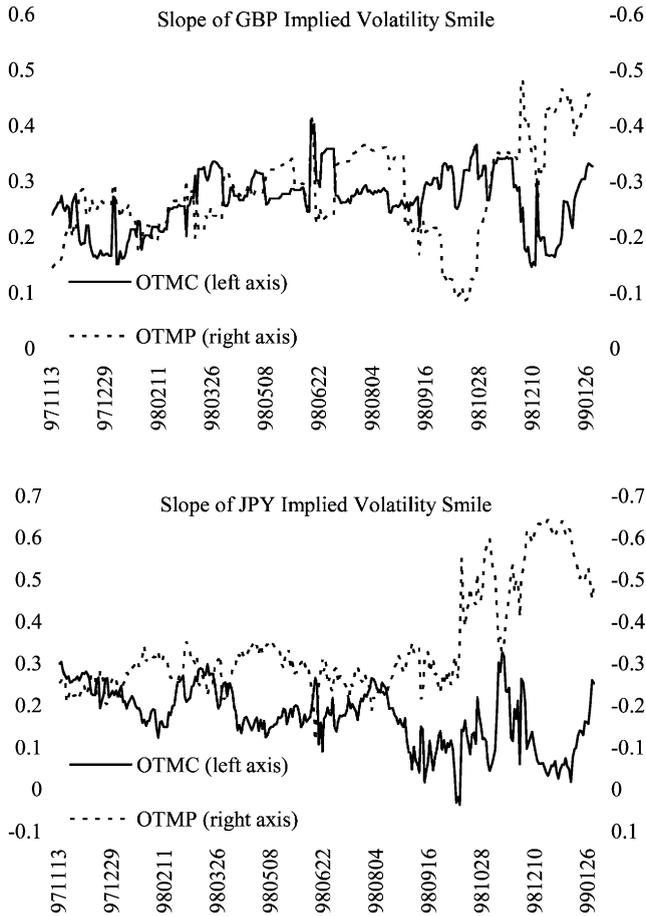


Fig. 2. Slopes of dealer quoted volatility smiles. Depicted are the daily slopes of the dealer quoted volatility (DQV) smiles for options on the GBP and JPY. The sample consists of options with maturity of one month quoted daily from 11/13/97–1/29/99, a total of 304 trading days. The slopes are estimated from the following OLS regression: $\sigma_i = \alpha + \beta M_i + \varepsilon_i$ where σ denotes the DQV and M denotes moneyness, defined as: $M_i = X_i/F - 1$, where X is the exercise price of the option and F is the forward exchange rate. Two slopes are computed each day. The OTMP slope uses seven options ranging from the ATM call to the deep OTM put. The OTMC slope uses seven options ranging from the ATM call to the deep OTM call.

3.1. Review of the regime-switching model

A regime-switching model is characterized by two or more data generating processes, each of which can govern the underlying variable at any point in time². The

² Hamilton (1988, 1989, 1990) pioneered the use of regime-switching models in financial economics.

Table 1
Summary statistics^a

	GBP			JPY	
	OTMP	OTMC		OTMP	OTMC
Mean	-0.2821	0.2660	Mean	-0.3424	0.1745
Standard deviation	0.0848	0.0537	Standard deviation	0.1219	0.0717
Maximum	-0.0866	0.4081	Maximum	-0.1165	0.3256
Minimum	-0.4841	0.1471	Minimum	-0.6445	-0.0354
#<0	304	0	#<0	304	3

^a Listed are summary statistics of the daily slopes of the dealer quoted volatility (DQV) smiles for options on the GBP and JPY. The sample consists of options with maturity of one month quoted daily from 11/13/97 to 1/29/99, a total of 304 trading days. The slopes are estimated from the following OLS regression: $\sigma_i = \alpha + \beta M_i + \varepsilon_i$, where σ denotes the DQV and M denotes moneyness, defined as: $M_i = X_i/F - 1$, where X is the exercise price of the option and F is the forward exchange rate. Two slopes are computed each day. The OTMP slope uses seven options ranging from the ATM call to the deep OTM put. The OTMC slope uses seven options ranging from the ATM call to the deep OTM call.

data generating processes can have identical structure but different parameter values, or, as in Gray (1996), they can have different structures as well. We use a two-regime model in which the data generating processes of both regimes are normal distributions. The means of the two distributions are restricted to be equal, but the volatilities may differ. The underlying variable is defined as daily log changes in foreign exchange rates, $r_t = \ln(S_t/S_{t-1})$, where S represents a foreign exchange rate in US dollars per GBP or US cents per JPY. The two-regime model studied here can be described in terms of the conditional distribution of log exchange rate changes as

$$r_t | \Phi_{t-1} \sim \begin{cases} N(\mu, \sigma_1) & \text{if } R_t = 1 \\ N(\mu, \sigma_2) & \text{if } R_t = 2 \end{cases} \tag{12}$$

where R_t denotes the regime operating at time t and Φ_{t-1} represents information available at time t . In our analysis we restrict the information set to be past observations of the exchange rate. The variable R_t evolves according to a first-order Markov process with transition probability matrix

$$\Pi = \begin{bmatrix} \pi_1 & 1 - \pi_1 \\ 1 - \pi_2 & \pi_2 \end{bmatrix}, \tag{13}$$

where $\pi_1 = \Pr(R_{t+1} = 1 | R_t = 1)$ is the probability of staying in regime 1 from one period to the next, and $\pi_2 = \Pr(R_{t+1} = 2 | R_t = 2)$ is the probability of staying in regime 2. These probabilities are referred to as the persistence of each regime.

While regime-switching models were originally written in terms of the persistence parameters, Hamilton (1994) and Gray (1996) show that estimation can be simplified by rewriting the model in terms of the probability that each regime is governing the variable at a particular date. These regime probabilities can be written as $\Pr(R_t =$

$i|\Phi_{t-1}$), and can be constructed recursively in a Bayesian framework as described in Gray (1996).

In most applications of regime-switching models, the mean of the distribution is allowed to differ across regimes, sometimes switching with the volatility, as in Engel and Hamilton (1990), or switching independently, as in Bollen et al. (2000). Different regime means are naturally interpreted in the foreign exchange market as periods of appreciation and depreciation. When valuing currency options in a risk-neutral framework, however, the mean of both conditional distributions is determined by relevant risk-free interest rates, as described next.

Regime-switching models are usually estimated by maximum likelihood, which is convenient and straightforward, given the relatively simple form of the distribution of the data³.

3.2. Option valuation in regime-switching models

We use the pentanomial lattice developed by Bollen (1998) to value our set of European-style options. Five branches emanate from each node; two branches correspond to one regime, the remaining three branches represent the other. This structure produces efficient recombining in the lattice. The branch probabilities and step sizes are selected to match the mean and variance of each regime. See Appendix A for additional detail and pseudo-code for implementing the valuation procedure.

The two-regime model used here is a stochastic volatility model, since volatility changes with regime. This additional source of uncertainty precludes the standard arbitrage argument supporting risk-neutral valuation. Naik (1993) demonstrates how to account for a volatility risk premium in a regime-switching model via an adjustment in the persistence parameters of the volatility regimes. Using our notation, the risk-adjusted transition probability matrix Π^* takes the following form:

$$\Pi^* = \begin{bmatrix} \pi_1(1 + \eta_1) & 1 - \pi_1(1 + \eta_1) \\ 1 - \pi_2(1 + \eta_2) & \pi_2(1 + \eta_2) \end{bmatrix} \quad (14)$$

where the η terms denote the risk coefficient in each regime. In this paper, we focus on computing implied model parameters from option prices. Rather than incorporate an additional parameter to reflect a volatility risk premium, we interpret the parameters implied from option prices as being net of the risk adjustment. In other words, we estimate the matrix Π^* as

$$\Pi^* = \begin{bmatrix} \pi_1^* & 1 - \pi_1^* \\ 1 - \pi_2^* & \pi_2^* \end{bmatrix}. \quad (15)$$

The transition parameters we infer from option prices may or may not equal the “real-world” parameters depending on whether there is a volatility risk premium. In either case, as shown in Naik (1993), the mean of both conditional distributions is

³ See Gray (1996) for details.

the risk-neutral mean, which for currencies equals the domestic risk-free interest rate less the foreign risk-free interest rate.

3.3. Review of the GARCH model

The regime-switching model described above provides one explanation for volatility clustering present in the data: different regimes correspond to different levels of volatility. An alternative explanation is provided by the GARCH class of models of Engle (1982) and Bollerslev (1986). In this paper, we use a GARCH(1,1) model (hereafter GARCH), which specifies a variance that is a function of the squared prior innovation and the prior observation of variance:

$$\sigma_t^2 = \alpha + \beta \varepsilon_{t-1}^2 + \gamma \sigma_{t-1}^2 \tag{16}$$

where the innovation is defined as $\varepsilon_t = r_t - \mu$, and μ is the constant mean of the random variable \tilde{r} .

Under the assumption of normality, the parameters of the GARCH model can be estimated in a maximum likelihood framework. To ease parameter estimation in the GARCH option valuation model, let us define the long-run or unconditional variance of the GARCH process as

$$\phi = \frac{\alpha}{(1-\beta-\gamma)}. \tag{17}$$

In the estimation of the GARCH parameters, this new parameter is estimated instead of the parameter α in Eq. (16), above.

3.4. Option valuation in GARCH models

The method used to value options consistent with a GARCH process is based on the lattice described in Ritchken and Trevor (1999), who invoke the theoretical work of Duan (1995) to establish a local risk-neutral measure under which options can be valued as expected payoffs discounted at the risk-free rate. A risk adjustment is made to one of the GARCH volatility parameters to account for a volatility risk premium. In particular, Eq. (16) takes the form

$$\sigma_t^2 = \alpha + \beta(\varepsilon_{t-1} - \eta)^2 + \gamma \sigma_{t-1}^2. \tag{18}$$

As with the regime-switching model, we do not explicitly incorporate a risk-adjustment parameter, but rather interpret the parameter values implied from option prices as being net of any risk premium that may be built into the option price. In other words, we estimate the following approximation to Eq. (18):

$$\sigma_t^2 = \alpha + \beta^* \varepsilon_{t-1}^2 + \gamma \sigma_{t-1}^2. \tag{19}$$

As with the regime-switching model, the parameters we infer from option prices may or may not equal the “real-world” parameters depending on whether there is a premium for the volatility risk. See Appendix B for additional detail and pseudo-code for implementing the valuation procedure.

3.5. Review of jump-diffusion models

Neither the GARCH model nor the regime-switching model can accommodate discontinuities in the evolution of exchange rates. In this paper, we use Bates' (1996b) application of a jump-diffusion model to exchange rates as an alternative specification which does. Following his notation, we assume that the log of the exchange rate is governed by the following process:

$$\frac{dS}{S} = (\mu - \lambda \bar{k})dt + \sigma dZ + (e^{\gamma} - 1)dq. \quad (20)$$

The volatility of log exchange rate changes is assumed to be constant over the life of the option. In practice, it is reasonable to assume that traders allow volatility to change over time. In contrast to the regime-switching and GARCH models, however, the jump-diffusion model does not provide structure to time variation in volatility. In the estimation of the jump-diffusion parameters, we fix the jump-related parameters over the entire data set, but allow the volatility parameter to change on each date.

3.6. Option valuation in jump-diffusion models

Bates (1991) develops a closed-form valuation formula for options when underlying asset returns are governed by a jump-diffusion, and applies it to currency options (1996b). The volatility of exchange rate returns is assumed constant, but the undiversifiable risk of jumps requires special care in the valuation. As with the regime-switching and GARCH option valuation models, a risk-adjustment is made to some of the parameters of the stochastic process to reflect the risk premium associated with the undiversifiable risk. The risk-neutral process is expressed as:

$$\frac{dS}{S} = (r_d - r_f - \lambda^* \bar{k}^*)dt + \sigma dZ^* + \bar{k}^* dq \quad (21)$$

where: r_d is the domestic interest rate and r_f is the foreign interest rate, and the other parameters are defined as before except now they reflect the risk-adjustment. The option valuation formula in Bates (1996b) for European options is similar to the Black–Scholes formula, but involves a summation of conditional option values, where the conditioning information is the number of jumps that occurred over the life of the option.

4. Option valuation performance of alternative foreign exchange rate models

This section determines the extent to which OTC option prices are consistent with the competing models. There are two parts to the analysis. First, parameters for all three models are implied from option prices in the same manner that Black–Scholes implied volatilities can be recovered from option prices. The calibrated models are

compared by their ability to fit quoted option prices both in-sample and out-of-sample. Second, a hedging experiment tests the true out-of-sample performance of the option valuation models. The first half of the option data set is withheld and used to estimate parameters of the three models implicit in option prices. The calibrated models are then used to hedge positions in options in the second half of the data set. The hedging performance of the alternative models is compared to the performance of an ad hoc strategy.

4.1. Model parameters implicit in option prices

In this subsection, parameters of the regime-switching, GARCH, and jump-diffusion models are inferred from option prices. This empirical task is closely related to the process of inferring the risk-neutral probability density function (pdf) of the underlying exchange rates. Campa et al. (1998) use OTC currency option prices to infer exchange rate pdfs three ways, one of which, the mixture of lognormals, is similar to the regime-switching model studied here. However, the mixture of lognormals model specifies constant weights on the two regimes, whereas the regime-switching model allows for time-varying weights (regime probabilities) that are a function of past observations. In addition, Campa et al. (1998) infer a different pdf on each date using current option prices, whereas we infer a single set of parameter values in order to price all options in our sample.

Once every two weeks in our sample (for a total of 31 dates), five one-month options are included in a subset of options for this experiment: the three closest to the money and the two furthest from the money. This particular cross-section was selected in order to capture the smile with a relatively small subset of options. Model parameters are estimated using nonlinear least squares as follows. Let θ denote a vector of model parameters, let $O_{i,t}$ denote the i th observed option price on date t , and let $V_{i,t}(\theta)$ denote the corresponding model value as a function of the parameter vector θ . Let T denote the number of dates on which option prices are observed in the sample. The nonlinear least squares estimator of the parameter vector θ minimizes S , the sum of squared errors (SSE) between the observed option prices and the model values:

$$S = \sum_{t=1}^T \sum_{i=1}^5 [O_{i,t} - V_{i,t}(\theta)]^2. \quad (22)$$

The computer program that minimizes S in Eq. (22) uses the IMSL routine DBCPOL, which performs a direct search over θ using a geometric complex.

The regime-switching parameters inferred from option prices include the variance and persistence of the two volatility regimes. Consistent with the regime-switching model, the option inference procedure restricts regime volatilities and persistence parameters to be constant for the entire sample of options. For each parameter vector considered in the estimation, the regime probabilities are constructed using Bayesian updating and a time series of exchange rates. Panel A of Table 2 lists the implied regime-switching parameters. For both currencies, the volatility regimes are quite

Table 2
Model parameters implicit in option prices

Panel A. Regime-switching estimates		
Parameter	GBP	JPY
σ_1	0.003842	0.006651
σ_2	0.005796	0.010513
π_1	0.993422	0.998905
π_2	0.991829	0.990056
Panel B. GARCH estimates		
Parameter	GBP	JPY
ϕ	0.000039	0.000044
β	0.041584	0.007378
γ	0.953223	0.985626
Panel C. Jump-diffusion estimates		
Parameter	GBP	JPY
λ	1.713639	17.609943
k	-0.001329	-0.002156
δ	0.034979	0.034140

Risk-neutral parameters of a two-regime model, a GARCH model, and a jump-diffusion model are estimated by minimizing the sum of squared deviations between model values and OTC option prices. Options are valued once every two weeks over the period 11/13/97–2/4/99. A total of 155 options on 31 dates are included. Regime-switching option values are computed using a lattice described in Bollen (1998). GARCH option values are computed using a lattice described in Ritchken and Trevor (1999). Jump-diffusion option values are computed using the closed-form solution in Bates (1996b).

distinct, with one regime's volatility approximately 50% higher than that of the other regime. Note also that the implied persistence parameters are all close to one, suggesting that the option prices do not imply a significant likelihood of switching regimes.

For the GARCH model, the three variance parameters are estimated, and are held fixed over the option sample. For each parameter vector considered in the estimation, the GARCH volatilities are constructed using the GARCH recursion and a time series of exchange rates. Panel B of Table 2 lists the results. The sum of the coefficients on squared lagged innovations and lagged variance, β and γ , respectively, are close to one, indicating a nearly integrated GARCH process with highly persistent volatility shocks. This is analogous to the persistent volatility regimes of the regime-switching model.

Panel C of Table 2 lists the implied jump-diffusion parameters that are held fixed over the sample. The λ parameter, which measures the annual frequency of jumps, differs substantially across the exchange rates, equaling 1.7 for the GBP and 17.6 for the JPY. The other parameters are similar, however, with the average jump being close to zero for both exchange rates, indicating that both positive and negative jumps are likely to occur, with a volatility of about 3.5%. The frequency of jumps for the JPY seems higher than one might expect, in the sense that jumps are usually interpreted as rare discontinuities in a time series. However, we find that several

parameter combinations achieve almost identical measures of in-sample fit, including some with lower jump frequencies. As noted by Bates (1996b), what is important is the distribution implied by a particular parameter combination, and presumably different parameter combinations can generate similar distributions.

In order to compare the three models, the rest of this section illustrates their respective abilities to fit option prices, both in and out-of-sample, as well as their abilities to hedge positions in options.

4.2. *Fitting observed option prices*

For both exchange rates, several measures of fit are computed at each option delta by comparing market prices to model values using the implied parameters. Panels A and B of Table 3 report the average percentage pricing error and the root mean squared error for the three models estimated in subsection 4.1., and a smile model, across the options on the 31 dates used in the estimation of model parameters. The smile model estimates an implied volatility smile for each exchange rate by choosing a volatility at each delta that minimizes the SSE between the market prices and the modified Black–Scholes price. The smile is held fixed over the option sample, so it cannot reflect time-variation in volatility in the same manner as the other models. In practice, one could use the prior date's smile to value options, but fixing the smile allows us to interpret any difference between the smile model and the other models as resulting from either their ability to capture time variation in volatility or the inclusion of jumps⁴.

As seen in Table 3, the regime-switching, GARCH, and jump-diffusion models offer a marked improvement over the smile model. The RMSE of the GBP ATM calls, for example, is 0.0015 for the regime-switching model, 0.0014 for the GARCH model, 0.0002 for the jump-diffusion model, and 0.0024 for the smile model. Similarly, the RMSE of the JPY ATM calls is 0.0014 for the regime-switching and GARCH models, 0.0006 for the jump-diffusion model, and 0.0028 for the smile model. The GARCH model and regime-switching model appear to fit option prices about equally well, but the jump-diffusion model clearly dominates. This is especially clear in the OTMP options, where the average percentage error is between three and six times larger for the GARCH and regime-switching models. Panel C summarizes the results by reporting the SSE for the three models. The GARCH model SSE is about 11% lower than the regime-switching model for the GBP, whereas the two models perform about the same for the JPY. Both models dramatically reduce the SSE for both exchange rates relative to the smile model. Again, the jump-diffusion model dominates: its SSE is about 1/20th the size of the regime-switching and GARCH SSE for the GBP options and 1/3rd the size for the JPY.

The dominance of the jump-diffusion model is perhaps not surprising, since the

⁴ As pointed out by the referee, the performance of the smile model would improve if we allowed the shape of the smile to change over time. This is one approach used in practice to value and hedge options. In the hedging experiment of section 4, the time-varying smile is used as an alternative to the GARCH and regime-switching models.

Table 3
Analysis of in-sample valuation errors

Panel A. GBP		Panel B. YEN		Panel C. SSE	
	%Error		%Error		%Error
	RS		RS		RS
OTMP	-30.43%	OTMP	-44.42%	OTMP	0.000183
NTMP	2.30%	NTMP	-1.44%	NTMP	0.000185
ATMC	2.05%	ATMC	1.69%	ATMC	0.000185
NTMC	2.05%	NTMC	3.40%	NTMC	0.000185
OTMC	-30.38%	OTMC	-21.02%	OTMC	0.000185
	GARCH		GARCH		GARCH
	18.13%		-47.87%		0.000162
	0.71%		-1.51%		0.000186
	0.53%		1.47%		0.000186
	0.60%		3.16%		0.000186
	17.84%		-24.52%		0.000186
	JUMP		JUMP		JUMP
	-4.93%		11.14%		0.000008
	0.08%		0.56%		0.000058
	0.48%		-1.03%		0.000058
	-0.14%		1.59%		0.000058
	-10.11%		23.90%		0.000058
	Smile		Smile		Smile
	-11.38%		-24.86%		0.000484
	2.90%		7.19%		0.000484
	2.25%		4.71%		0.000484
	2.99%		5.45%		0.000484
	-13.28%		-17.74%		0.000484
	RMSE		RMSE		RMSE
	RS		RS		RS
	0.000430		0.000599		0.000516
	0.000354		0.001462		0.000909
	0.001268		0.001395		0.000641
	0.001377		0.001193		0.000540
	0.000237		0.000340		0.000296
	0.001226		0.000332		0.000296
	0.000317		0.000332		0.000296

Panels A and B list the average valuation errors of five sets of currency options for four models: RS denotes a regime-switching model, GARCH denotes a GARCH model, JUMP denotes a jump-diffusion model, and Smile denotes the Black-Scholes model with a fixed volatility across the sample that differs across the five sets of options. The five sets of options differ by delta: OTMP and OTMC are out of the money puts and calls with Black-Scholes deltas of -5 and 5% respectively; NTMP and NTMC are near the money puts and calls with deltas of -35 and 35% respectively; and ATMC are at the money calls with deltas of 50%. The options are valued on 31 dates between 11/13/97-1/31/99. Listed are the average percentage error and the root mean squared error. Panel C lists the sum of squared valuation errors for the 155 options in the sub-sample across the four models.

volatility on each date is allowed to vary as a free parameter, whereas the GARCH and regime-switching models posit a specific dynamic for volatility. The out-of-sample tests described next set the models on more of an even footing.

Table 4 lists measures of fit between model values and observed option prices computed across the 273 dates in the sample that were not used in the estimation of model parameters. So this is, in a sense, an out-of-sample measure of fit. The regime-switching, GARCH, and jump-diffusion models again provide significant improvement over the fixed smile model. Here, though, the jump-diffusion model is not as dramatically superior to the other two as it was in the in-sample test. As seen in Panel C, the GARCH model generates an SSE for the GBP that is about 17% lower than the regime-switching model, and 8% lower for the JPY. The jump-diffusion model lowers the SSE for the GBP by an additional 15% below the GARCH SSE, and displays an almost-identical level of fit for the JPY.

Closer inspection of the percentage pricing errors in Table 4 shows how the different models fit option prices at various deltas. For the GBP options, the regime-switching model slightly overvalues the three close-to-the-money options but significantly undervalues the out-of-the-money options. The GARCH model fits the close-to-the-money options quite well, but overvalues the out-of-the-money options. The jump-diffusion model has errors in between the other two models for the close-to-the-money options, but fits the out-of-the-money options much more tightly. For the JPY options, the regime-switching and GARCH models generate a reasonable fit for the close-to-the-money options but dramatically undervalue the out-of-the-money options. The jump-diffusion model slightly overvalues the OTMP options and significantly overvalues the OTMC options. Fig. 3 graphically displays this information by plotting the implied volatility smiles of the regime-switching, GARCH, and jump-diffusion models versus the actual dealer quoted volatility smiles of the option prices.

All options in the data set were valued using the implied parameters of the three models. Then, Black–Scholes implied volatilities were inferred from these fitted model prices. The implied volatilities were then averaged over each delta. The dealer quoted volatility smiles are identical to those depicted in Fig. 1, except that only five deltas are included. For both exchange rates, the regime-switching model generates an implied volatility smile that is too flat, so that the out-of-the-money options are severely undervalued. The GARCH model also generates a flat smile for the JPY, but does much better than the regime-switching model for the GBP in capturing the observed smile. For both exchange rates, the jump-diffusion model is able to capture the implied volatility smile quite well.

The failure of the regime-switching and GARCH models to fit OTM option prices is likely caused by an inability to capture the higher-order moments of the implied return distribution of exchange rate changes⁵. This provides support for the jump-diffusion model, as well as the semi-nonparametric approach of Gallant et al. (1991).

⁵ The authors thank the referee for this observation.

Table 4
Analysis of out-of-sample valuation errors

Panel A. GBP									
	%Error								
	RS	GARCH	JUMP	Smile	RMSE	GARCH	JUMP	JUMP	Smile
OTMP	-30.31%	16.90%	-1.12%	-10.17%	0.000448	0.000352	0.000301	0.000551	
NTMP	2.37%	0.45%	1.20%	3.36%	0.001509	0.001372	0.001292	0.002217	
ATMC	2.15%	0.40%	1.22%	2.58%	0.001671	0.001545	0.001412	0.002447	
NTMC	2.05%	0.26%	0.91%	3.37%	0.001507	0.001398	0.001279	0.002289	
OTMC	-29.61%	16.73%	-6.89%	-12.59%	0.000457	0.000342	0.000298	0.000573	
Panel B. YEN									
	%Error								
	RS	GARCH	JUMP	Smile	RMSE	GARCH	JUMP	JUMP	Smile
OTMP	-48.21%	-51.09%	6.79%	-28.65%	0.000640	0.000637	0.000560	0.000752	
NTMP	-4.22%	-3.77%	-1.69%	4.53%	0.002021	0.001881	0.001910	0.003159	
ATMC	-0.28%	-0.15%	-2.71%	2.84%	0.002023	0.001938	0.002010	0.003235	
NTMC	0.33%	0.72%	-0.78%	2.62%	0.001709	0.001685	0.001680	0.002722	
OTMC	-26.41%	-28.14%	19.15%	-23.45%	0.000381	0.000394	0.000387	0.000557	
Panel C. SSE									
	%Error								
	RS	GARCH	JUMP	Smile	RMSE	GARCH	JUMP	JUMP	Smile
GBP	0.002115	0.001765	0.001496	0.005062					
YEN	0.003180	0.002919	0.002996	0.008520					

Panel A and B list the average valuation errors of five sets of currency options for four models: RS denotes a regime-switching model, GARCH denotes a GARCH model, JUMP denotes a jump-diffusion model, and Smile denotes the Black-Scholes model with a fixed volatility across the sample that differs across the five sets of options. The five sets of options differ by delta: OTMP and OTMC are out of the money puts and calls with Black-Scholes deltas of -5 and 5% respectively; NTMP and NTMC are near the money puts and calls with deltas of -35 and 35% respectively; and ATMC are at the money calls with deltas of 50%. The options are valued on the 273 dates between 11/13/97-1/31/99 that are withheld from the full sample when implying model parameters. Listed are the average percentage error and the root mean squared error. Panel C lists the sum of squared valuation errors for the 1365 options in the sub-sample across the four models.

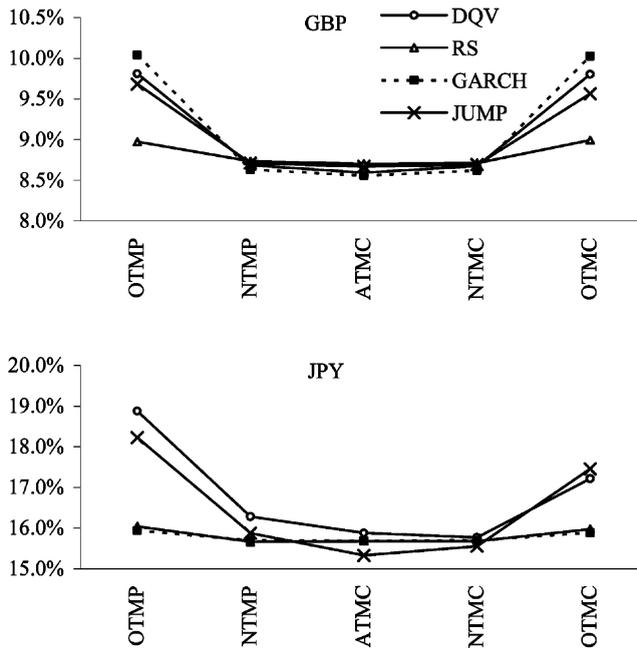


Fig. 3. **Implied volatility smiles.** Four volatility smiles are displayed for both currencies: DQV is the average dealer quoted volatility at each delta, RS is the average Black–Scholes volatility at each delta based on fitted regime-switching option prices using implied regime-switching parameters, GARCH is the average implied Black–Scholes volatility at each delta based on fitted GARCH option prices using implied GARCH parameters, and JUMP is the average Black–Scholes volatility at each delta based on jump-diffusion option prices using implied jump-diffusion parameters. A cross-section of five one month options is selected to imply model parameters and to compute implied volatility smiles. OTMP and OTMC are out of the money puts and calls with Black-Scholes deltas of -5 and 5% respectively. NTMP and NTMC are near the money puts and calls with deltas of -35 and 35% respectively. ATMC are at the money calls with deltas of 50% . The full sample of options are quoted daily from 11/13/97 to 1/29/99, a total of 304 trading days. The subset used to imply model parameters consists of the five options in the cross-section selected once every two weeks over the sample (31 dates). The subset used to compute the implied volatility smiles consists of the five options in the cross-section each day in the sample.

4.3. Hedging experiment

A true out-of-sample test of model performance requires that information used to value options or imply model parameters at time t be available at time t . In the following experiment, we assume the role of a market maker who sells puts and calls on foreign currency and delta-hedges the resulting short positions in options using one of four models. The models are the regime-switching model, the GARCH model, the jump-diffusion model, and a smile model that is updated daily. Parameters of the regime-switching, GARCH, and jump-diffusion models are implied from the first half of the option data set following the same procedure outlined above. The remainder of the option data set is used in the hedging experiment to judge how

effectively the different models can measure the sensitivity of option price to changes in the underlying exchange rate.

The options used to imply model parameters are the same cross-section of five options used previously, and are again selected every two weeks. However, option prices on only the first 16 of the original 31 dates are included in this subset. The options used to test hedging performance are selected weekly for the next 25 weeks of the data set, and are the same cross-section of five one month options. This is a true out-of-sample test since the model parameters used to hedge options are available at the time the options are hedged. A total of 125 options are sold for each currency⁶. Positions are assumed to be held to maturity and are hedged daily.

Each option is assumed to represent the right to buy or sell \$1,000,000 worth of the underlying currency. The number of options at a particular delta sold on a particular date therefore equals 1,000,000 divided by the spot exchange rate (in \$ per unit of foreign currency) on that date. Option proceeds are assumed to be invested at the prevailing Eurodollar rate for the life of the option. Proceeds are based on the midpoint prices used in the study, so they underestimate the revenue that would be generated in practice. Since we are concerned primarily with the relative performance of different strategies, this will not qualitatively affect the results. On the first day the option is held, the hedge ratio is determined by one of four methods. For the regime-switching, GARCH, and jump-diffusion models, the hedge ratio equals the derivative of the fitted option value, using implied model parameters, with respect to the underlying exchange rate. Derivatives are estimated numerically for these models. For the smile model, the hedge ratio is given by the analytic formulae in Eqs. (5) and (6) using the dealer quoted volatility. The short call positions require long positions in the foreign currency. The size of the long position equals the call's hedge ratio times the number of options sold. We assume that US dollars are borrowed at the prevailing one month Eurodollar rate to purchase the foreign currency, which is in turn deposited in an account bearing interest at the appropriate one month Euro rate. Similarly the short put positions require short positions in the foreign currency. We assume that the foreign currency is borrowed at the prevailing one month Euro rate, converted into U.S. dollars at the prevailing spot exchange rate, and deposited in an account bearing interest at the prevailing one month Eurodollar rate.

Each trading day of the option's life, up to but not including the expiration of the option, the hedge ratio is recomputed. For the regime-switching and GARCH models, regime probabilities and GARCH volatilities are first updated using their respective recursive schemes. For the jump-diffusion model, the most recently implied instantaneous volatility is used to compute the hedge ratio. For the smile model, the volatility used to compute the hedge ratio is given by the prevailing dealer quoted volatility at that delta. This strategy is used in order to simulate how a trader might use available information to set a hedge ratio. The number of units of foreign currency

⁶ A larger set of options could be included in the hedging experiment by selling options daily instead of weekly. However this would probably not provide any additional insight as hedging errors would likely be highly correlated.

is adjusted appropriately. On the option expiration date, positions in foreign currency are closed and converted to US dollars. The unhedged net cash flow of the option equals the option proceeds less any liability from option exercise. The hedged net cash flow is the sum of the unhedged net cash flow and the profit or loss on the hedging activity.

Table 5 summarizes the results of the hedging experiment. For each model, several statistics are computed for each exchange rate. These statistics are computed over all options, as well as over each delta separately. Listed in Table 5 are the average net cash flow and standard deviation of net cash flow resulting from the sale of a single option. Also listed is the percentage of option positions that result in positive net cash flows. The standard deviation of the cash flow generated by selling options is a measure of the hedge effectiveness. The lower the standard deviation, the better the hedge.

Panel A shows the results for the regime-switching model. The average profit is \$1092.7 per GBP option, but $-\$1162.6$ per JPY option. All deltas have positive profits on average for the GBP, whereas the out-of-the-money calls lose money for the JPY. The JPY appreciated significantly versus the dollar towards the end of the sample. Apparently the hedging strategy did not provide sufficient insurance (in the form of long positions in the JPY) to counter losses due to the exercise of calls. Panel B shows that the GARCH model generates average profits and standard deviations that are very similar to the regime-switching model. This result suggests that in practice, the regime-switching and GARCH option valuation models generate very similar hedge ratios.

As listed in Panel C, the jump-diffusion model generates an average profit for the GBP options of \$1140.5, similar to the other two models, but has a standard deviation almost 25% smaller. Since standard deviation is the measure of hedging effectiveness, the jump-diffusion model is superior to the other two. For the JPY, the average loss for the jump-diffusion model is about 30% smaller than for the other two models, though the standard deviation is slightly larger. In sum, the jump-diffusion model performs better than the other two.

Interestingly, as listed in Panel D, the smile strategy clearly dominates the regime-switching and GARCH models for the GBP options, with an average profit about 11% higher and a standard deviation over 20% lower. The smile strategy performs about the same as the jump-diffusion model for the GBP options. For the JPY options, however, the smile model is dominated by the other three models.

These results suggest that the optimal hedging strategy differs across the exchange rates. For the GBP, the jump-diffusion and smile models perform comparably, both offering substantial improvements over the regime-switching and GARCH models. For the JPY, the smile model is the worst out of the four competitors considered here. Similarly, in an out-of-sample volatility forecasting experiment, Ederington and Guan (1999) find that no single time series model performs consistently well across markets and forecasting horizons. They find that an exponentially-weighted moving average (EWMA) volatility estimator generally outperforms a number of GARCH models. This result suggests that, in the hedging experiment described above, the

Table 5
Results of hedging experiment

Panel A. Regime-switching		JPY		JPY		JPY	
GBP		Average	% Position	Standard deviation	% Position	Standard deviation	% Position
All	1,092.7	-1,162.6	69%	2,707.7	65%	6,464.3	65%
OTMP	362.1	1,102.7	76%	544.0	100%	430.8	100%
NTMP	2,326.3	940.1	60%	3,882.4	60%	3,794.3	60%
ATMC	1,268.7	134.0	64%	3,152.8	60%	5,365.7	60%
NTMC	1,017.9	-2,762.7	72%	2,721.0	48%	7,951.3	48%
OTMC	488.3	-5,227.1	72%	1,219.2	56%	8,486.6	56%
Panel B. GARCH		JPY		JPY		JPY	
GBP		Average	% Position	Standard deviation	% Position	Standard deviation	% Position
All	1,091.6	-1,135.0	70%	2,686.1	66%	6,496.7	66%
OTMP	345.6	1,100.4	80%	534.8	100%	441.7	100%
NTMP	2,333.6	1,129.7	60%	3,848.5	60%	3,684.8	60%
ATMC	1,268.7	254.5	64%	3,141.1	60%	5,138.4	60%
NTMC	1,016.1	-2,741.8	72%	2,730.4	52%	7,761.6	52%
OTMC	493.8	-5,417.6	76%	1,078.7	56%	8,811.1	56%
Panel C. Jump-diffusion		JPY		JPY		JPY	
GBP		Average	% Position	Standard deviation	% Position	Standard deviation	% Position
All	1,140.5	-957.3	79%	2,058.6	69%	7,053.1	69%
OTMP	518.0	1,229.7	92%	479.3	100%	376.7	100%
NTMP	2,441.9	1,313.0	80%	2,588.4	68%	3,771.0	68%
ATMC	1,544.8	1,001.8	72%	2,279.1	60%	5,100.6	60%
NTMC	878.0	-2,260.7	64%	2,444.8	60%	7,935.8	60%
OTMC	320.0	-6,070.4	88%	316.6	56%	10,186.1	56%

(continued on next page)

Table 5 (continued)

Panel D. Smile		GBP			JPY		
		Average	Standard deviation	% Position	Average	Standard deviation	% Position
All		1,215.3	2,100.7	77%	-1,346.9	6,743.1	58%
OTMP		424.0	557.1	80%	537.4	886.9	64%
NTMP		2,573.3	2,579.7	80%	52.6	4,057.5	60%
ATMC		1,794.9	2,216.6	76%	163.8	5,875.3	64%
NTMC		1,123.8	2,445.9	72%	-2,247.4	8,711.2	56%
OTMC		160.2	525.7	76%	-5,240.6	8,713.9	48%

Listed are the results of a hedging experiment that tests the performance of four hedging strategies. The first half of the option sample is used to imply parameters of regime-switching, GARCH, and jump-diffusion models. Once every week in the second half of the option sample a cross-section of five one month options is sold: OTMP and OTMC are out of the money puts and calls with Black-Scholes deltas of -5 and 5% respectively; NTMP and NTMC are near the money puts and calls with deltas of -35 and 35% respectively; and ATMC are at the money calls with deltas of 50%. The positions are held to maturity and delta-hedged daily. The hedge ratios are updated daily using four models: regime-switching, GARCH, and jump-diffusion, all using implied parameters, and a "Smile" model that is the Black-Scholes using a different volatility at each delta. For the regime-switching strategy, regime probabilities are updated daily throughout the options' lives using daily exchange rate returns. For the GARCH strategy, GARCH volatilities are updated daily as well using daily exchange rate returns. For the jump-diffusion strategy, instantaneous volatilities are updated every two weeks and are inferred from option prices. For the Smile strategy, hedge ratios are determined each day based on that day's option prices and the Black-Scholes model. Listed for each model are the average and standard deviation of net cash flow generated per option, and the percentage of options for which the cash flow is positive.

jump-diffusion model may perform even better if the instantaneous volatility is updated using some form of EWMA scheme.

5. Conclusions

This paper compares the performance of three competing option valuation models in the foreign exchange market, one based on a regime-switching process for exchange rate changes, one based on a GARCH process, and the other based on a jump-diffusion. Three tests are executed based on the parameter values implicit in option prices. The first test compares in-sample fit of option prices. The GARCH model is superior to the regime-switching model for the GBP, but the two models perform about the same for the JPY. Both models dominate an ad hoc valuation model based on Black–Scholes, but both are dominated by the jump-diffusion model. The second test compares out-of-sample fit of option prices. Here the GARCH model performs better than the regime-switching model for both exchange rates, and again both are superior to an ad hoc valuation model. The jump-diffusion model offers only modest improvement over the others for the GBP options. The third test simulates the role of a market maker who sells options and uses the competing valuation models to hedge the positions. In this test, the GARCH and regime-switching models perform almost identically. An ad hoc strategy is superior for the GBP and is inferior for the JPY. The jump-diffusion model performs as well as the ad-hoc strategy for the GBP and is superior to all others for the JPY options.

In summary, the regime-switching, GARCH, and jump-diffusion option valuation models are superior to simpler models when fitting option prices. The jump-diffusion model dominates, perhaps due to the degrees of freedom awarded to it here in fitting time-varying volatility. As simulated in the hedging experiment, however, simpler models can generate comparable performance in practice. In addition, since the results appear to differ across the exchange rates, optimal performance may require stress-testing a variety of models on each exchange rate separately.

Acknowledgements

Most of this paper was written while the first author was an assistant professor at the David Eccles School of Business, University of Utah. The authors thank Goldman Sachs for supplying the option quotes used in this study. David Bates kindly shared his computer code for valuing options in a jump-diffusion model. Bob Whaley, Tongshu Ma, and an anonymous referee provided many helpful comments and suggestions.

Appendix A. Option Valuation in the Regime-Switching Model

This appendix shows how to construct the five-branch (pentanomial) lattice used to value options when log exchange rate changes, r , are governed by the regime-switching process studied in this paper. The process is:

$$r_t | \Phi_{t-1} \sim \begin{cases} N(\mu, \sigma_1) & \text{if } R_t = 1 \\ N(\mu, \sigma_2) & \text{if } R_t = 2 \end{cases} \quad (1)$$

where Φ denotes the information set and R denotes regime, which is governed by the transition probability matrix Π . This matrix is given by:

$$\Pi = \begin{bmatrix} \pi_1 & 1 - \pi_1 \\ 1 - \pi_2 & \pi_2 \end{bmatrix} \quad (2)$$

where $\pi_1 = \Pr(R_{t+1} = 1 | R_t = 1)$ is the probability of staying in regime 1 from one period to the next, and $\pi_2 = \Pr(R_{t+1} = 2 | R_t = 2)$ is the probability of staying in regime 2. The regime-switching lattice matches the first two moments implied by each regime's distribution, and reflects the regime uncertainty and switching possibilities of the regime-switching model.

To determine step sizes, one first calculates the binomial lattice branch probabilities and step sizes for both regimes. Two equations and two unknowns are relevant. The mean equation sets the expected return implied by the lattice equal to the distribution's mean:

$$pe^\delta + (1-p)e^{-\delta} = e^{\mu dt} \quad (3)$$

where p denotes the probability of traveling along the upper branch, δ is the continuously compounded rate of change after traveling along the upper branch, and dt corresponds to the duration of one time step in the lattice. The variance equation sets the variance implied by the lattice equal to the distribution's variance:

$$p(\delta)^2 + (1-p)(-\delta)^2 - \mu^2 dt^2 = \sigma^2 dt. \quad (4)$$

These equations can be solved to yield the following values for p and δ :

$$p = \frac{e^{\mu dt} - e^{-\delta}}{e^\delta - e^{-\delta}} \quad \delta = \sqrt{\sigma^2 dt + \mu^2 dt^2}. \quad (5)$$

The computer code for these steps is:

```
STEPH=SQRT((VOLH**2.0)*DT+(RATE*DT)**2.0)
STEPL=SQRT((VOLL**2.0)*DT+(RATE*DT)**2.0)
```

where VOLH (VOLL) indicates the volatility of the high (low) volatility regime and STEPH (STEPL) is the corresponding step size.

One of the regime's step sizes is now increased such that the step sizes of the two regimes are in a 1:2 ratio. The smaller step size will be increased if it is less than one-half the larger step size; otherwise the larger step size will be increased. The additional restriction imposed on the lattice is accommodated by adding a third branch to the adjusted regime. Suppose, for example, that the low volatility step size is less than one-half the high volatility step size, so the low volatility regime's step size will be increased. The conditional branch probabilities of the high volatility regime remain unchanged from the binomial results. The step size of the low vola-

tility regime is set equal to one-half the step size of the high volatility regime, $\delta_l = \delta_h/2$. Next, the three conditional branch probabilities of the low volatility regime are determined by matching the moments implied by the lattice to the moments implied by the distribution of the low volatility regime. For the low volatility regime, the mean equation is:

$$p_{l,u}e^{\delta_h/2} + p_{l,m}e^0 + (1-p_{l,u}-p_{l,m})e^{-\delta_h/2} = e^{\mu dt} \quad (6)$$

where the subscripts u , m , and d denote the up, middle, and down branches. The corresponding variance equation is:

$$p_{l,u}(\delta_h/2)^2 + (1-p_{l,u}-p_{l,m})(-\delta_h/2)^2 - \mu^2 dt^2 = \sigma_l^2 dt. \quad (7)$$

The solutions to these two equations are:

$$p_{l,u} = \frac{e^{\mu dt} - e^{-\delta_h/2} - p_{l,m}(1 - e^{-\delta_h/2})}{e^{\delta_h/2} - e^{-\delta_h/2}}, p_{l,m} = 1 - \left(\frac{4}{\delta_h^2}\right)(\sigma_l^2 dt + \mu^2 dt^2), p_{l,d} = 1 \quad (8)$$

$$-p_{l,u} - p_{l,m}$$

It can be shown that for small enough dt , these probabilities will lie between 0 and 1. (Proof available from the author.) The computer code for these steps is:

```

IF(STEPL.GE.(.5*STEPH))THEN
  CASE=1
ELSE
  CASE=2
ENDIF
IF(CASE.EQ.1)THEN
  DELTA =STEPL
  BPH(2)=1.0-.25*((STEPH/STEPL)**2.0)
  BPH(1)=(EXP(RATE*DT)-EXP(-2.0*STEPL)-BPH(2)*
    (1.0-EXP(-2.0*STEPL)))/(EXP(2.0*STEPL)-
    EXP(2.0*STEPL))
  BPH(3)=1.0-BPH(1)-BPH(2)
  BPL(1)=(EXP(MU*DT)-EXP(-STEPL))/(EXP(STEPL)-EXP(-STEPL))
  BPL(2)=0.0
  BPL(3)=1.0-BPL(1)
ELSE
  DELTA =.5*STEPH
  BPL(2)=1.0-4.0*((STEPL/STEPH)**2.0)
  BPL(1)=(EXP(MU*DT)-EXP(-.5*STEPH)-BPL(2)*
    (1.0-EXP(-.5*STEPH)))/
    (EXP(.5*STEPH)-EXP(-.5*STEPH))
  BPL(3)=1.0-BPL(1)-BPL(2)
  BPH(1)=(EXP(MU*DT)-EXP(-STEPH))/(EXP(STEPH)-EXP(-STEPH))
  BPH(2)=0.0
  BPH(3)=1.0-BPH(1)

```

ENDIF

When using this code to value currency options, MU is set to the difference between the domestic and foreign interest rates, since in a risk-neutral framework this is the expected rate of return on a unit of foreign currency.

In the regime-switching model, time-varying regime probabilities and the possibility of switching regimes complicate the computation of expected option payoffs in the lattice. To simplify matters, two conditional option values are calculated at each node, where the conditioning information is the governing regime. Options are valued in the standard way, iterating backward from the terminal array of nodes. For the terminal array, the two conditional option values are the same at each node. They are simply the maximum of zero and the option's exercise proceeds. For earlier nodes, conditional option values will depend on regime persistence since the persistence parameters are equivalent to future regime probabilities in a conditional setting. The computer code for these steps for a European-style call option is:

```

N=STEP
DO I=1,4*N+1
  EOPT(I,1)=MAX(0,S*EXP((2*N+1-I)*DELTA)-X)
  EOPT(I,2)=EOPT(I,1)
ENDDO
DO J=1,N
  DO I=1,(N-J)*4+1
    EEOH =BPH(1)*EOPT(I,1)+BPH(2)*EOPT(I+2,1)+
          BPH(3)*EOPT(I+4,1)
    EEOL =BPL(1)*EOPT(I+1,2)+BPL(2)*EOPT(I+2,2)+
          BPL(3)*EOPT(I+3,2)
    TEEOH=TPH*EEOH + (1.0-TPH)*EEOL
    TEEOL=TPL*EEOL + (1.0-TPL)*EEOH
    TEOPT(I,1)=EXP(-RF*DT)*TEEOH
    TEOPT(I,2)=EXP(-RF*DT)*TEEOL
  ENDDO
  DO I=1,(N-J)*4+1
    EOPT(I,1)=TEOPT(I,1)
    EOPT(I,2)=TEOPT(I,2)
  ENDDO
ENDDO
OPT(1)=EOPT(1,1)
OPT(2)=EOPT(1,2)

```

In this code, TPH (TPL) is the persistence of the high (low) volatility regime, and RF is set to the domestic risk-free interest rate.

When the current regime is known, one of the two conditional option values at the seed node in the lattice is correct. When regimes are not known with certainty, the European-style option value is a weighted average of the two conditional option

values at the seed node, where initial regime probabilities are the weights. The backward iteration technique is simply a way of computing probabilities of terminal option payoffs consistent with initial regime probabilities and the probability of switching regimes at each intermediate node.

Appendix B. Option Valuation in the GARCH Model

This appendix shows how to construct a lattice used to value options when log exchange rate changes, r , are governed by a GARCH(1,1) process. The variable NSTEP is the number of time steps in the option's life. Each time step is further separated into N sub steps. Each sub step represents the evolution of the underlying variable with a trinomial branch. Thus, from a particular node at a particular time step, there are $2N+1$ adjacent nodes in the following time step. Within each time step, the step size of each sub step is restricted to be an integer multiple ETA of a fixed quantity equal to GAMMAN in the code below, which is determined by the variance at the time the option is valued. ETA is selected within each time step in order to generate branch probabilities between zero and one.

The first task in the GARCH option valuation procedure is to establish the topology of the lattice by going forward in time. Consistent with the discrete-time nature of the model, variance is constant over each time step in the lattice, then is updated according to the GARCH scheme depending on the change in the exchange rate, as represented by the position in the lattice. In the GARCH process, a path-dependency arises because the updated variance is a function of which particular path was taken over a given time step. The solution used in Ritchken and Trevor (1999) is to record at each node the maximum and minimum variance that is possible given all the paths that lead to the node. This procedure is executed while the topology of the lattice is being established going forward. Then, when proceeding to the next array of nodes in the following time step, K equally spaced variances between the maximum and minimum are used as initial variances at each node.

In the computer code that follows, the topology of the lattice is established. The variable H0 is the level of variance on the date the option is valued. The parameters B0, B1, and B2 are the GARCH parameters. The parameter MAXV is a user-defined maximum variance that caps the level of variance achieved in the lattice. This is necessary to avoid explosive variance. The subroutine GETETA is called at each time step to find the minimum value of the integer variable ETA which results in branch probabilities that lie between zero and one. The matrix H stores intermediate levels of variance in the lattice. The vectors MUP and MDOWN store the position of the extreme upper and lower levels of the lattice at each time step. The computer code is:

```
GAMMA=SQRT(H0)
CALL GETETA(ETA)
GAMMAN=GAMMA/SQRT(N)
DO I=N,-N,-1
```

```

EPSILON=I*ETA*GAMMAN-(RF-.5*H0)
H(1,I*ETA,1)=MIN(MAXV,B0+B1*H0+B2*(EPSILON)**2)
H(1,I*ETA,K)=H(1,I*ETA,1)
ENDDO
MUP(1)=N*ETA
MDOWN(1)=N*ETA
DO I=1,NSTEP-1
  DO J=MUP(I),-MDOWN(I),-1
    DELTA=(H(I,J,K)-H(I,J,1))/(K-1)
    DO L=2,K-1
      H(I,J,L)=H(I,J,1)+L*DELTA
    ENDDO
    DO L=1,K
      HH=H(I,J,L)
      CALL GETETA(ETA)
      IF(J+N*ETA.GT.MUP(I+1))THEN
        MUP(I+1)=J+N*ETA
      ENDIF
      IF(J-N*ETA.LT.-MDOWN(I+1))THEN
        MDOWN(I+1)=-(J-N*ETA)
      ENDIF
      DO M=N,-N,-1
        EPSILON=M*ETA*GAMMAN-(RF-.5*HH)
        TH=MIN(MAXV,B0+B1*HH+B2*(EPSILON)**2)
        IF(H(I+1,J+M*ETA,1).EQ.0.0)THEN
          H(I+1,J+M*ETA,1)=TH
        ELSE
          IF(TH.LT.H(I+1,J+M*ETA,1))THEN
            H(I+1,J+M*ETA,1)=TH
          ENDIF
        ENDIF
        IF(H(I+1,J+M*ETA,K).EQ.0.0)THEN
          H(I+1,J+M*ETA,K)=TH
        ELSE
          IF(TH.GT.H(I+1,J+M*ETA,K))THEN
            H(I+1,J+M*ETA,K)=TH
          ENDIF
        ENDIF
      ENDDO
    ENDDO
  ENDDO
ENDDO

```

Option values are computed with a backward induction technique. At each node, K conditional option values are computed corresponding to K variance levels. Expec-

tations are taken over the $2N+1$ possible adjacent nodes in the next time step. Given that a particular node was reached, the GARCH variance is computed in order to determine the resulting conditional option value. Since only K conditional option values are stored, the conditional expectation involves interpolation using the subroutine GETOPVAL. The subroutine GETPROBS computes the probability of traversing from a particular node in a particular time to each of the $2N+1$ adjacent nodes in the next time step. The computer code is:

```

DO I=MUP(NSTEP),-MDOWN(NSTEP),-1
  S=EXP(LOG(S0)+GAMMAN*I)
  DO L=1,K
    OVAL(I,L)=MAX(0,S-X)
  ENDDO
ENDDO
DO I=NSTEP-1,1,-1
  DO J=MUP(I),-MDOWN(I),-1
    S=EXP(LOG(S0)+GAMMAN*J)
    DO L=1,K
      HH=H(I,J,L)
      CALL GETETA(ETA)
      DO M=N,-N,-1
        EPSILON=(M*ETA*GAMMAN-(RF-.5*HH))/SQRT(HH)
        TH=MIN(MAXV,B0+B1*HH+B2*HH*(EPSILON)**2)
        CALL GETOPVAL(TH,OVAL,C)
      ENDDO
      CALL GETPROBS(HH,ETA,PR)
      DO M=N,-N,-1
        TOVAL(J,L)=TOVAL(J,L)+DEXP(-RF)*PR(M)*C(M)
      ENDDO
    ENDDO
  ENDDO
  DO J=MUP(I),-MDOWN(I),-1
    DO L=1,K
      OVAL(J,L)=TOVAL(J,L)
      TOVAL(J,L)=0.0
    ENDDO
  ENDDO
ENDDO

```

References

- Bank for International Settlements, 1999. Press release.
 Bates, D.S., 1996a. Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *Review of Financial Studies* 9 (1), 69–107.

- Bates, D.S., 1996b. Dollar jump fears, 1984–1992: Distributional abnormalities in currency futures options. *Journal of International Money and Finance* 15 (1), 65–93.
- Bates, D.S., 2000. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics* 94, 181–238.
- Bekaert, G., Hodrick, R.J., 1993. On biases in the measurement of foreign exchange risk premiums. *Journal of International Money and Finance* 12, 115–138.
- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81 (3), 637–659.
- Black, F., 1976. The pricing of commodity options. *Journal of Financial Economics* 3, 167–179.
- Bollen, N.P., Gray, S.F., Whaley, R.E., 2000. Regime-switching in foreign exchange rates: Evidence from currency option prices. *Journal of Econometrics* 94, 239–276.
- Bollen, N.P., 1998. Valuing options in regime-switching models. *Journal of Derivatives* 6 (1), 38–49.
- Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31, 307–327.
- Campa, J.M., Chang, P.H., 1995. Testing the expectations hypothesis on the term structure of volatilities in foreign exchange options. *Journal of Finance* 50 (2), 529–547.
- Campa, J.M., Chang, P.H., 1998. The forecasting ability of correlations implied in foreign exchange options. *Journal of International Money and Finance* 17 (6), 855–880.
- Campa, J.M., Chang, P.H., Reider, R.L., 1998. Implied exchange rate distributions: evidence from OTC option markets. *Journal of International Money and Finance* 17 (1), 117–160.
- Duan, J., 1995. The GARCH option pricing model. *Mathematical Finance* 5 (1), 13–32.
- Duan, J., Popova, I., Ritchken, P., 1999. Option pricing under regime shifting. Working paper.
- Ederington, L., Guan, W., 1999. Forecasting volatility. Working paper.
- Engel, C.E., Hamilton, J.D., 1990. Long swings in the dollar: are they in the data and do markets know it? *American Economic Review* 80, 689–713.
- Engle, R.F., 1982. Autoregressive conditional heteroskedasticity with estimates of UK inflation. *Econometrica* 50, 987–1008.
- Gallant, A.R., Hsieh, D.A., Tauchen, G.E., 1991. On fitting a recalcitrant series: The pound/dollar exchange rate, 1974–1983. In: Barnett, W.A., Powell, J., Tauchen, G.E. (Eds.), *Nonparametric and Semiparametric Methods in Econometrics and Statistics*. Cambridge University Press, Cambridge, pp. 199–240.
- Garcia, R., 1998. Asymptotic null distribution of the likelihood ratio test in Markov-switching models. *International Economic Review* 39, 763–788.
- Garman, M.B., Kohlhagen, S.W., 1983. Foreign currency option values. *Journal of International Money and Finance* 2, 231–237.
- Grabbe, J.O., 1983. The pricing of call and put options on foreign exchange. *Journal of International Money and Finance* 2, 239–254.
- Gray, S.F., 1996. Modeling the conditional distribution of interest rates as a regime-switching process. *Journal of Financial Economics* 42, 27–62.
- Hamilton, J.D., 1988. Rational expectations econometric analysis of changes in regime: an investigation of the term structure of interest rates. *Journal of Economic Dynamics and Control* 12, 385–423.
- Hamilton, J.D., 1989. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57, 357–384.
- Hamilton, J.D., 1990. Analysis of time series subject to changes in regime. *Journal of Econometrics* 45, 39–70.
- Hamilton, J.D., 1994. *Time series analysis*. Princeton University Press, Princeton, NJ.
- Hamilton, J.D., Susmel, R., 1994. Autoregressive conditional heteroskedasticity and changes in regime. *Journal of Econometrics* 64, 307–333.
- Hansen, B.E., 1992. The likelihood ratio test under nonstandard conditions: Testing the markov switching model of GNP. *Journal of Applied Econometrics* 7, S61–S82.
- Jorion, P., 1988. On jump processes in the foreign exchange and stock markets. *Review of Financial Studies* 1 (4), 427–445.
- Jorion, P., 1995. Predicting volatility in the foreign exchange market. *Journal of Finance* 50 (2), 507–528.

- Naik, V., 1993. Option valuation and hedging strategies with jumps in the volatility of asset returns. *Journal of Finance* 48 (5), 1970–1984.
- Rubinstein, M., 1994. Implied binomial trees. *Journal of Finance* 49 (3), 771–817.
- Ritchken, P., Trevor, R., 1999. Pricing options under generalized GARCH and stochastic volatility processes. *Journal of Finance* 54 (1), 337–402.
- West, K.D., Cho, D., 1995. The predictive ability of several models of exchange rate volatility. *Journal of Econometrics* 69, 367–391.