

Stochastic Skew Models for FX Options

PETER CARR

Bloomberg LP and Courant Institute, NYU

LIUREN WU

Zicklin School of Business, Baruch College

Special thanks to Bruno Dupire, Harvey Stein, Arun Verma, and as
always, Dilip Madan.

Overview

- There is a huge market for foreign exchange (FX), much larger than the equity market ...
As a result, an understanding of FX dynamics is economically important.
- FX option prices can be used to understand risk-neutral FX dynamics, i.e. how the market prices various path bundles.
- Despite their greater economic relevance, FX options are not as widely studied as equity index options, probably due to the fact that the FX options market is now primarily OTC.
- Nonetheless, we obtained OTC options data on 2 underlying currency pairs (JPYUSD, GBPUSD) over 8 years.

Overview (Con'd)

- We used our data to study the variation of FX option prices in the cross section and over calendar time.
- Like equity options, FX option implied volatilities vary stochastically over calendar time, and there is a smile in FX option implieds i.e. the convexity measure is always positive.
- This suggests that stochastic volatility is needed to explain risk-neutral currency dynamics, as shown for example by Bates (1996).
- However, unlike equity options, there is a substantial variation in the skewness measure as well. For both currency pairs, the skewness measure switches signs several times over our 8 year history.
- This suggests that stochastic *skewness* is also needed to explain risk-neutral currency dynamics.

What We Do

- FX option prices apparently reflect at least 3 sources of uncertainty:
 1. stochastic FX rate (i.e. random risk-neutral mean)
 2. stochastic volatility (i.e. implieds are affected by at least a 2nd source of uncertainty besides the FX rate)
 3. stochastic skewness (i.e. the asymmetry of the risk-neutral distribution changes randomly over calendar time).
- The classic Black Scholes model handles 1).
- SV models such as Heston or Bates handles 1) and 2).
- We explore several potential modelling approaches for capturing 1), 2), and 3).

What We Do (Con'd)

- We find that the most tractable approach for pricing standard European options is to employ (stochastically) time-changed Lévy processes (more on this later).
- Since our new models capture stochastic mean, stochastic volatility, *and* stochastic skewness, we christen these models as stochastic skew models (SSM).
- We estimate our SSM's and compare them to older SV models such as Heston (1993) and Bates (1996).
- We find substantially improved pricing performance over Bates for the same number of parameters.

OTC FX Option Market Conventions

- We focus on the valuation of standard European options, although barrier options are also liquid. American options are almost never traded.
- Perhaps due to the dominance of European options, the OTC FX options market uses the Black Scholes (BS) formula in several ways:
 1. Quotes are in terms of BS model implied volatilities rather than on option prices directly.
 2. Quotes are provided at a fixed BS delta rather than a fixed strike. In particular, the liquidity is mainly at 5 levels of delta:
10 δ Put, 25 δ Put, 0 δ straddle, 25 δ call, 10 δ call.
 3. Trades include both the option position and the underlying, where the position in the latter is determined by the BS delta.

A Review of the Black-Scholes Formulae

- BS call and put pricing formulae:

$$\begin{aligned}c(K, \tau) &= e^{-rf\tau} S_t N(d_+) - e^{-rd\tau} K N(d_-), \\p(K, \tau) &= -e^{-rf\tau} S_t N(-d_+) + e^{-rd\tau} K N(-d_-),\end{aligned}$$

with

$$d_{\pm} = \frac{\ln(F/K)}{\sigma\sqrt{\tau}} \pm \frac{1}{2}\sigma\sqrt{\tau}, \quad F = Se^{(rd-rf)\tau}.$$

- BS Delta

$$\delta(c) = e^{-rf\tau} N(d_+), \quad \delta(p) = -e^{-rf\tau} N(-d_+).$$

$|\delta|$ is *roughly* the probability that the option will expire in-the-money.

- BS Implied Volatility (IV): the σ input in the BS formula that matches the BS price to the market quote.

Data

- We have 8 years of weekly data from January 1996 to January 2004 (419 weeks).
- At each date, we have 8 maturities: 1w, 1m, 2m, 3m, 6m, 9m, 12m, 18m.
- At each maturity, we have five option quotes.
- Thus all together, we have 16,760 quotes for each currency pair.
- We also have two currency pairs: JPYUSD and GBPUSD.

OTC Currency Option Quotes

- The five option quotes at each maturity are:
 1. Implied volatility of a delta-neutral straddle (ATMV)
 - A straddle is the sum of a call with a put at the same strike.
 - Delta-neutral means $\delta(c) + \delta(p) = 0 \Rightarrow N(d_+) = 0.5 \Rightarrow d_+ = 0$.
 - $ATMV \equiv IV(50 \delta c)$ ($= IV(-50 \delta p)$ by put call parity).
 2. 25-delta risk reversal (RR25)
 - $RR25 \equiv IV(25\delta c) - IV(25\delta p)$.
 - Captures the slope of the smile, which proxies the **skewness** of the risk-neutral return distribution.
 3. 25-delta strangle margin (a.k.a butterfly spread) (SM25)
 - A strangle is the sum of a call and a put at two different strikes.
 - $SM25 \equiv (IV(25\delta c) + IV(25\delta p)) / 2 - ATMV$.
 - Captures the curvature of the smile (**kurtosis of the distribution**).
 4. 10-delta RR and 10-delta SM.

Convert Quotes to Option Prices

- Convert the quotes into implied volatilities at the five deltas:


$$\begin{aligned}IV(0\delta_s) &= ATMV; \\IV(25\delta_c) &= ATMV + RR25/2 + SM25; \\IV(25\delta_p) &= ATMV - RR25/2 + SM25; \\IV(10\delta_c) &= ATMV + RR10/2 + SM10; \\IV(10\delta_p) &= ATMV - RR10/2 + SM10.\end{aligned}$$

- Download LIBOR and swap rates on USD, JPY, and GBP to generate the relevant yield curves (r_d, r_f) .
- Convert deltas into strike prices

$$K = F \exp \left[\mp IV(\delta, \tau) \sqrt{\tau} N^{-1} (\pm e^{r_f \tau} \delta) + \frac{1}{2} IV(\delta, \tau)^2 \tau \right].$$


- Convert implied volatilities into out-of-the-money option prices using the BS formulae.

Time Series of Implied Volatilities



Stochastic volatility—Note the impact of the 1998 hedge fund crisis on dollar-yen. During the crisis, hedge funds bought call options on yen to cover their yen debt.


Stylized Evidence I



1m (solid), 3m (dashed), 12m (dash-dotted)

- The mean implied volatility smile is relatively symmetric ...
- The smile (kurtosis) persists with increasing maturity.

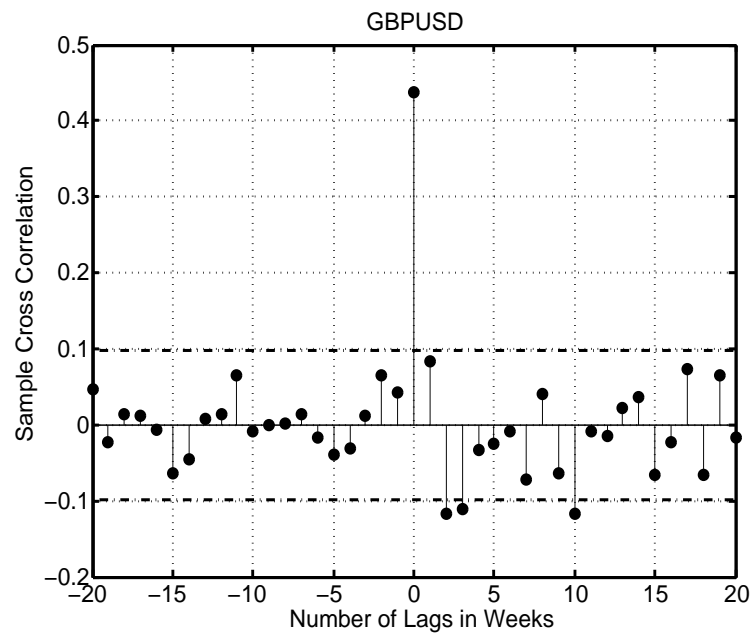
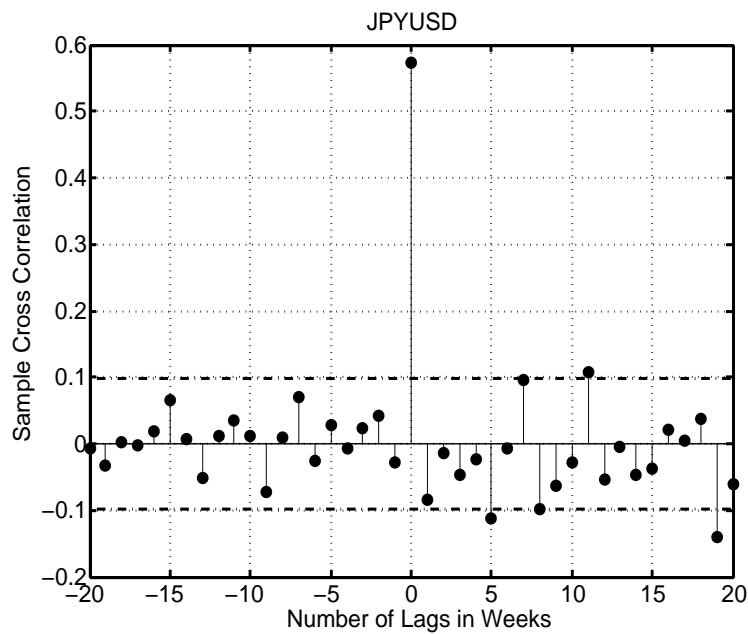
Stylized Evidence II



RR10 (solid), SM10(dashed)

- The strangle margin (kurtosis measure) is stable over time at $\approx 10\%$ of ATMV.
- But the risk reversal (skewness measure) varies greatly over time
 \Rightarrow Stochastic Skew.

Stylized Evidence III



- Changes in risk reversals are positively correlated with contemporaneous currency returns ...
- But there are no lead-lag effects.

How Has the Literature Priced FX Options?

- The literature documents the fact that FX option implied volatilities vary randomly over calendar time and display convexity across the moneyness measure.
- The literature has used two ways to generate models consistent with these observations:

1) Jump-diffusion model: e.g., Merton (1976)

$$\frac{dS_t}{S_{t-}} = (r_d - r_f)dt + \sigma dW_t + \int_{-\infty}^{\infty} (e^x - 1)[\mu(dx, dt) - \lambda n_x(\mu_j, \sigma_j)dt].$$

- The arrival of jumps is controlled by a Poisson process with arrival rate λ .
- Given that a jump has occurred, the jump in the return (log price relative) x is normally distributed, $n_x(\mu_j, \sigma_j) \equiv \frac{e^{-\frac{1}{2}\left(\frac{x-\mu_j}{\sigma_j}\right)^2}}{\sqrt{2\pi}\sigma_j}$.
- Nonzero σ_j generates a smile (curvature) and nonzero μ_j generates skew (slope) ...

How has the Literature Priced FX Options?

- Recall that FX option implied volatilities vary randomly over calendar time and display convexity across the moneyness measure.
- A second way that the literature has captured these observations is stochastic diffusion volatility: e.g., Heston (1993)

$$\begin{aligned}\frac{dS_t}{S_t} &= (r_d - r_f)dt + \sqrt{v_t}dW_t, \\ dv_t &= \kappa(\theta - v_t)dt + \sigma_v\sqrt{v_t}dZ_t, \quad dW_t dZ_t = \rho dt\end{aligned}$$

- Positive vol of vol (σ_v) generates a smile,
- Nonzero correlation (ρ) generates skew of the same sign...

How Do The Two Methods Differ?

- Jump diffusions induce short term smiles and skews that dissipate quickly with increasing maturity due to the central limit theorem.
- Stochastic volatility induces smiles and skews that increase as maturity increases over the horizon of interest.
- Recall that the strangle margin (convexity measure) is more or less constant as maturity increases.
- The Bates (1996) model is a generalization of both Heston and Merton which accomodates this observation.
 - The Bates model generalizes Heston by adding Merton's IID jumps.
 - Alternatively, the Bates model generalizes Merton's jump diffusion model by making the diffusion volatility stochastic.
- At short maturity, the desired U-shaped volatility profile is generated by jumps, while at longer maturities, it is generated by SV.

Consistency with Stochastic Skewness

- Recall that the literature has used jump diffusions and/or stochastic volatility to capture the fact that FX option implied volatilities vary randomly over calendar time and display convexity across the moneyness measure.
- Both approaches can also generate deterministic skewness, but neither approach generates stochastic skewness.
- Since neither the Merton model nor the Heston model can generate stochastic skewness, neither does the Bates model.

Adapting Existing Approaches?

- For a jump diffusion model such as Merton, one can induce stochastic skewness by randomizing the mean jump size.
- Alternatively, for an SV model such as Heston or Bates, one can induce stochastic skewness by randomizing the correlation between returns and increments in volatility.
- Neither approach is tractable to our knowledge.
- The key to developing a tractable approach for handling stochastic skew is to regard the pricing problem from a more general perspective.

Be Wise: Generalize (Part 1)

- We may generalize jump diffusions such as Merton to the wider class of Lévy processes.
- Lévy processes comprise all continuous time stochastic process with IID increments. Common examples include arithmetic Brownian motion, compound Poisson processes, and their sum (which is a jump diffusion).
- Lévy processes also include infinite activity pure jump processes such as Madan's Variance Gamma (VG) model, infinite variation pure jump processes such as stable processes, and generalizations of both such as CGMY.
- The main advantage of placing the infinite activity in the jump component rather than the continuous component is that the latter are more flexible - for example, under finite variation, one can model the up moves and the down moves separately and independently.

The Characteristic Function of a Lévy Process

The Lévy-Khintchine Theorem describes *all* Lévy processes via their characteristic function:

$$\phi_x(u) \equiv E [e^{iuX_t}] = e^{-t\psi_x(u)}, \quad t \geq 0,$$

where the characteristic exponent $\psi_x(u)$, $u \in \mathbb{R}$, is given by

$$\psi_x(u) = -iub + \frac{u^2\sigma^2}{2} - \int_{\mathbb{R}^0} (e^{iux} - 1 - iux1_{|x|<1}) k(x)dx.$$

The triplet (b, σ^2, k) defines the Lévy process X :

- b describes the constant drift of the process.
- σ^2 describes the constant variance rate of the continuous martingale component.
- $k(x)$ describes the jump structure and determines the arrival rate of jumps of size x .

$$\lim_{t \downarrow 0} \frac{P(X_t \in dx | X_0 = 0)}{t} = k(x)dx, \quad x \neq 0.$$

Be Wise: Generalize (Part 2)

- We may generalize SV models such as Heston by regarding them as stochastically time-changed stochastic processes.
- Recall that the usual lognormal volatility has dimensions of one over the square root of time.
- To randomize volatility, we can randomize the denominator rather than the numerator.
- Since only “total variance” matters for European option prices, stochastic volatility can be induced by randomizing the interval over which “total variance” is measured.
- To induce randomness in volatility over any time interval, one randomizes the clock on which the process is run.

The CF of a Time-Changed Brownian Motion

- Since an arithmetic Brownian motion, $dX_t = bdt + \sigma dW_t$, is a Lévy process (the only continuous one), its CF is given by:

$$\phi_x(u) \equiv E [e^{iuX_t}] = e^{-t\psi_x(u)}, \quad t \geq 0,$$

where $X_0 = 0$ and the characteristic exponent $\psi_x(u)$, $u \in \mathbb{R}$, is given by:

$$\psi_x(u) = -iub + \frac{u^2\sigma^2}{2}.$$

- Under a deterministic time change $t \mapsto T(t)$, where $T'(t) > 0$, $dX_t = bT'(t)dt + \sigma\sqrt{T'(t)}dW_t$, the CF generalizes to:
$$\phi_x(u) \equiv E [e^{iuX_t}] = e^{-T(t)\psi_x(u,s)ds}, \quad t \geq 0.$$
- When we stochastically time-change arithmetic Brownian motion, $dX_t = bV_tdt + \sigma\sqrt{V_t}dW_t$, where $V \perp W$ is non-negative, then from an argument due to Hull and White, the CF generalizes to:
$$\phi_x(u) \equiv E [e^{iuX_t}] = E e^{-\psi_x(u)T_t}, \quad t \geq 0, \quad \text{where } T_t = \int_0^t V_s ds.$$
- This argument also works for the more general class of Lévy processes.

Time-Changed Lévy Processes

- A (stochastically) time-changed Lévy process is defined as a Lévy process run on a stochastic clock.
- By randomizing the clock on which the Lévy process runs, we introduce volatility clustering and slow the smile flattening caused by the CLT.
- By time-changing a Lévy process rather than a Brownian motion, we get short term skews and smiles that persist over calendar time.
- CGMY (2003) first introduce time-changed Lévy processes using a stochastic clock which is independent of the Lévy process, and Carr/Wu (2004) show how to handle dependence.

Time-Changed Lévy Process Perspective

The class of time-changed Lévy processes is very general and includes many tractable special cases. In particular, the Merton model and Heston model are both tractable special cases:

- Merton (1976) uses a Lévy martingale $L_t = \sigma W_t + M_t^j$, where M_t^j is a pure jump martingale, *but has no time change*:

$$\ln \left(\frac{S_t}{S_0} \right) = (r_d - r_f)t + L_t - \xi t,$$

where $\xi > 0$ is a concavity correction due to the log.

- Heston (1993) time-changes a Brownian motion, *but has no jumps*:

$$\begin{aligned} \ln \left(\frac{S_t}{S_0} \right) &= (r_d - r_f)t + W_{T_t} - \frac{1}{2}T_t, & v_t &\equiv \partial T_t / \partial t, \\ dv_t &= \kappa(\theta - v_t)dt + \sigma_v \sqrt{v_t} dZ_t, & dW_t dZ_t &= \rho dt. \end{aligned}$$

Stochastic Skew Models

$$\ln \left(\frac{S_t}{S_0} \right) = (r_d - r_f)t + \left(L_{T_t^R}^R - \xi^R T_t^R \right) + \left(L_{T_t^L}^L - \xi^L T_t^L \right),$$

- L_t^R is a Lévy martingale that generates +ve skewness (diffusion + positive jumps).
- L_t^L is a Lévy martingale that generates -ve skewness (diffusion + negative jumps).
- $[T_t^R, T_t^L]$ are random clocks underlying the two Lévy martingales with:
 - $[T_t^R + T_t^L]$ determining total volatility: stochastic
 - $[T_t^R - T_t^L]$ determining skewness (risk reversal): ALSO stochastic.
- We assume that both clocks are continuous over time, so that we can correlate the Brownian motions driving them to the diffusion components of the corresponding Lévy process. This allows (+ve) correlation between ΔRR and returns.
- ξ^R and ξ^L are due to the concavity of the log and are determined by the parameters of the 2 Lévy processes and the 2 random clocks that they run on.

SSMs in the Language of Merton and Heston

$$\begin{aligned}\frac{dS_t}{S_{t-}} &= (r_d - r_f)dt \leftarrow \text{risk-neutral drift} \\ &+ \sigma \sqrt{v_t^R} dW_t^R + \int_0^\infty (e^x - 1) [\mu^R(dx, dt) - k^R(x) dx v_t^R dt] \leftarrow \text{right skew} \\ &+ \sigma \sqrt{v_t^L} dW_t^L + \int_{-\infty}^0 (e^x - 1) [\mu^L(dx, dt) - k^L(x) dx v_t^L dt]. \leftarrow \text{left skew}\end{aligned}$$

- The counting measure μ^R assigns mass to positive jumps only and likewise, the Lévy density $k^R(x)$ has support on $x \in (0, \infty)$.
- The counting measure μ^L assigns mass to negative jumps only and likewise, the Lévy density $k^L(x)$ has support on $x \in (-\infty, 0)$.
- $[v_t^R, v_t^L]$ follow mean reverting square root processes. We refer to them as activity rates and they are chosen here to have long run means of one.

Option Pricing via Fourier Inversion

- Breeden and Litzenberger's results imply that the European call value as a function C of its strike K is obtained by twice integrating the risk-neutral PDF π of the final FX rate S_T :

$$C(K) = e^{-rT} \int_K^\infty \int_L^\infty \pi(M) dM dL.$$

- Letting $k \equiv \ln(K/S_0)$, $\gamma(k) \equiv C(K)$, and $q(\ell)$ be the risk-neutral PDF of the log price relative $X = \ln(S_T/S_0)$, we also have:

$$\gamma(k) = e^{-rT} \int_k^\infty e^\ell \int_\ell^\infty q(m) dm d\ell.$$

- Carr and Madan (1999) derive the following Fourier analog:

$$\mathcal{F}[\gamma](u, T) = e^{-rT} \frac{F_0 \mathcal{F}[q](u - i, T)}{(i - u)u}.$$

- Hence, the Fourier Transform of a call price $\mathcal{F}[\gamma](u, T)$ is analytically related to $\mathcal{F}[q]$, which is the Characteristic Function (CF) of the return.

Option Pricing via FFT

- Recall that:

$$\mathcal{F}[\gamma](u, T) = e^{-rT} \frac{F_0 \mathcal{F}[q](u - i, T)}{(i - u)u}.$$

- When the CF $\mathcal{F}[q]$ is available in closed form, then so is the Fourier Transform of the call.
- The Fast Fourier Transform (FFT) can be used to quickly invert for a strike structure of call values.
- Hence the option valuation problem is effectively reduced to finding the CF of the return in closed form.
- This result is completely general.

Option Pricing Under SSM

- If the 2 the Lévy processes were each independent of their stochastic clocks, then a closed form expression for the CF of the return is obtained as follows:
 1. Derive the CF of the 2 Lévy processes in closed form by a wise choice of the Lévy density. The log of the CF at unit time is called the characteristic exponent.
 2. Derive the Laplace transform (LT) of the 2 random clocks in closed form by a wise choice of the activity rate processes.
 3. Using a conditioning argument similar to one in Hull White (1987), one can show that the CF of a time-changed Lévy process is just the LT of the stochastic clock, evaluated at the characteristic exponent.
- Since the CF of the Lévy process and the LT of the clock are both known in closed form, so is the CF of the time-changed Lévy process.

Option Pricing Under SSM (Con'd)

- Recall that when the clock is independent of the Lévy process, then the CF of the return is just the LT of the clock evaluated at the characteristic exponent.
- Our SSM models assume that each Lévy process is correlated with its stochastic clock. Carr&Wu (2004) show that when calculating CF's of returns, this correlation induces a new measure for calculating the LT of the random clock.

$$e^{iu(r_d-r_f)t} E^{\mathbb{Q}} \left[e^{iu \left(L_{T_t^R}^R - \xi^R T_t^R \right) + iu \left(L_{T_t^L}^L - \xi^L T_t^L \right)} \right] = e^{iu(r_d-r_f)t} E^{\mathbb{M}} \left[e^{-\psi^\top T_t} \right].$$

- This new *leverage-neutral* measure \mathbb{M} is complex-valued.
- The RHS is proportional to the bivariate Laplace transform of the two new clocks $[T_t^R, T_t^L]$, evaluated at the characteristic exponent $[\psi^R, \psi^L]$.
- Just as the change from \mathbb{P} to \mathbb{Q} absorbs risk aversion into the probabilities, the change from \mathbb{Q} to \mathbb{M} absorbs correlation into the probabilities.

Option Pricing Under SSM (Con'd)

- To summarize, the CF of a time-changed Lévy process is always the LT of the random time, evaluated at the characteristic exponent of the Lévy process.
- When there is correlation between the new clock and the Lévy process, then the LT is evaluated using a complex-valued measure. Under no correlation, this measure \mathbb{M} reduces to the usual risk-neutral probability measure \mathbb{Q} .
- Once we have the CF of the return in closed form, we automatically have the FT of the call value in closed form. We then apply FFT to numerically obtain the strike structure of option prices.
- This operation is very fast, allowing the calculation of about a thousand option prices per second.

Our Jump Specification

- Recall that the characteristic exponent $\psi_x(u)$, $u \in \mathbb{R}$, of a Lévy process is:

$$\psi_x(u) \equiv \ln Ee^{iuX_1} = -iub + \frac{u^2\sigma^2}{2} - \int_{\mathbb{R}^0} (e^{iux} - 1 - iux1_{|x|<1}) k(x)dx.$$

- The **Lévy density** $k(x)$ specifies the arrival rate as a function of the jump size:

$$k(x) \geq 0, x \neq 0, \quad \int_{\mathbb{R}^0} (x^2 \wedge 1)k(x)dx < \infty.$$

- To obtain tractable models, choose the Lévy density so that the above integral can be done in closed form.

Our Jump Specification (Con'd)

- For our SSM models, we modelled the Lévy density for the right and left jump processes by an exponentially dampened power law (DPL):

$$k^R(x) = \begin{cases} \lambda \frac{e^{-\frac{|x|}{v_j}}}{|x|^{\alpha+1}}, & x > 0, \\ 0, & x < 0. \end{cases}, \quad k^L(x) = \begin{cases} 0, & x > 0, \\ \lambda \frac{e^{-\frac{|x|}{v_j}}}{|x|^{\alpha+1}}, & x < 0. \end{cases}$$

- The specification originates in Carr, Géman, Madan, Yor (2002), and captures much of the stylized evidence on both equities and currencies (Wu, 2004).
- This 3 parameter specification is both general and intuitive with many interesting special cases:
 1. $\alpha = -1$: Kou's double exponential model (KJ), $< \infty$ **activity**.
 2. $\alpha = 0$: Madan's variance-gamma model (VG), ∞ **activity**, $< \infty$ **variation**.
 3. $\alpha = 1$: Cauchy dampened by exponential functions (CJ), ∞ **variation**.
- The parameter α determines the fine structure of the sample paths.

Characteristic Exponents For Dampened Power Law

Model	Right (Up) Component	Left (Down) Component
KJ	$\psi^D - iu\lambda \left[\frac{1}{1-iv_j} - \frac{1}{1-v_j} \right]$	$\psi^D + iu\lambda \left[\frac{1}{1+iv_j} - \frac{1}{1+v_j} \right]$
VG	$\psi^D + \lambda \ln(1 - iv_j) - iu\lambda \ln(1 - v_j)$	$\psi^D + \lambda \ln(1 + iv_j) - iu\lambda \ln(1 + v_j)$
CJ	$\psi^D - \lambda(1/v_j - iu) \ln(1 - iv_j)$ $+iu\lambda(1/v_j - 1) \ln(1 - v_j)$	$\psi^D - \lambda(1/v_j + iu) \ln(1 + iv_j)$ $+iu\lambda(1/v_j + 1) \ln(1 + v_j)$
CG	$\psi^D + \lambda\Gamma(-\alpha) \left[\left(\frac{1}{v_j}\right)^\alpha - \left(\frac{1}{v_j} - iu\right)^\alpha \right]$ $-iu\lambda\Gamma(-\alpha) \left[\left(\frac{1}{v_j}\right)^\alpha - \left(\frac{1}{v_j} - 1\right)^\alpha \right]$	$\psi^D + \lambda\Gamma(-\alpha) \left[\left(\frac{1}{v_j}\right)^\alpha - \left(\frac{1}{v_j} + iu\right)^\alpha \right]$ $-iu\lambda\Gamma(-\alpha) \left[\left(\frac{1}{v_j}\right)^\alpha - \left(\frac{1}{v_j} + 1\right)^\alpha \right]$

$$\psi^D = \frac{1}{2}\sigma^2 (iu + u^2) \quad \psi_x(u) \equiv \ln Ee^{iuX_1} = \psi^D - \int_{\mathbb{R}^0} (e^{iux} - 1 - iux1_{|x|<1}) k(x)dx.$$

CF of Return as LT of Clock

$$\begin{aligned}
 \phi_s(u) &\equiv E^{\mathbb{Q}} e^{iu \ln(S_T/S_0)} \\
 &= e^{iu(r_d - r_f)t} E^{\mathbb{Q}} \left[e^{iu \left(L_{T_t}^R - \xi^R T_t^R \right) + iu \left(L_{T_t}^L - \xi^L T_t^L \right)} \right] \\
 &= e^{iu(r_d - r_f)t} E^{\mathbb{M}} \left[e^{-\psi^\top T_t} \right] \equiv e^{iu(r_d - r_f)t} \mathcal{L}_T^{\mathbb{M}}(\psi),
 \end{aligned}$$

- The new measure \mathbb{M} is absolutely continuous with respect to the risk-neutral measure \mathbb{Q} and is defined by a complex-valued exponential martingale,

$$\left. \frac{d\mathbb{M}}{d\mathbb{Q}} \right|_t \equiv \exp \left[iu \left(L_{T_t}^R - \xi^R T_t^R \right) + iu \left(L_{T_t}^L - \xi^L T_t^L \right) + \psi^R T_t^R + \psi^L T_t^L \right].$$

- Girsanov's Theorem yields the (complex) dynamics of the relevant processes under the complex-valued measure \mathbb{M} .

The Laplace Transform of the Stochastic Clocks

- We chose our two new clocks to be continuous over time:

$$T_t^j = \int_0^t v_s^j ds, \quad j = R, L,$$

where for tractability, the activity rates v^j are mean-reverting square root processes with unit long run mean.

- As a result, the Laplace transforms are exponential affine:

$$\mathcal{L}_T^{\mathbb{M}}(\psi) = \exp \left(-b^R(t)v_0^R - c^R(t) - b^L(t)v_0^L - c^L(t) \right),$$

where:

$$b^j(t) = \frac{2\psi^j(1-e^{-\eta^j t})}{2\eta^j - (\eta^j - \kappa^j)(1-e^{-\eta^j t})},$$

$$c^j(t) = \frac{\kappa}{\sigma_v^2} \left[2 \ln \left(1 - \frac{\eta^j - \kappa^j}{2\eta^j} \left(1 - e^{-\eta^j t} \right) \right) + (\eta^j - \kappa^j)t \right],$$

and:

$$\eta^j = \sqrt{(\kappa^j)^2 + 2\sigma_v^2\psi^j}, \quad \kappa^j = \kappa - iu\rho^j\sigma\sigma_v, \quad j = R, L.$$

- Hence, the CFs of the currency return are known in closed form for our models.

Estimation

- We estimated 6 models: HSTSV, MJDSV, KJSSM, VGSSM, CJSSM, CGSSM.
- We used quasi-maximum likelihood with unscented Kalman filtering (UKF).
 - State propagation equation: The time series dynamics for the 2 activity rates

$$dv_t = \kappa^P(\theta^P - v_t)dt + \sigma_v\sqrt{v_t}dZ, \quad (2 \times 1)$$

- Measurement equations: $y_t = O(v_t; \Theta) + e_t, \quad (40 \times 1)$
- y — Use out-of-the-money option prices scaled by the BS vega of the option.
- UKF generates efficient forecasts and updates on states, measurements, and covariances.
- We maximize the following likelihood function to obtain parameter estimates:

$$\begin{aligned} \mathcal{L}(\Theta) &= \sum_{t=1}^N l_{t+1}(\Theta) \\ &= -\frac{1}{2} \sum_{t=1}^N \left[\log |\bar{A}_t| + \left((y_{t+1} - \bar{y}_{t+1})^\top (\bar{A}_{t+1})^{-1} (y_{t+1} - \bar{y}_{t+1}) \right) \right]. \end{aligned}$$

Model Performance Comparison

	HSTSV	MJDSV	KJSSM	VGSSM	CJSSM	CGSSM
JPY: rmse	1.014	0.984	0.822	0.822	0.822	0.820
$\mathcal{L}, \times 10^3$	-9.430	-9.021	-6.416	-6.402	-6.384	-6.336
GBP: rmse	0.445	0.424	0.376	0.376	0.376	0.378
$\mathcal{L}, \times 10^3$	4.356	4.960	6.501	6.502	6.497	6.521

- SSM models with different jump structures perform similarly.
- Under CG, $\alpha = 1.6$ for JPYUSD and 1.18 for GBPUSD, which is weak evidence in favor of **infinite-variation** of the sample paths of the jump components.
- All SSM models perform much better than MJDSV, which performs better than HSTSV.

Likelihood Ratio Tests

Under the null $H_0 : E[l_i - l_j] = 0$, the statistic (\mathcal{M}) is asymptotically $N(0, 1)$.

Curr	\mathcal{M}	HSTSV	MJDSV	KJSSM	VGSSM	CJSSM	CGSSM
JPY	HSTSV	0.00	-2.55	-4.92	-4.88	-4.75	-4.67
JPY	MJDSV	2.55	0.00	-5.39	-5.33	-5.22	-5.07
JPY	KJSSM	4.92	5.39	0.00	-1.11	-0.86	-1.20
JPY	VGSSM	4.88	5.33	1.11	0.00	-0.72	-1.21
JPY	CJSSM	4.75	5.22	0.86	0.72	0.00	-1.59
JPY	CGSSM	4.67	5.07	1.20	1.21	1.59	0.00
GBP	HSTSV	0.00	-2.64	-4.70	-4.68	-4.63	-4.71
GBP	MJDSV	2.64	0.00	-3.85	-3.86	-3.89	-4.19
GBP	KJSSM	4.70	3.85	0.00	-0.04	0.34	-0.37
GBP	VGSSM	4.68	3.86	0.04	0.00	0.56	-0.39
GBP	CJSSM	4.63	3.89	-0.34	-0.56	0.00	-0.51
GBP	CGSSM	4.71	4.19	0.37	0.39	0.51	0.00

Out of Sample Performance Comparison

HSTSV MJDSV KJSSM VGSSM CJSSM CGSSM HSTSV MJDSV KJSSM VGSSM CJSSM CGSSM

JPYUSD

GBPUSD

In-Sample Performance: 1996-2001

rmse	1.04	1.02	0.85	0.85	0.85	0.85	0.47	0.44	0.41	0.41	0.41	0.41
\mathcal{L}/N	-23.69	-23.03	-16.61	-16.60	-16.57	-16.47	8.36	10.06	12.27	12.27	12.26	12.28
\mathcal{M}												
HSTSV	0.00	-2.14	-4.44	-4.41	-4.33	-4.17	0.00	-3.34	-4.42	-4.39	-4.24	-4.33
MJDSV	2.14	0.00	-4.74	-4.70	-4.61	-4.42	3.34	0.00	-3.40	-3.39	-3.33	-3.33
KJSSM	4.44	4.74	0.00	-0.49	-0.51	-0.84	4.42	3.40	0.00	0.08	0.36	-0.42
VGSSM	4.41	4.70	0.49	0.00	-0.51	-0.89	4.39	3.39	-0.08	0.00	0.51	-0.42
CJSSM	4.33	4.61	0.51	0.51	0.00	-1.14	4.24	3.33	-0.36	-0.51	0.00	-0.55
CGSSM	4.17	4.42	0.84	0.89	1.14	0.00	4.33	3.33	0.42	0.42	0.55	0.00

Similar to the whole-sample performance

Out of Sample Performance Comparison

HSTSV MJDSV KJSSM VGSSM CJSSM CGSSM HSTSV MJDSV KJSSM VGSSM CJSSM CGSSM

JPYUSD

GBPUSD

Out-of-Sample Performance: 2001-2002


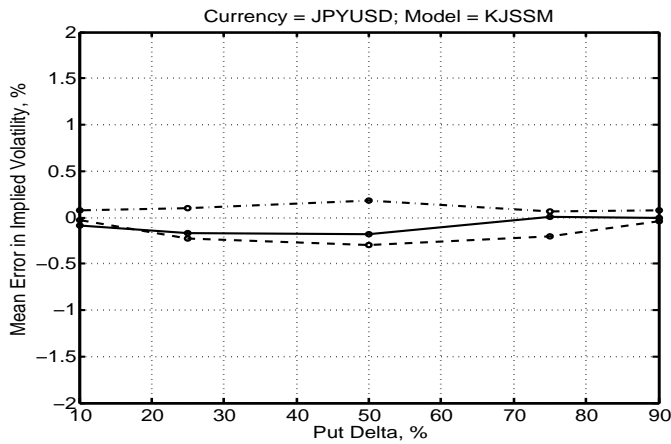
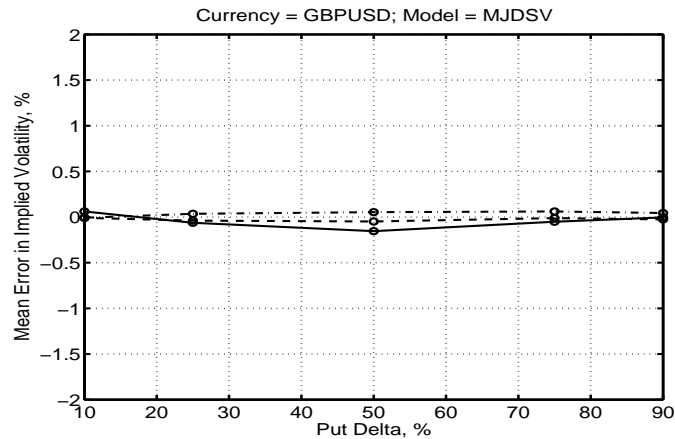
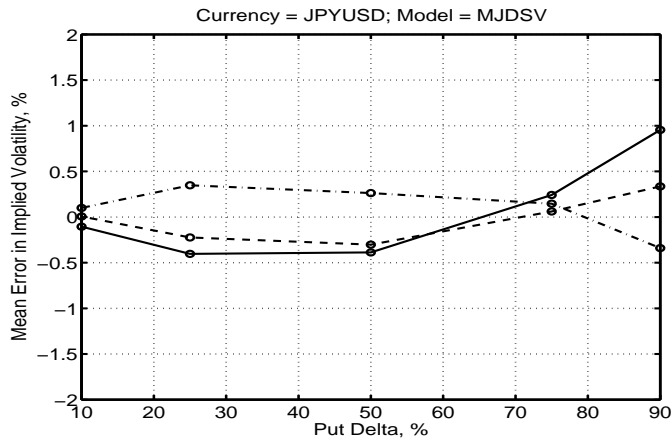
rmse	1.06	1.00	0.90	0.90	0.89	0.89	0.39	0.37	0.27	0.27	0.27	0.27
\mathcal{L}/N	-24.01	-21.75	-18.47	-18.35	-18.23	-18.11	14.36	15.85	23.30	23.29	23.26	23.25
\mathcal{M}												
HSTSV	0.00	-6.01	-5.90	-6.01	-6.08	-6.12	0.00	-4.88	-7.06	-7.06	-7.05	-7.05
MJDSV	6.01	0.00	-3.11	-3.23	-3.32	-3.48	4.88	0.00	-5.98	-5.99	-5.99	-5.97
KJSSM	5.90	3.11	0.00	-7.76	-6.81	-5.27	7.06	5.98	0.00	0.64	1.47	4.51
VGSSM	6.01	3.23	7.76	0.00	-4.39	-3.67	7.06	5.99	-0.64	0.00	1.63	4.19
CJSSM	6.08	3.32	6.81	4.39	0.00	-3.11	7.05	5.99	-1.47	-1.63	0.00	0.23
CGSSM	6.12	3.48	5.27	3.67	3.11	0.00	7.05	5.97	-4.51	-4.19	-0.23	0.00

Similar to the in-sample performance





Infinite variation jumps are more suitable to capture large smiles/skews.

Mean Pricing Error


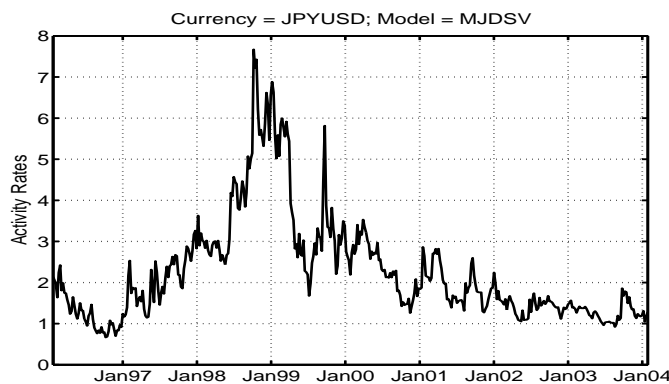
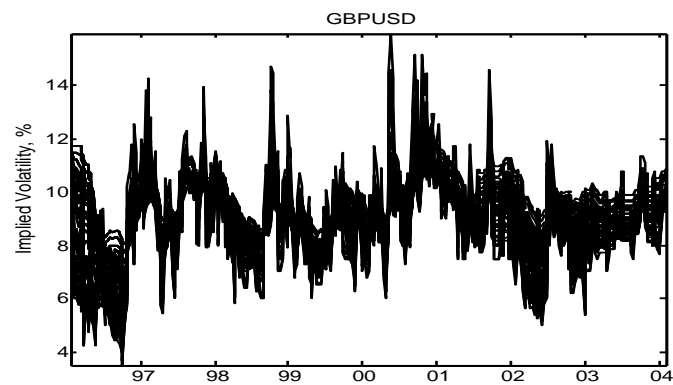
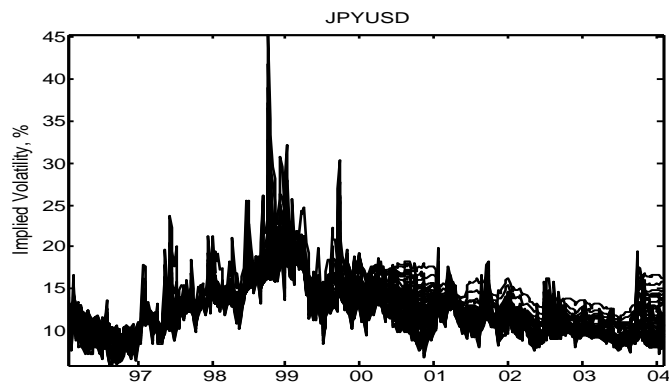
1 Month = Solid Line; 3 Month = Dashed Line; 12 Month = Dash Dotted Line



Mean Absolute Pricing Error







Bates Model Captures Stochastic Volatility





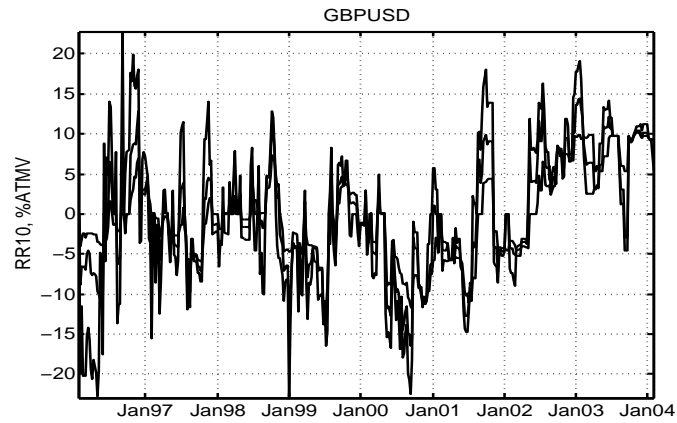

Demand for calls on yen drives up the yen volatility during the hedge fund crisis.

SSM Captures Stochastic Volatility, and ...



The demand for **yen calls** only drives up the activity rate (volatility) of **upward yen moves** (solid line), but not the volatility of downward yen moves.

SSM Also Captures Stochastic Skew



Activity Rate Dynamics: JPYUSD

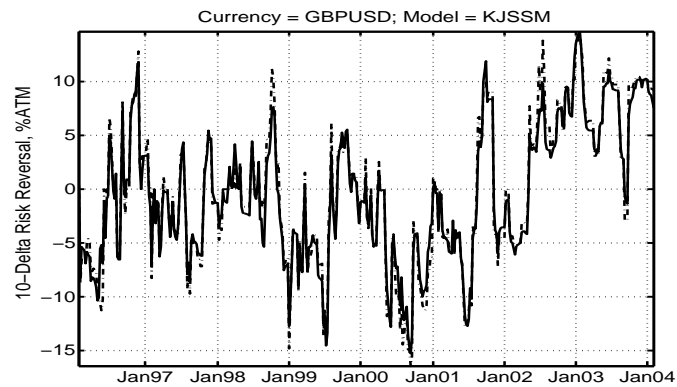



Θ_B	Θ_S	HSTSV	MJDSV	KJSSM	VGSSM	CJSSM	CGSSM
κ	κ	0.559 (0.006)	0.569 (0.011)	0.387 (0.005)	0.394 (0.006)	0.421 (0.007)	0.465 (0.010)
σ_v	σ_v	1.837 (0.023)	1.210 (0.022)	1.675 (0.027)	1.657 (0.028)	1.582 (0.027)	1.566 (0.031)
ρ	ρ^R	0.076 (0.005)	0.123 (0.065)	0.395 (0.017)	0.393 (0.018)	0.400 (0.022)	0.424 (0.056)
μ_j	ρ^L	— (—)	-0.210 (0.024)	-0.739 (0.034)	-0.758 (0.036)	-0.851 (0.040)	-1.000 (0.144)
κ^P	κ^P	0.745 (0.396)	0.258 (0.114)	0.522 (0.289)	0.502 (0.288)	0.544 (0.251)	0.586 (0.261)
—	α	— —	— —	— —	— —	— —	1.602 (0.126)

Activity Rate Dynamics: GBPUSD

Θ_B	Θ_S	HSTSV	MJDSV	KJSSM	VGSSM	CJSSM	CGSSM
κ	κ	1.532 (0.007)	1.044 (0.007)	1.205 (0.006)	1.206 (0.006)	1.211 (0.006)	1.180 (0.008)
σ_v	σ_v	2.198 (0.026)	1.737 (0.023)	1.429 (0.039)	1.447 (0.040)	1.505 (0.017)	1.492 (0.018)
ρ	ρ^R	-0.023 (0.003)	-0.061 (0.017)	0.848 (0.040)	0.848 (0.043)	0.849 (0.017)	0.836 (0.016)
μ_j	ρ^L	— (—)	0.002 (0.001)	-1.000 (0.047)	-0.999 (0.050)	-1.000 (0.000)	-1.000 (0.004)
κ^P	κ^P	1.276 (0.345)	0.800 (0.236)	2.062 (0.213)	2.092 (0.213)	1.158 (0.006)	3.296 (0.223)
—	α	— —	— —	— —	— —	— —	1.180 (0.155)

Theory and Evidence on Stochastic Skew

Three-month ten-delta risk reversals: data (dashed lines), model (solid lines).



Conclusions

- Using currency option quotes, we find that under a risk-neutral measure, currency returns display not only stochastic volatility, but also stochastic skew.
- State-of-the-art option pricing models (e.g. Bates 1996) capture stochastic volatility and static skew, but not stochastic skew.
- Using the general framework of time-changed Lévy processes, we propose a class of models (SSM) that captures both stochastic volatility *and* skewness.
- The models we propose are also highly tractable for pricing and estimation. The pricing speed is comparable to the speed of the Bates model.