

Essays on Foreign Exchange Option Markets

D I S S E R T A T I O N
of the University of St. Gallen,
School of Management,
Economics, Law, Social Sciences
and International Affairs
to obtain the title of
Doctor of Philosophy in Management

submitted by

Ralf Büsser

from

Amden (St. Gallen)

Approved on the application of

Prof. Dr. Manuel Ammann

and

Prof. Paul Söderlind, PhD

Dissertation no. 4002

Difo-Druck GmbH, Bamberg, 2012

The University of St. Gallen, School of Management, Economics, Law, Social Sciences and International Affairs hereby consents to the printing of the present dissertation, without hereby expressing any opinion on the views herein expressed.

St. Gallen, May 11, 2012

The President:

Prof. Dr. Thomas Bieger

*To Annina and my parents,
Ursula & Theo Büsser*

Acknowledgement

I am highly indebted to a number of people who have encouraged, inspired and supported me prior to and throughout my time as a doctoral student. Without their help, I would not have accomplished the successful completion of the present dissertation.

First, I wish to thank my supervisor Prof. Manuel Ammann for his advice and support, and for giving me the opportunity to pursue doctoral studies despite my having already spent an extended period in the professional world. I am grateful to my co-supervisor Prof. Paul Söderlind for many helpful comments and detailed suggestions which have greatly improved the individual papers. The dissertation has benefited from lively discussions with many colleagues at the Swiss Institute of Banking and Finance and Algotin AG. A number of individuals at UBS Investment Bank deserve mentioning. I like to thank Tim Weithers for having spurred my interest in the mathematics of financial options. Benjamin Anderegg has been an excellent mentor and teacher on many practical aspects of foreign exchange options trading during my early days at UBS. Finally, I am indebted to Hans-Ueli Berger for giving me the chance to work on the foreign exchange exotics trading desk and for providing me with the data which forms the backbone of this thesis. Undoubtedly, the dissertation in its current form would not have been possible without his support.

I dedicate this thesis to my parents, Ursula and Theo Büsser, and my beloved Annina. Throughout my life, the former have always encouraged me to remain engaged in education and have provided me with all their help and support. Lastly, Annina deserves my utmost gratitude for all her love, unconditional support and understanding throughout these sometimes difficult times.

Zurich, June 2012

Ralf Büsser

Contents

I	Biases in Foreign Exchange Density Forecasts	1
1	Introduction	2
2	Forecasting Methodology	4
2.1	The Malz Approach	5
2.2	The Vanna-Volga Method	5
2.3	The Kalman Filter Extension of the Vanna-Volga Method	6
3	Testing the Unconditional Forecasting Density	9
3.1	QQ Plots	9
3.2	A Formal Test of Currency Density Forecasts	12
3.3	Robustness Check Using Deep OTM Options	15
4	Forecasting Performance and Risk Factors	18
4.1	Variance Risk Factor	18
4.2	Jump Risk Factor	21
5	Conclusion	24
6	Appendix	25
A	Quotation Convention Adapted Version of the Malz Interpolation Method	25
B	The Vanna-Volga Method	26
C	Maximum Likelihood Estimation of the Kalman Filter	27
II	Variance Risk Premiums in Foreign Exchange Markets	31
1	Introduction	32
2	Static Hedging and Model-Free Variance Forecasting	34
3	Data and Methodology	37
4	Time Series Dynamics of Variance Risk Premiums	40
4.1	Evidence on the Sign and Size of Variance Risk Premiums	40
4.2	Variance Risk Premiums from Intraday Data	43
4.3	Realized Variance at Alternative Frequencies	45

4.4	Are Variance Risk Premiums Time-varying?	48
5	How Does the Financial Crisis of 2008 Affect the Results?	50
6	What Drives Variance Risk Premiums in Foreign Exchange Markets?	53
6.1	S&P 500, the VIX and the TED Spread	53
6.2	Do Variance Risk Premiums Subsume Fear of Jump Risk?	54
7	Conclusion	56
8	Appendix	58
A	Derivation of Risk-Neutral Variance Forecasts	58
B	Generalization of the Malz Interpolation Method	60
III	The Pricing of Foreign Exchange One-Touch Options	64
1	Introduction	65
2	The Risk-Neutral Dynamics of Currency Pairs	67
2.1	Pure Diffusion Model	67
2.2	Pure Jump and Jump-Diffusion Model	68
2.3	Stochastic Volatility	70
3	Calibrating the Models to the Vanilla Market	72
3.1	Data	72
3.2	Calibration	73
3.3	In-Sample Evidence on the Structure of Currency Returns	75
4	Evidence from One-Touch Options	77
4.1	Data and Biases	78
4.2	Calibration, Pricing and Numerical Issues	81
4.3	Results	83
5	Conclusion	88
6	Appendix	90
A	Measure Change and Laplace Transform	90
B	Option Pricing via Fourier-cosine Expansion	91

C	One-Touch Option Prices when Currency Pairs Follow a Geometric Brownian Motion	92
D	Simulating Jumps	93

List of Figures

I	Biases in Foreign Exchange Density Forecasts	1
1	QQ Plots for the 1-month Horizon	10
2	QQ Plots for the 3-month Horizon	12
3	Alternative QQ Plots for the 1-month Horizon	16
4	Alternative QQ Plots for the 3-month Horizon	17
II	Variance Risk Premiums in Foreign Exchange Markets	31
1	Implied Volatility Functions	39
2	Variance Risk Premiums at Alternative Frequencies	47
3	Time-Variation in Variance Risk Premiums	49
4	Variance Risk Premiums During the Financial Crisis	51
III	The Pricing of Foreign Exchange One-Touch Options	64
1	Model Overview	72
2	Exemplary Implied Volatility Functions	73
3	One-Touch Option Data	79
4	Black-Scholes Model Biases in One-Touch Options	80
5	EURUSD Pricing Errors	83
6	USDJPY Pricing Errors	84

List of Tables

I	Biases in Foreign Exchange Density Forecasts	1
1	Tests of the Quality of 1-month Density Forecasts	14
2	Tests of the Quality of 3-month Density Forecasts	15
3	Wald Tests for the Alternative Density Forecasts	18
4	Variance Risk Factor and Forecasting Quality at the 1-month Horizon .	20
5	Variance Risk Factor and Forecasting Quality at the 3-month Horizon .	21
6	Jump Risk Factors and Forecasting Quality at the 1-month Horizon . .	22
7	Jump Risk Factors and Forecasting Quality at the 3-month Horizon . .	23
II	Variance Risk Premiums in Foreign Exchange Markets	31
1	Average Variance Risk Premiums from Daily Data	41
2	Average Log Variance Risk Premiums from Daily Data	42
3	Average Variance Risk Premiums from Intraday Data	43
4	Average Log Variance Risk Premiums from Intraday Data	44
5	Average Realized Volatility at Different Sampling Frequencies	46
6	Expectation Hypothesis Regression	50
7	Variance Risk Premiums for the Full Sample	52
8	Classic Risk Factor Regression	54
9	Jump Risk Factor Regression	55
III	The Pricing of Foreign Exchange One-Touch Options	64
1	EURUSD Parameter Estimates and Model Performance	76
2	USDJPY Parameter Estimates and Model Performance	77
3	Pricing Errors by Moneyness Bucket	86
4	Performance Comparison	87

Zusammenfassung

Die vorliegende Doktorarbeit gliedert sich in drei Teile, welche als eigenständige Forschungsarbeiten konzipiert sind. Gemeinsam ist den drei Teilen, dass sie unterschiedliche Aspekte von Devisenoptionsmärkten beleuchten. Im ersten Teil wird die Genauigkeit von risiko-neutralen Dichtevorhersagen untersucht. Zur Konstruktion von Dichten werden eine Reihe von Interpolationstechniken angewendet, darunter eine von den Autoren vorgeschlagene dynamische Methode. Wir stellen fest, dass zwischen der risiko-neutralen und der physischen Dichte eines Währungspaars signifikante Diskrepanzen bestehen. Diese treten unabhängig von der gewählten Informationsmenge auf. Des Weiteren weisen wir empirisch einen Zusammenhang zwischen der Genauigkeit von risiko-neutralen Dichtevorhersagen und geeigneten Massen für Varianz- und Kurs-sprungrisiko nach.

Der zweite Teil untersucht Varianzrisikoprämien in Devisenmärkten anhand eines modell-unabhängigen Ansatzes. Zur Berechnung der realisierten Varianz werden Daten unterschiedlicher Frequenz herangezogen. Bei tiefer Frequenz können signifikant negative Varianzrisikoprämien nachgewiesen werden. Im Gegensatz dazu gilt dies nicht für den Gebrauch von Hochfrequenzdaten, was mutmasslich von Marktmikrostruktur-Effekten herrührt. Weitere Untersuchungen zeigen, dass Varianzrisikoprämien nur bedingt auf klassische Risikofaktoren oder Angst vor Kurssprüngen zurückgeführt werden können. Hingegen kann ein signifikanter Zusammenhang zwischen dem Erfolg einer Varianzswapstrategie und dem VIX, dem TED Spread und der impliziten Volatilität im jeweiligen Devisenpaar nachgewiesen werden. Insgesamt kommen wir zum Schluss, dass in Devisenmärkten die Varianz einen eigenständigen Risikofaktor mit zeitabhängiger Prämie darstellt.

Im dritten Teil wird die Struktur von risiko-neutralen Devisenrenditen anhand von zwei Datensätzen, nämlich den Preisen für Vanilla und One-touch Optionen, untersucht. Verschiedene Modelle werden am Vanilla Markt kalibriert und anschliessend auf den One-touch Datensatz angewendet. Dieses Vorgehen erlaubt eine Aussage hinsichtlich der Kohärenz zwischen den beiden Märkten, ermöglicht aber auch fundierte Schlussfolgerungen über die Dynamik von Devisenrenditen. Die Untersuchungen zeigen, dass ein komplexes Modell mit stochastischer Volatilität und Kurssprüngen den Vanilla Optionsmarkt am besten abzubilden vermag. Im Gegensatz dazu erzielt ein vergleichsweise einfaches stochastisches Volatilitätsmodell das beste Ergebnis am Markt für One-touch Optionen. Mutmasslich lässt sich dieses Resultat darauf zurückführen, dass sich Optionshändler bei der Preisstellung am Black-Scholes Modell orientieren.

Summary

This dissertation comprises three parts. Each part has been set up as a self-contained research project and deals with a specific aspect of foreign exchange option markets. Part I examines the accuracy of option-implied density forecasts in predicting future realizations of the spot rate. To produce density forecasts, a range of interpolation techniques is used, among them a novel method that dynamically updates the information content of currency options. We observe that the risk-neutral density generally provides a biased estimate of the physical return distribution. This finding is robust to different choices of information sets. In Part I, we further establish empirically a link between forecasting accuracy and surrogates for variance and jump risk.

Part II examines the variance risk premiums in foreign exchange markets using a model-free approach. When realized variance is computed from intraday data with low frequency, variance risk premiums are significantly negative. In contrast, estimates based on high-frequency data provide a somewhat different picture. This latter finding is likely owed to microstructure effects. Further investigations suggest that variance risk premiums are essentially unexplained by classic risk factors or fear of jump risk. However, we find a significant relationship between the success of a variance swap strategy and the VIX, the TED spread and the shape of the implied volatility function. Overall, we conclude that foreign exchange markets feature a separately priced variance risk factor with time-varying risk premium.

In Part III, the structure of currency returns is examined at the case of two unspanned information sets, namely the prices for vanilla and one-touch options. A variety of models with increasing complexity is calibrated to the vanilla market and subsequently applied to quotes for one-touch options. This approach allows to draw conclusions on the coherence between the two markets and ultimately to infer risk-neutral currency dynamics from a richer data set. Evidence suggests that vanilla options are most accurately priced using a model that comprises both stochastic volatility and jumps with time-varying intensity. In contrast, one-touches imply little jump risk. A comparatively simple stochastic volatility model attains the best performance. This finding may be a result of market makers' use of Black-Scholes prices as a reference point.

Part I

Biases in Foreign Exchange Density Forecasts

Abstract

In this article, we examine biases in foreign exchange risk-neutral density forecasts. To determine the biases, we propose a novel method that accounts for the full information content in the history and cross-section of option prices. Unlike in previous studies, our approach is consistent with the specific quotation convention of currency markets. To investigate whether risk aversion may explain the biases in density forecasts, we relate them to suitable surrogates for variance and jump risk.

Our results suggest that risk-neutral density forecasts are biased estimates of the statistical return distribution. This finding is robust to alternative information sets comprised of deep OTM options, and holds for a short and longer forecasting horizon. Moreover, we find that both the variance and jump risk factors have a significant impact on the forecasting accuracy.

1 Introduction

Option-implied density forecasts are popular in policy making and risk management because they provide forward-looking information. However, there is considerable doubt with regard to their accuracy.¹ In this paper, we analyze the biases in density forecasts implied in currency options. By bias, we refer to the difference between the risk-neutral and statistical distribution.² We contribute to the literature in several ways: First, we measure biases in density forecasts using a novel method that dynamically updates the information content of currency options. Second, we combine options with different moneyness to assess whether they alter the biases in the tails of the predicted density. Third, we examine whether forecasting accuracy is related to commonly priced risk factors in currency markets.

Several methods to construct option-implied density forecasts have been proposed. For example, Jackwerth and Rubinstein (1996) use a minimization criterion to adjust a prior log-normal distribution to account for the information contained in option prices. Bates (1996) and Melick and Thomas (1997) assume a log-normal mixture density and then use option prices to estimate the unknown parameters. Ait-Sahalia and Lo (1998) and Ait-Sahalia and Duarte (2003) propose non-parametric estimation techniques. All of these methods have their shortcomings when applied to currency markets: First, it is preferred not to have a prior regarding the currency return distributions. Second, the methods need to build a distribution function on the basis of a few quotes only. Third, density forecasting requires extrapolation beyond the available strike domain.

A comprehensive analysis on the accuracy of density forecasts from foreign exchange options has been conducted by Christoffersen and Mazzotta (2005). They report significant differences between the statistical and risk-neutral distribution for currencies. To generate the forecasts, Christoffersen and Mazzotta resort to an interpolation method proposed by Malz (1997). However, as shown in Reiswich and Wystup (2010), forecasts from the Malz approach are generally flawed since the method is inconsistent with the specific quotation convention used in currency markets. Also, to the best of our knowledge, there has been no study to date that empirically investigates the determinants of the biases in foreign exchange density forecasts.

In this article, we address these issues. To ensure that identified biases are not subject to a particular choice of methodology, we use three alternative approaches: First,

¹See Ait-Sahalia and Lo (2000), Jackwerth (2000) and Christoffersen and Mazzotta (2005) to name a few.

²As a matter of definition, we may synonymously speak of the quality or accuracy of risk-neutral density forecasts.

we propose an adoption of the Malz approach consistent with the quotation conventions. Second, we consider density forecasts from the vanna-volga method, which is introduced in Castagna and Mercurio (2007). Both the Malz and vanna-volga method rely on three option quotes only and are thus well-suited to cope with the scarce information in the strike domain of currency option markets. We apply the approaches to different combinations of option quotes, including far out-of-the money (OTM) options. OTM options potentially provide different information on the tails of the distribution, where Christoffersen and Mazzotta (2005) report most of the biases. Finally, we propose a Kalman filter extension of the vanna-volga method (henceforth KFVV) that dynamically updates the information content of currency options. Moreover, the KFVV extends the information set to all options in the strike domain. In all cases, we aim to examine whether our findings are robust to altering the information set. The KFVV may be viewed as the foreign exchange equivalent to the technique presented in Bedendo and Hodges (2009). They use the Kalman filter to produce one-day volatility forecasts for S&P 500 quarterly future options. Although not the purpose of this paper, such a task is straightforward to perform using our approach.

Ait-Sahalia and Lo (2000) and Jackwerth (2000) suggest that possible differences between the risk-neutral and statistical densities arise as a result of market participants' risk aversion. We take up this hypothesis and investigate whether there is a systematic relationship between the accuracy of density forecasts and suitable proxies for variance and jump risk. Our conjecture is the following: If the variance and jump risk factors are priced, the premiums are observable in option prices and hence in the option-implied density while the statistical distribution is unaffected. As a result, risk factor realizations determine at least to some extent the biases in density forecasts. Similar reasoning has led Bakshi and Kapadia (2003) to examine whether delta-hedged gains can be attributed to priced risk factors. Effectively, delta-hedged gains arise because the statistical and risk-neutral density functions differ.

Our results suggest that although the biases are not entirely unrelated to the choice of methodology, similar and persistent biases are observed for all three interpolation techniques. The biases tend to aggravate at a longer forecasting horizon, and they persist no matter what options are chosen as the underlying information set. Hence, deep OTM options do not seem to provide different information on the tails of the statistical distribution. Finally, the conjecture that the variance risk factor affects several aspects of the quality of risk-neutral density forecasts is confirmed. This result also applies in the presence of jump risk proxies, which themselves have a significant impact on the accuracy of density forecasts.

The remainder of this paper is organized as follows. In section 2, we present the forecasting methodology used throughout this paper. Furthermore, we briefly present the Malz and vanna-volga method and introduce the Kalman filter approach. In section 3, we present results on the forecasting biases under the various methods. This section also includes robustness checks that combine various option quotes into different information sets. In section 4, we establish a test environment to investigate the relationship between the biases in density forecasts and commonly priced risk factors. Section 5 concludes.

2 Forecasting Methodology

Throughout this paper, we follow the common quotation convention and measure the price of a foreign currency (FOR) in domestic (DOM) units. We investigate the biases in risk-neutral density forecasts for EURUSD, GBPUSD, USDJPY and EURGBP, quoted as FORDOM. In foreign exchange, traders quote volatilities in delta space. From UBS, a major investment bank and market maker in foreign exchange, we obtain quotes on delta-neutral straddles (DN) and 5-delta, 10-delta and 25-delta call (DC) and put (DP) options. When these volatility quotes are mapped to their respective delta, special attention must be paid to the premium convention, since the delta depends on it. For EURUSD and GBPUSD, the domestic currency, i.e. the USD, is the premium currency. Therefore, a so-called regular delta convention applies. In contrast, USDJPY and EURGBP are quoted with a foreign currency premium. This implies that the spot delta must be adjusted by an amount that reflects the premium paid (premium-adjusted delta convention). For details, we refer to Reisch and Wystup (2010).

To construct option-implied density forecasts, we rely on the seminal Breeden and Litzenberger (1978) result. Given a complete stochastic basis $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{Q})$ with risk-neutral measure \mathbb{Q} , the price $c(X, \tau)$ of a call option on the foreign currency is expressed as

$$c(X, \tau) = E^{\mathbb{Q}}[e^{-r_d \tau} c(X, 0) \mid \mathcal{F}_t], \quad (1)$$

where X , $\tau = (T - t)$ and r_d denote the exercise price, the time to maturity and the domestic interest rate, respectively. Twice differentiating (1) with respect to the exercise price yields an expression for a currency pair's probability density function

$$q(X) = e^{r_d \tau} \frac{\partial^2 c(X, \tau)}{\partial X^2}. \quad (2)$$

In practice, we choose a small interval h and approximate the probability density function at X with its central difference estimator,³

$$\hat{q}(X) = e^{r_{a\tau}} \left[\frac{c(X-h, \tau) - 2c(X, \tau) + c(X+h, \tau)}{h^2} \right] h. \quad (3)$$

To obtain an estimate of a currency pair's density function, we require call option prices $c(\mathbf{X}, \tau)$ for a wide range of exercise prices. For this reason, we resort to one of the interpolation methods introduced in what follows.

2.1 The Malz Approach

Malz (1997) proposes an interpolation of implied volatilities in delta space. To fix matters, define a risk reversal and butterfly for a call and put option with delta x as

$$\sigma^{xrr} := \sigma^{xDC} - \sigma^{xDP}, \quad \sigma^{xbf} := \frac{\sigma^{xDC} + \sigma^{xDP}}{2} - \sigma^{DN}. \quad (4)$$

To generate an implied volatility function, Malz specifies a quadratic polynomial of the form

$$\hat{\sigma}_t^\Delta = a_0 \sigma_t^{DN} + a_1 \sigma_t^{xrr} (\Delta - 0.50) + a_2 \sigma_t^{xbf} (\Delta - 0.50)^2, \quad (5)$$

where Δ refers to the delta of the required option and a_0, a_1 and a_2 are a set of parameters that can be solved for by requiring (5) to hold for $\sigma^{DN}, \sigma^{xrr}$ and σ^{xbf} . The Malz approach refers to a forward delta under the regular quotation convention, since only then the delta-neutral quote has a delta of 0.50. In all other instances, the Malz approach produces market inconsistent forecasts. To cope with this problem, we replace 0.50 with the appropriate Δ_{DN} and solve for a different set of parameters b_0, b_1, b_2 . To arrive at the required call prices, we first move from delta-volatility to exercise price-volatility space. Next, we obtain call prices through the Garman and Kohlhagen (1983) formula for currency options. Details have been relegated to Appendix A.

2.2 The Vanna-Volga Method

The vanna-volga method presented in Castagna and Mercurio (2007) and Shkolnikov (2009) interpolates between market prices by representing options as replicating portfolios. When volatility is stochastic, the Itô expansion of an option price is given by

$$dc(X, \tau, S_t, \sigma_t) = \left[\frac{\partial c(\cdot)}{\partial t} + \frac{1}{2} \sigma_t^2 S_t^2 \frac{\partial^2 c(\cdot)}{\partial S_t^2} \right] dt + \frac{\partial c(\cdot)}{\partial S_t} dS_t + \frac{\partial c(\cdot)}{\partial \sigma_t} d\sigma_t + \frac{\partial^2 c(\cdot)}{\partial S_t \partial \sigma_t} dS_t d\sigma_t + \frac{1}{2} \frac{\partial^2 c(\cdot)}{\partial \sigma_t^2} d\sigma_t^2, \quad (6)$$

³We multiply the difference estimator by h to make $\hat{q}(\cdot)$ an Arrow-Debreu security (See Arrow (1964) and Debreu (1959)).

where $\partial c(\cdot)/\partial\sigma_t$, $\partial^2 c(\cdot)/\partial S_t \partial\sigma_t$ and $\partial^2 c(\cdot)/\partial\sigma_t^2$ refer to the vega, vanna and volga of an option.⁴ Hedging an option with second-order accuracy requires that the stochastics in $d\sigma_t$, $dS_t d\sigma_t$ and $d\sigma_t^2$ are taken care of. Since there are three volatility greeks, an accurate replication is achieved by holding a portfolio of three options, on top of the spot and money market positions inherited from a Black and Scholes (1973) setting.

To establish a formal relationship, suppose we observe a portfolio of three pivot options. Presume that the corresponding volatility greeks are concatenated in the 3×3 matrix A .⁵ For an arbitrary option with strike price X , stack its volatility greeks into a 3×1 vector \mathbf{y} . To offset the risk arising from stochastic volatility, we must choose portfolio weights \mathbf{x} according to

$$\mathbf{y} = \mathbf{A}\mathbf{x}, \quad \rightarrow \quad \mathbf{x} = \mathbf{A}^{-1}\mathbf{y}. \quad (7)$$

To arrive at the required price $c(X, \tau, S_t, \sigma_t^X)$, define a 3×1 vector \mathbf{c}^{me} of market excess prices, i.e. a vector that subsumes the market prices of the options in the hedging portfolio in excess of their Black-Scholes values. By the usual replication arguments, it must hold that an option price reflects the hedging costs associated with spot diffusion risk and the extra costs associated with stochastic volatility. Since the former are given by the Black-Scholes value, we formally obtain the dichotomized relationship

$$c(X, \tau, S_t, \sigma_t^X) = c(X, \tau, S_t, \sigma_t^{BS}) + (\mathbf{c}^{\text{me}})' \mathbf{x}. \quad (8)$$

Equation (8) is directly employed in (3). To effectively construct density forecasts, we follow industry practice and set $\sigma^{BS} = \sigma^{DN}$.

2.3 The Kalman Filter Extension of the Vanna-Volga Method

Density forecasts from the Malz and vanna-volga method rely on three option quotes only, whereas the data set comprises 7 quotes in the strike domain. Moreover, the Malz and vanna-volga method disregard potentially useful historical information. To address these concerns, we propose a Kalman filter extension of the vanna-volga method.⁶ To establish the methodology, we rewrite (8) as

$$c(X, \tau, S_t, \sigma_t^X) - c(X, \tau, S_t, \sigma_t^{BS}) = \mathbf{y}'_t \alpha_t, \quad (9)$$

⁴We henceforth refer to their union as the 'volatility greeks'.

⁵We refer to Appendix B for a more detailed exposition.

⁶To draw an analogy to the time series estimation of volatility, consider a GARCH model (Bollerslev, 1986) which relies on historical information to obtain an update of current and future volatility. Similarly, the KFVV takes the history of option prices into account.

where $\alpha_t = (\mathbf{A}_t^{-1})' \mathbf{c}_t^{\text{me}}$ is a 3×1 vector which can be interpreted as a vector which entails prices for vega, vanna and volga. The vanna-volga method therefore maintains that for an option with exercise price X , the market price in excess of the Black-Scholes value is a result of the option's vega, vanna and volga multiplied by the respective hedging costs. Importantly, this interpretation is consistent with the way currency option traders perceive second-order risk.

Suppose we stack all observed excess prices into a (7×1) vector $\mathbf{c}_t^{\text{me},X}$. Since (9) must hold for any option, we can estimate the loadings from

$$\mathbf{c}_t^{\text{me},X} = \mathbf{Y}_t' \alpha_t + \varepsilon_t, \quad (10)$$

where Y_t is a 3×7 matrix of volatility greeks and ε_t comprises error terms. To allow for time-varying loadings, we estimate (10) using the Kalman filter. In a state-space representation, the *observation equation* is given by

$$\mathbf{c}_t^{\text{me},X} = \mathbf{Y}_t' \xi_t + \mathbf{Y}_t' \bar{\alpha} + \varepsilon_t. \quad (11)$$

The *state equation* which governs the time-varying component of the greek loadings $\xi_t = \alpha_t - \bar{\alpha}$ is expressed as

$$\xi_t = \mathbf{F} \xi_{t-1} + \omega_t. \quad (12)$$

F is a 3×3 transition matrix and ω_t entails error terms. The instantaneous error covariance matrices are defined as

$$\mathbf{E}(\varepsilon_t \varepsilon_t') =: \mathbf{R}, \quad \mathbf{E}(\omega_t \omega_t') =: \mathbf{Q}, \quad \text{with } \mathbf{E}(\varepsilon_t \omega_t') = \mathbf{0}. \quad (13)$$

Applying the Kalman filter allows us to update the greek loadings by using the information content of all currency options in the cross section. The construction of density forecasts then follows along the lines of the regular vanna-volga method but using (9) instead of (8). This proximity to the vanna-volga method ensures that for typical market parameters, the option-implied density satisfies no-arbitrage conditions.⁷

Estimation of the greek loadings boils down to recursively calculate linear least square forecasts of the state vector

$$\hat{\xi}_{t+1|t} = \hat{\mathbf{E}}(\xi_{t+1} | \mathcal{I}_t), \quad (14)$$

on the basis of observed data $\mathcal{I}_t := (c_t^{\text{me},X}, c_{t-1}^{\text{me},X}, \dots, c_1^{\text{me},X}, Y_t, Y_{t-1}, \dots, Y_1)$ up to time t . Associated with these forecasts is a mean square error matrix

$$\mathbf{P}_{t+1|t} := \mathbf{E}[(\xi_{t+1} - \hat{\xi}_{t+1|t})(\xi_{t+1} - \hat{\xi}_{t+1|t})']. \quad (15)$$

⁷See Castagna and Mercurio (2007) for a note on this.

Suppose we have a current estimate for the time-varying component of greek loadings $\hat{\xi}_{t|t-1}$. With the arrival of new information $\{\mathbf{c}_t^{\text{me},\mathbf{X}}, \mathbf{Y}_t\}$, an update for ξ is obtained through

$$\hat{\xi}_{t|t} = \hat{\xi}_{t|t-1} + \{\mathbf{P}_{t|t-1} \mathbf{Y}_t [\mathbf{Y}_t' \mathbf{P}_{t|t-1} \mathbf{Y}_t + \mathbf{R}]^{-1} [\mathbf{c}_t^{\text{me},\mathbf{X}} - \mathbf{Y}_t' \hat{\xi}_{t|t-1} - \mathbf{Y}_t' \bar{\alpha}]\}. \quad (16)$$

Similarly,

$$\mathbf{P}_{t|t} = \mathbf{P}_{t|t-1} - \{\mathbf{P}_{t|t-1} \mathbf{Y}_t [\mathbf{Y}_t' \mathbf{P}_{t|t-1} \mathbf{Y}_t + \mathbf{R}]^{-1} \mathbf{Y}_t' \mathbf{P}_{t|t-1}\}. \quad (17)$$

The next step estimates are then given by

$$\hat{\xi}_{t+1|t} = \mathbf{F} \hat{\xi}_{t|t}, \quad (18)$$

and

$$\mathbf{P}_{t+1|t} = \mathbf{F} \mathbf{P}_{t|t} \mathbf{F}' + \mathbf{Q}. \quad (19)$$

To make the KFVV approach work for our purposes, we establish the following algorithm.

- a) We estimate the parameters $\{\bar{\alpha}, \mathbf{F}, \mathbf{Q}, \mathbf{R}\}$ over an initial one-year period.
- b) We propagate the variables ξ and \mathbf{P} in the system of equations (16) to (19), starting with the first observation of the estimation period. We do so until we reach the end of a one-month out-of-sample period that immediately follows the estimation window. Risk-neutral density forecasts are produced along the way for the out-of-sample period only.
- c) By the time the end of the one-month out-of-sample period is reached, we go back to step a) and re-estimate the parameters, adding the latest one-month set of data to the estimation period. The choice of a one-month out-of-sample period is one of convenience. Presumably more accurate forecasts result if the parameters were re-estimated on a daily basis. However, this comes at the cost of additional computing time.

To facilitate estimation, we impose some structure on the parameters \mathbf{F} and \mathbf{R} . Specifically, we assume \mathbf{F} to be a diagonal matrix. Spillover effects from one greek loading to another are therefore precluded. Furthermore, to achieve a better smoothing effect for our forecasts, the elements in \mathbf{F} are restricted between 0 and 1. Concerning \mathbf{R} , we presume that correlations are constant, where the correlations are obtained from a prior-stage OLS regression. Hence, only the diagonal elements in \mathbf{R} need to be estimated. No assumptions are made about $\bar{\alpha}$ and \mathbf{Q} . Initial estimates for $\bar{\alpha}$ are also from an OLS regression. The parameters are estimated by maximum likelihood. A brief description is found in Appendix C.

3 Testing the Unconditional Forecasting Density

We examine the biases in density forecasts for the 1-month and 3-month horizon using option quotes with corresponding maturities. Our data set covers a sample period from January 2003 to August 2009. As a result of the initial estimation stage for the KFVV, forecasts are effectively produced as of January 2004.

The so-called peso problem is commonly put forward as an argument for the difference between the risk-neutral and statistical distributions (see e.g. Ait-Sahalia, Wang and Yared, 2001). Accordingly, option prices account for rare jumps which do not occur in the sample time series simply because the time frame considered is too short. To deal with this assertion, we deliberately include in our analysis the full sample until August 2009. This period covers the collapse of Lehman Brothers and the ensuing market turbulence in autumn 2008. During this time, all currency pairs have exhibited jump-like behavior. With high confidence therefore, the peso problem can be ruled out as an explanation for potential forecasting biases.

Our choice of currency pairs reflects the following considerations: First, we include only major currencies to prevent any issues with liquidity. Second, the currency pairs represent the different implied volatility functions commonly observed in foreign exchange. For example, USDJPY is highly skewed and thus more equity-like, whereas EURUSD is fairly symmetric over the sample period. Finally, we include EURGBP to examine whether our findings apply to so-called cross-biproducts. To enable a comparison of our results with those from Christoffersen and Mazzotta (2005), we base the main analysis on the 25-delta option quotes. Deep OTM options are considered in a separate section on robustness. To complement the data set, we obtain spot and interest rates from Bloomberg.

3.1 QQ Plots

To quantify the biases in the density forecasts for the various methods, we define a probability transform variable

$$U_{t,\tau} := \int_{-\infty}^{S_{t+\tau}} \hat{q}_{t,\tau}(u) du = \hat{Q}_{t,\tau}(S_{t+\tau}), \quad (20)$$

where $\hat{q}_{t,\tau}(\cdot)$ is the probability density function from (2) for options with tenor τ , observed on the forecasting date t . $S_{t+\tau}$ is the observed spot rate at the end of the forecasting horizon. The random variable $U_{t,\tau}$ expresses the probability that the spot rate at the end of the forecasting horizon is lower than the realization $S_{t+\tau}$. Under the joint-premise of a correctly specified forecasting methodology and equivalence between

the physical and risk-neutral currency return density, the variable $U_{t,\tau}$ is uniformly distributed on the interval $[0, 1]$.

For the remainder of section 3, we employ the testing framework of Berkowitz (2001) and Christoffersen and Mazzotta (2005). We apply the standard normal inverse cumulative density function to transform $U_{t,\tau}$ into a normal transform variable

$$Z_{t,\tau} := \Phi^{-1}(U_{t,\tau}) = \Phi^{-1}(\hat{Q}_{t,\tau}(S_{t+\tau})). \quad (21)$$

For the same reason $U_{t,\tau}$ is uniform on the interval $[0, 1]$, $Z_{t,\tau}$ is standard normally distributed. In Figure 1, we plot for the 1-month horizon the empirically observed $Z_{t,\tau}$ against the null hypothesis of standard normality. The so-called QQ plots are particularly useful to identify tail biases. To enhance inference, we augment the plots with 5% confidence bands from a Kolmogorov-Smirnov test.

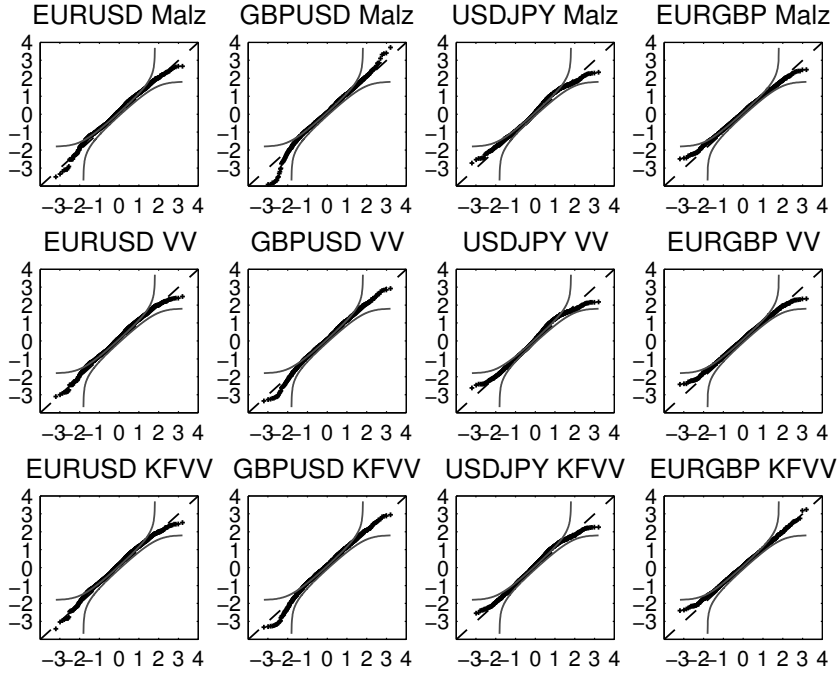


Figure 1: QQ Plots for the 1-month Horizon

The figure reveals QQ plots for the 1-month forecasting horizon for all currency pairs and methods under consideration. In case of the Malz and vanna-volga method, forecasts have been produced on the basis of 25-delta options. The horizontal and vertical axes refer to the realized and forecast quantile respectively. The data covers the period from January 2004 to August 2009.

To enhance interpretation of the QQ plots, consider the GBPUSD forecast from the Malz approach (second panel from left, top row). Under the null hypothesis that the

Malz approach produces unbiased forecasts of the statistical density, the very bottom-left realization $S_{t+\tau}$ should occur with a likelihood that corresponds to a Z-score of -4 (vertical axis). Empirically however, such an event turns out to be more common, attaining a Z-score of -3 (horizontal axis). In general, QQ plots coincide with the 45° line when the null hypothesis holds true. If, for a given quantile, the QQ plot lies below the diagonal, the risk-neutral density underestimates the amount of realizations in that quantile.

The results for the 1-month forecasting horizon reveal a strikingly uniform bias pattern across the methods. For EURUSD, all methods provide a fairly accurate fit of the statistical density. Concerning the GBPUSD forecasts from the Malz approach, the lower tail is biased, suggesting that more realizations occur in the lower tail than predicted. This bias, albeit smaller, is still present when forecasts are produced from the vanna-volga and the KFVV method. For USDJPY, we observe an S-shaped pattern that looks very similar for all methods. This finding suggests that the option-implied tail densities are too thick. A similar conclusion applies for EURGBP, except for the KFVV: In contrast to the simple methods, we find an almost perfect fit in the upper tail.

We present QQ plots for the 3-month forecasting horizon in Figure 2. Comparing them with the plots from Figure 1, we immediately find that the difference between the statistical and the risk-neutral density is more pronounced at the 3-month horizon. For EURUSD, too many observations populate the lower tail density relative to what is implied in option prices. In contrast, the upper tail of the risk-neutral density is too thick. These findings apply irrespective of the chosen forecasting method, though arguably, the lower tail bias is smaller for the vanna-volga and KFVV methods. For GBPUSD, we find similar results. In particular, all methods share a significant lower tail bias. For USDJPY, the S-shaped patterns observed at 1-month horizon prevail. Considerable biases in the lower tail are also observed for EURGBP. Concerning the upper tail, the forecasts from the vanna-volga method are fairly accurate, while the other two methods both fail to predict enough realizations.

Overall, we confirm the results in Christoffersen and Mazzotta (2005) in that we observe substantial deviations from the 45° line. Since similar biases are found for all methods, we conjecture that the statistical and risk-neutral densities in foreign exchange markets differ. To confirm this notion, we proceed with a formal test.




Figure 2: QQ Plots for the 3-month Horizon

The figure reveals QQ plots for the 3-month forecasting horizon for all currency pairs and methods under consideration. In case of the Malz and vanna-volga method, forecasts have been produced on the basis of 25-delta options. The horizontal and vertical axis refer to the realized and forecast quantile respectively. The data covers the period from January 2004 to August 2009.

3.2 A Formal Test of Currency Density Forecasts

We exploit the fact that under the null hypothesis of unbiased density forecasts, the moments of the normal transform variable $Z_{t,\tau}$ are given by

$$E[Z_{t,\tau}] = 0 \quad E[Z_{t,\tau}^2] = 1 \quad E[Z_{t,\tau}^3] = 0 \quad E[Z_{t,\tau}^4] = 3 \quad .$$

Since on each day, we produce density forecasts with either a 1- or 3-month horizon, $Z_{t,\tau}$ is serially correlated. To cope with this problem, we make use of the generalized method of moments (GMM) proposed by Hansen (1982). Specifically, we set up the

sample moment conditions

$$\bar{g}(\beta) = \frac{1}{T} \sum_{t=1}^T \begin{bmatrix} Z_{t,\tau} - \beta_1 \\ Z_{t,\tau}^2 - 1 - \beta_2 \\ Z_{t,\tau}^3 - \beta_3 \\ Z_{t,\tau}^4 - 3 - \beta_4 \end{bmatrix} = \mathbf{0}_{4 \times 1}. \quad (22)$$

Under the null hypothesis of unbiased density forecasts, $\beta_0 = (\beta_1 \dots \beta_4)' = \mathbf{0}$ and therefore

$$\sqrt{T}(\hat{\beta} - \beta_0) \xrightarrow{d} N(0, V). \quad (23)$$

We set V equal to the spectral density matrix of Newey and West (1987). To perform t-tests on β_i , we select the diagonal elements in V . To perform a joint-test on β_0 , we compute a Wald-statistic, where under the null hypothesis it holds that

$$\bar{g}(\beta)' \left(\frac{\hat{V}}{T} \right)^{-1} \bar{g}(\beta) \xrightarrow{d} \chi_4^2. \quad (24)$$

In Table 1, we present the results for the 1-month forecasting horizon. For EURUSD and GBPUSD, we do not reject the null hypothesis that option-implied density forecasts provide an unbiased estimate of the statistical return distribution. This finding holds for all methods, and both when we look at the individual moment parameters and a joint-test on all of them. The only significant parameter is the mean for the KFVV methodology, and only on a 10% confidence level. For USDJPY, our findings are different. The kurtosis parameter is significantly negative on the 5% and 10% level for the Malz and vanna-volga method, which indicates that risk-neutral densities have too much probability mass in the tails. Furthermore, the Wald-statistics of 63.32 and 85.54 strongly indicate rejection of the null hypothesis. While the kurtosis parameter is not rejected for the KFVV, its Wald-statistic of 17.67 still suggests rejection of the null hypothesis on the 1% confidence level. For EURGBP, we reject the option-implied density forecasts from the Malz and vanna-volga method both on the basis of the kurtosis parameter and the Wald-statistic. The Wald-statistic for the KFVV is significant only on the 10% level. This apparent difference is consistent with our findings from Figure 1, where the upper tail of the density function exhibits no bias when the KFVV methodology is used.

Results for the 3-month forecasting horizon are presented in Table 2. Despite the evidence of more pronounced biases from Figure 2, hardly any of the individual moment parameters are significantly different from zero. The reason for this finding is that

	EURUSD		GBPUSD		USDJPY		EURGBP	
	β	t-stat	β	t-stat	β	t-stat	β	t-stat
Malz								
Mean	0.163	1.584	0.109	1.018	0.107	1.088	0.096	1.091
Variance	0.018	0.151	0.135	0.804	-0.002	-0.025	-0.146	-1.588
Skew	0.298	0.874	0.048	0.085	0.223	0.985	0.197	0.912
Kurtosis	0.017	0.025	1.926	1.197	-0.779	-2.445**	-1.036	-2.809***
Wald-test	3.77		5.91		63.32***		24.01***	
VV								
Mean	0.162	1.564	0.105	0.991	0.081	0.792	0.096	1.074
Variance	0.023	0.199	0.101	0.704	0.056	0.621	-0.115	-1.226
Skew	0.308	0.977	0.097	0.226	0.076	0.310	0.206	0.929
Kurtosis	-0.210	-0.362	0.763	0.777	-0.557	-1.661*	-0.951	-2.613***
Wald-test	6.22		3.06		85.54***		31.36***	
KFVV								
Mean	0.170	1.659*	0.125	1.186	0.119	1.194	0.104	1.157
Variance	0.013	0.117	0.107	0.722	0.033	0.348	-0.101	-1.001
Skew	0.290	0.919	0.074	0.163	0.125	0.504	0.350	1.354
Kurtosis	-0.222	-0.374	0.927	0.872	-0.407	-0.961	-0.618	-1.213
Wald-test	6.40		5.01		17.67***		10.77*	

Table 1: Tests of the Quality of 1-month Density Forecasts

For each method and currency pair, the table presents the estimated moment parameters and their t-statistics. In case of the Malz and vanna-volga method, forecasts have been produced on the basis of 25-delta options. The rejection of the null hypothesis is indicated with *, ** and *** for the 10%, 5% and 1% significance level. For each panel, the last row shows the Wald-statistic for a joint-test on all parameters.

robust inference for the 3-month horizon requires the augmentation of the spectral density matrix with more lags, which generally reduces the power of the tests. We therefore focus on the Wald tests. For EURUSD, the Wald-statistic indicates rejection of the null hypothesis for the vanna-volga and the KFVV method on a 5% and 1% confidence level. The Malz approach has a Wald-statistic that is significant on the 10% level. For GBPUSD, USDJPY and EURGBP, the 3-month option-implied density forecasts are all rejected on the 1% confidence level, irrespective of the chosen method. The rejection of the USDJPY risk-neutral density forecasts appears to be rooted in a kurtosis mismatch, in particular for the Malz approach. For EURGBP, the skew parameter is positive on a 10% confidence level for all methods. A positive parameter suggests that the risk-neutral densities predict too many observations in the lower tail, which is what we also observe from the QQ plots. In sum, the findings strongly suggest that the 3-month risk-neutral densities provide a biased estimate of the physical return distribution.

	EURUSD		GBPUSD		USDJPY		EURGBP	
	β	t-stat	β	t-stat	β	t-stat	β	t-stat
Malz								
Mean	0.230	1.241	0.086	0.431	0.154	0.932	0.208	1.277
Variance	0.117	0.443	0.206	0.532	-0.047	-0.393	-0.110	-0.475
Skew	-0.230	-0.258	-1.247	-0.896	0.337	0.951	1.093	1.731*
Kurtosis	1.702	0.678	3.969	0.856	-1.032	-1.979**	0.232	0.135
Wald-test	10.95*		17.17***		48.51***		161.91***	
VV								
Mean	0.235	1.278	0.092	0.472	0.113	0.644	0.204	1.250
Variance	0.096	0.410	0.145	0.456	0.037	0.270	-0.117	-0.563
Skew	-0.017	-0.023	-0.803	-0.797	0.072	0.175	0.918	1.871*
Kurtosis	0.713	0.452	1.826	0.673	-0.694	-1.225	-0.536	-0.550
Wald-test	11.75**		19.58***		28.16***		169.70***	
KFVV								
Mean	0.240	1.319	0.107	0.549	0.160	0.948	0.236	1.345
Variance	0.086	0.368	0.165	0.497	-0.006	-0.043	0.040	0.134
Skew	-0.082	-0.115	-0.868	-0.814	0.224	0.548	1.574	1.820*
Kurtosis	0.651	0.406	2.102	0.720	-0.706	-1.148	1.804	0.733
Wald-test	17.80***		23.57***		15.59***		135.64***	

Table 2: Tests of the Quality of 3-month Density Forecasts

For each method and currency pair, the table presents the estimated moment parameters and their t-statistics. In case of the Malz and vanna-volga method, forecasts have been produced on the basis of 25-delta options. The rejection of the null hypothesis is indicated with *, ** and *** for the 10%, 5% and 1% significance level. For each panel, the last row shows the wald-statistic for a joint-test on all parameters.

3.3 Robustness Check Using Deep OTM Options

The Malz and vanna-volga forecasts need not be constructed on the basis of 25-delta options. Any set of pivot options works, though it is sensible to maintain symmetry, i.e. to select a call and put option with the same moneyness. Since we have 5-delta and 10-delta options available, we combine them with the delta-neutral quote to produce alternative density forecasts. Presumably, these forecasts convey different information on the tails of the physical return density and therefore alter some of the biases observed in Figure 1 and Figure 2. Figure 3 presents the results for the 1-month horizon. Since the KFVV method incorporates all option quotes in the first place, we refrain from displaying its results again.

A comparison of Figure 3 with Figure 1 suggests that there is little difference between the density forecasts constructed from the 25-delta and the deep OTM options: Fairly unbiased forecasts are obtained for EURUSD. For GBPUSD, forecasts from the Malz approach still exhibit a lower tail bias, which is somewhat smaller when the




Figure 3: Alternative QQ Plots for the 1-month Horizon

The figure reveals QQ plots for the 1-month forecasting horizon for the Malz and vanna-volga method. Forecasts have been constructed from 10-delta and 5-delta options. The horizontal and vertical axes refer to the realized and forecast quantile respectively. The data covers the period from January 2004 to August 2009.

vanna-volga method is used. For USDJPY and EURGBP, the S-shaped patterns, i.e. the overestimation of tail events, prevail. On the basis of Figure 3, we dismiss the notion that at the 1-month forecasting horizon, deep OTM options convey different information on the tails of the physical return distribution.

Figure 4 presents the corresponding results for the 3-month horizon. Again, we observe that the methods produce similar biases as previously for the 25-delta options. For example, the EURUSD and GBPUSD lower tail density is still biased for both methods, although the bias is larger when the Malz approach is used. For USDJPY, the S-shaped patterns mostly prevail, except for the 5-delta forecast from the Malz approach, which provides a surprisingly good fit of the upper tail density. For EURGBP,




Figure 4: Alternative QQ Plots for the 3-month Horizon

The figure reveals QQ plots for the 3-month forecasting horizon for the Malz and vanna-volga method. Forecasts have been constructed from 10-delta and 5-delta options. The horizontal and vertical axes refer to the realized and forecast quantile respectively. The data covers the period from January 2004 to August 2009.

the two methods overestimate the likelihood of a lower tail event, while they reveal slightly different biases in the upper tail. This holds equivalently for the forecasts constructed from 10-delta and 5-delta options.

Table 3 presents results from a formal test. For brevity, we report only the results from a joint-test on all moment parameters. The upper panel refers to the 1-month forecasting horizon. It confirms our findings from Table 1: The accuracy of density forecasts is not rejected for EURUSD and GBPUSD, while it is strongly rejected for USDJPY and EURGBP. For the EURUSD forecasts at the 3-month horizon, the Wald-statistics consistently indicate rejection of the null hypothesis on the 5% level. For GBPUSD, USDJPY and EURGBP, the null hypothesis of unbiased forecasts is rejected

	EURUSD	GBPUSD	USDJPY	EURGBP
1-month Forecast				
10D Malz	4.26	5.01	113.00***	38.34***
5D Malz	5.95	4.47	141.25***	57.04***
10D VV	6.28	3.08	79.87***	29.24***
5D VV	5.55	3.53	60.55***	26.57***
3-month Forecast				
10D Malz	12.24**	18.66***	97.04***	191.01***
5D Malz	13.70**	20.41***	1.21	223.11***
10D VV	11.76**	19.39***	25.62***	169.91***
5D VV	11.35**	19.35***	19.35***	176.15***

Table 3: Wald Tests for the Alternative Density Forecasts

For the 1-month and 3-month forecasting horizon, the table presents statistics from a Wald test on the estimated moment parameters and their t-statistics. The rejection of the null hypothesis is indicated with *, ** and *** for the 10%, 5% and 1% significance level.

on the 1% level. The one big exception to this finding are the density forecasts that result when the Malz approach is combined with the 5-delta options. The Wald-statistic reported in Table 3 is consistent with the accurate upper tail modelling of Figure 4, though it is not with the overestimation of realizations in the lower tail. Apart from this outlier, the results confirm the conjecture that the quality of density forecasts does not depend on the moneyness of the options employed in the forecasting process.

4 Forecasting Performance and Risk Factors

The previous section provides strong evidence that the risk-neutral currency return density is a biased estimator of the statistical distribution. In this section, we examine whether there is a systematic link between the biases in density forecasts and risk factors that in the equities literature have been identified to command a premium. Two factors are of particular interest: Lamoureux and Lastrapes (1993) conjecture that variance risk is priced. More recently, Bakshi and Kapadia (2003) and Carr and Wu (2009) report the presence of variance risk premiums in equity index and individual stock options. On the other hand, Jackwerth (2000) and Pan (2002), inter alia, find that for equities, jump risk is priced too.

4.1 Variance Risk Factor

We quantify forecasting biases by means of the normalized moments of the normal transform variable $Z_{t,\tau}$. Intuitively, the more the normalized moments differ from zero,

the worse the forecasting accuracy. Concerning variance risk, a good proxy is given by the prices for variance swap rates, since variance swaps are equivalent to model-free risk-neutral expectations of future variance.⁸

In their analysis of the conditional forecasting quality, Christoffersen and Mazzotta (2005) extend the framework in (22) and add (powers of) Black-Scholes implied volatilities as conditioning variables. Due to the non-stationary nature of variance swap rates, we fear that this approach leads to spurious results. For this reason, we consider first differences and interpret the results in terms of the impact a shock to the variance risk factor has on the various dimensions of forecasting accuracy. Formally, we define a vector of normalized moments $\mathbf{z}_t = (Z_{t,\tau}, Z_{t,\tau}^2 - 1, Z_{t,\tau}^3, Z_{t,\tau}^4 - 3)'$. Furthermore, denote with $RNV_{t,\tau}$ the variance swap rate observed at time t with expiry at $t + \tau$, and set the regressor $f_t = RNV_{t,\tau}$. We estimate the equation

$$\Delta \mathbf{z}_t = \beta \Delta \mathbf{f}_t + \epsilon_t, \quad (25)$$

by means of GMM. More precisely, we set up moment conditions

$$\mathbf{g}_t(\beta) = E(\Delta \mathbf{f}_t \otimes \epsilon_t) = E(\Delta \mathbf{f}_t \otimes (\Delta \mathbf{z}_t - \beta \Delta \mathbf{f}_t)) = \mathbf{0}_{4 \times 1}. \quad (26)$$

As previously for the unconditional setting, individual parameters are tested using (23), where V is replaced with V^*

$$\mathbf{V}^* = (\mathbf{D}_0' \mathbf{V}_0^{-1} \mathbf{D}_0)^{-1}, \quad \mathbf{D}_0 = \frac{1}{T} \sum_{t=1}^T (\Delta \mathbf{f}_t \Delta \mathbf{f}_t') \otimes \mathbf{I}_4. \quad (27)$$

\mathbf{I}_4 refers to the 4×4 identity matrix.

Since quotes on variance swaps are not available to us, we follow Carr and Wu (2009) and synthesize variance swap rates from a continuum of option prices in the strike domain. To obtain option prices in the first place, we interpolate the 25-delta quotes on a ± 6 standard deviations range around the current forward price using the Malz (1997) approach. In light of the subsequent jump risk analysis, we normalize variance swap rates to one.

Table 4 presents the results for the 1-month forecasting horizon. For brevity, we focus on the forecasting biases implied by the KFVV methodology. For EURUSD, a positive shock to the variance risk factor significantly reduces the variance parameter of the normalized transform variable. To interpret this causality, suppose the risk-neutral density forecast is unbiased, i.e. the normalized variance parameter is equal to zero. After the increase in variance, the variance parameter is negative, meaning that in a

⁸See Carr and Madan (1998) and Demeterfi, Derman, Kamal and Zou (1999).

	EURUSD		GBPUSD		USDJPY		EURGBP	
	β	t-stat	β	t-stat	β	t-stat	β	t-stat
Variance Factor								
Mean	0.036	0.355	0.278	4.723***	0.253	3.830***	-0.328	-6.176***
Variance	-0.342	-2.392***	-0.595	-1.941*	-0.161	-2.785***	-0.710	-2.705***
Skew	0.238	0.606	1.612	1.785*	0.420	2.376**	-2.058	-2.912***
Kurtosis	-2.174	-2.017**	-4.657	-1.861*	-0.417	-1.801*	-5.280	-2.550**

Table 4: Variance Risk Factor and Forecasting Quality at the 1-month Horizon

For each currency pair, the table shows parameter estimates and t-statistics from a regression of differences in the normalized moments of the normal transform variable $Z_{t,\tau}$ on changes in variance swap rates. The moments of $Z_{t,\tau}$ are obtained from KFVV density forecasts. Variance swap rates are constructed on the basis of the Malz approach in conjunction with 25-delta options. Significance of parameters is indicated with *, ** and *** for the 10%, 5% and 1% confidence level.

QQ plot, we observe a transition from the diagonal to a line with slope below 45° . In terms of a distribution function, this implies that market participants expect too much volatility in the statistical return distribution. This form of overshooting is observed for all currency pairs and is highly significant for USDJPY and EURGBP. A similar finding applies for the kurtosis parameter: A positive shock to the variance risk factor decreases the kurtosis parameter, which indicates that option-implied densities predict too many observations in the tails. Clearly, this behavior is consistent with risk-averse investors.

The mean parameter is not significant for EURUSD, but is highly significant for all other currency pairs. For GBPUSD and USDJPY, a positive shock to the variance factor shifts the diagonal in a QQ plot to the left. The implication is that relative to the statistical distribution, the risk-neutral density is shifted too much to the left. For EURGBP, the opposite holds true. The skew parameter is significant for GBPUSD, USDJPY and EURGBP with increasing confidence. The positive sign for GBPUSD and USDJPY indicates that after a shock to the variance factor, the risk-neutral density exhibits skew that is too negative compared to the statistical distribution. In contrast, it holds for EURGBP that after a variance shock, market participants tend to expect too many observations in the upper tail. The sign on both the mean and skew parameter suggest that for EURUSD, GBPUSD and USDJPY, investors perceive the downside as the bad state of the economy, just as is the case for equities. This interpretation is significant for GBPUSD and highly significant for USDJPY. For EURGBP, the opposite holds true.

Table 5 presents the corresponding results for the 3-month horizon. For each cur-

	EURUSD		GBPUSD		USDJPY		EURGBP	
	β	t-stat	β	t-stat	β	t-stat	β	t-stat
Variance Factor								
Mean	0.024	0.362	0.315	5.127***	0.303	2.866***	-0.293	-4.404***
Variance	-0.062	-0.626	-0.603	-1.543	-0.280	-2.107**	-0.749	-2.095**
Skew	0.172	0.720	2.124	2.183**	0.760	2.976***	-2.846	-1.862*
Kurtosis	-0.910	-1.005	-5.013	-1.679*	-0.874	-1.631	-10.61	-1.692*

Table 5: Variance Risk Factor and Forecasting Quality at the 3-month Horizon

For each currency pair, the table shows parameter estimates and t-statistics from a regression of differences in the normalized moments of the normal transform variable $Z_{t,\tau}$ on changes in variance swap rates. The moments of $Z_{t,\tau}$ are obtained from KFVV density forecasts. Variance swap rates are constructed on the basis of the Malz approach in conjunction with 25-delta options. Significance of parameters is indicated with *, ** and *** for the 10%, 5% and 1% confidence level.

currency pair, none of the parameters has changed sign, so the interpretations for the 1-month horizon still apply. In general, we observe that the parameters are somewhat less significant than before, which we again attribute to the use of robust standard errors. Taken together with the results from Table 4, Table 5 suggests that the variance risk factor strongly impacts the forecasting accuracy. More precisely, the quality of option-implied density forecasts is affected on several accounts, including the mean, variance, skew and kurtosis of the hypothesized density.

4.2 Jump Risk Factor

Similarly to variance risk, fear of jumps may too have a distorting effect on forecasting accuracy. To test this hypothesis, suitable jump surrogates are needed. Bakshi and Kapadia (2003) propose to use risk-neutral skew and kurtosis. The rationale behind this choice is that skew proxies for the mean jump size, whereas kurtosis is a surrogate for jump intensity.

In analyzing the impact of jump risk on the quality of density forecasts, we adhere to the econometric setting proposed in (25) to (27). Specifically, we augment the regressor f_t with a proxy for the risk-neutral skew $\theta_{t,\tau}$ and kurtosis $\kappa_{t,\tau}$, i.e. $\mathbf{f}_t = (RNV_{t,\tau}, \theta_{t,\tau}, \kappa_{t,\tau})'$. As a result, (26) comprises 12 moment conditions.

We construct the model-free jump risk metrics in analogy to the variance risk factor: First, we interpolate 25-delta option prices using the Malz (1997) approach. Second, we compute prices for so-called volatility, cubic and quartic contracts. Finally, we obtain the risk-neutral skew and kurtosis. For details, we refer to Bakshi, Kapadia and Madan (2003).

	EURUSD		GBPUSD		USDJPY		EURGBP	
	β	t-stat	β	t-stat	β	t-stat	β	t-stat
Variance Factor								
Mean	0.070	0.749	0.276	3.829***	0.308	5.730***	-0.268	-5.000***
Variance	-0.289	-1.997**	-0.542	-1.796*	-0.156	-2.117**	-0.670	-2.515**
Skew	0.348	0.859	1.607	1.703*	0.520	3.290***	-1.887	-2.582***
Kurtosis	-2.017	-1.862*	-4.284	-1.763*	-0.372	-1.232	-5.066	-2.434**
Skew Factor								
Mean	-3.564	-13.85***	-3.345	-6.398***	-2.237	-4.383***	-2.343	-10.07***
Variance	-0.730	-0.821	-0.942	-0.929	-1.071	-1.662*	-0.492	-0.713
Skew	-10.70	-6.291***	-8.689	-3.703***	-5.506	-3.312***	-4.700	-2.546**
Kurtosis	0.686	0.102	-10.48	-1.810*	-5.133	-1.793*	0.810	0.167
Kurtosis Factor								
Mean	-0.649	-3.124***	-1.153	-5.347***	-0.502	-2.517**	-0.040	-0.278
Variance	0.351	0.973	0.196	0.436	-0.381	-1.319	0.328	0.941
Skew	-1.873	-2.460**	-2.991	-3.025***	-1.462	-2.239**	0.555	0.716
Kurtosis	1.824	0.716	0.068	0.030	-1.683	-1.430	2.974	1.703*

Table 6: Jump Risk Factors and Forecasting Quality at the 1-month Horizon

For each currency pair, the table shows parameter estimates and t-statistics from a regression of differences in the normalized moments of the normal transform variable $Z_{t,\tau}$ on changes in variance swap rates, risk-neutral skew and risk-neutral kurtosis. The moments of $Z_{t,\tau}$ are obtained from KFVV density forecasts. Variance swap rates and risk-neutral higher-order moments are constructed on the basis of the Malz approach in conjunction with 25-delta options. Significance of parameters is indicated with *, ** and *** for the 10%, 5% and 1% confidence level.

Table 6 reveals the results for the 1-month horizon, where again we focus on density forecasts from the KFVV method. Despite the presence of jump risk factors, the impact of the variance risk factor is almost unchanged. In particular, the estimates on the moment parameters have the same sign and a similar magnitude as before. Concerning the jump size proxy, similar results are obtained for all currency pairs. The parameters on the first and third moment of the normal transform variable are significantly negative. In contrast, there is only weak evidence that the jump size factor impacts the variance and kurtosis of the bias measure. The parameter estimates suggest that a positive shock to the jump size proxy leads to a right shift of the risk-neutral density that is not matched by the statistical distribution. As a result, the likelihood of predicting too few observations up to a given quantile is increased. Similar findings apply to the skew moment parameter: When the jump size factor increases, there is a tendency to overestimate the upper tail density. Given that we approximate the mean jump size by risk-neutral skewness, this result is to be expected.

The jump intensity surrogate has a significant impact on the mean and skew parameter for EURUSD, GBPUSD and USDJPY, but otherwise does not affect forecasting

quality. In particular, we find little evidence that EURGBP density forecasts are affected by jump intensity. From those parameter estimates that are significant, we infer that the intensity factor works in the opposite direction of the variance risk factor. An estimated increase in the number of jumps results in a negative mean shift of the normal transform variable. This implies that in tendency, spot realizations come lower than predicted. Furthermore, from the estimates on the skew moment parameter, we conclude that after a shock to the jump intensity risk factor, the upper tail density tends to be too thick.

	EURUSD		GBPUSD		USDJPY		EURGBP	
	β	t-stat	β	t-stat	β	t-stat	β	t-stat
Variance Factor								
Mean	0.106	1.712*	0.352	3.937***	0.401	4.182***	-0.216	-4.361***
Variance	0.006	0.065	-0.556	-1.590	-0.161	-1.116	-0.691	-1.947*
Skew	0.427	1.303	2.117	2.186**	0.989	3.677***	-2.741	-1.762*
Kurtosis	-0.865	-0.831	-4.606	-1.814*	-0.530	-0.914	-10.71	-1.656*
Skew Factor								
Mean	-2.248	-5.363***	-2.233	-5.204***	-2.072	-5.102***	-0.942	-2.186**
Variance	-1.214	-1.476	-2.846	-2.805***	-2.346	-2.850***	0.299	0.301
Skew	-6.389	-4.768***	-1.018	-0.267	-4.850	-3.142***	1.728	0.501
Kurtosis	0.804	0.135	-21.87	-1.913*	-6.682	-2.118**	11.638	1.048
Kurtosis Factor								
Mean	-0.365	-3.021***	-0.486	-4.559***	-0.508	-4.769***	0.003	0.051
Variance	-0.083	-0.501	-0.621	-2.708***	-0.547	-2.755***	0.258	1.474
Skew	-0.929	-3.159***	-0.300	-0.363	-1.196	-3.253***	0.771	1.392
Kurtosis	0.471	0.502	-4.620	-2.061**	-1.541	-2.014**	2.645	1.632

Table 7: Jump Risk Factors and Forecasting Quality at the 3-month Horizon

For each currency pair, the table shows parameter estimates and t-statistics from a regression of differences in the normalized moments of the normal transform variable $Z_{t,\tau}$ on changes in variance swap rates, risk-neutral skew and risk-neutral kurtosis. The moments of $Z_{t,\tau}$ are obtained from KFVV density forecasts. Variance swap rates and risk-neutral higher-order moments are constructed on the basis of the Malz approach in conjunction with 25-delta options. Significance of parameters is indicated with *, ** and *** for the 10%, 5% and 1% confidence level.

In Table 7, we present the results for the 3-month forecasting horizon. In essence, they confirm what we have observed before. The variance risk factor retains its mostly significant role in the presence of the jump risk surrogates. For EURUSD and USDJPY, the jump size factor plays a similar role as for the 1-month horizon. For GBPUSD, the parameter estimate on the third moment of the normal transform variable is no longer significant. In contrast, we observe a significant negative relationship for the variance parameter. Therefore, a positive shock to the jump size factor leads to an overestimation of the volatility of the statistical return distribution. Concerning the

jump intensity risk factor, similar results as for the 1-month horizon are obtained for EURUSD and EURGBP. For GBPUSD and USDJPY, jump intensity has a more complex impact on forecasting accuracy. In particular, we find a significant negative relationship for the variance and kurtosis parameter of the normal transform variable. Therefore, both at the 1-month and 3-month horizon, we conjecture that fear of jump risk affects the quality of density forecasts.

5 Conclusion

In this paper, we have examined the biases in density forecasts implied in currency options. To ensure that our evidence is robust, we have constructed density forecasts from three alternative interpolation methods: Apart from the well-established Malz and vanna-volga method, we proposed to enhance the latter by applying a Kalman filter. This approach dynamically updates the information content of all currency options available in the strike domain. Our results suggest that option-implied densities provide biased estimates of the statistical return distribution. This finding is persistent for all methods, and it also holds when we select alternative information sets comprised of deep OTM options. Furthermore, we observe that the biases become more pronounced for a longer forecasting horizon.

To further investigate the nature of the biases, we have established a relationship between our measure of forecasting accuracy and commonly priced risk factors in currency markets. We have found that for all currency pairs, at least some moments of the bias measure are significantly affected by surrogates for variance and jump risk. Furthermore, the variance risk factor retains a dominant role in the presence of jump risk proxies. Our results support the interpretation that risk aversion determines the difference between the risk-neutral and statistical return distributions in currency markets.

6 Appendix

A Quotation Convention Adapted Version of the Malz Interpolation Method

Given a set of option quotes σ^{DP} , σ^{DN} and σ^{DC} , we retrieve strike prices and compute spot deltas Δ_P , Δ_{DN} and Δ_C that are consistent with a currency pair's premium convention. For details, see Reiswich and Wystup (2010).

Define the vector $B = (\sigma^{DN}, \sigma^{rr}(\Delta - \Delta_{DN}), \sigma^{bf}(\Delta - \Delta_{DN})^2)'$, where the risk reversal and butterfly are constructed as in (4). Next, we set up the matrix

$$\mathbf{A} = \begin{pmatrix} \sigma^{DN} & 0 & 0 \\ 0 & \sigma^{RR}(\Delta_C - \Delta_P) & \sigma^{BF}((\Delta_C - \Delta_{DN})^2 - (\Delta_P - \Delta_{DN})^2) \\ 0 & 0.5\sigma^{RR}(\Delta_C + \Delta_P - 2\Delta_{DN}) & 0.5\sigma^{BF}((\Delta_C - \Delta_{DN})^2 + (\Delta_P - \Delta_{DN})^2) \end{pmatrix} \quad (\text{A.1})$$

The rows of matrix A reflect the right hand side of (5) for the delta-neutral quote, the risk reversal and the butterfly. Set $\sigma^{\mathbf{p}\mathbf{v}} = (\sigma^{DN}, \sigma^{RR}, \sigma^{BF})'$ and define a vector of parameters $\mathbf{b} = (b_0, b_1, b_2)'$. Since the convention-adapted version of (5) must hold for any quote in delta space, the parameters follow from

$$\mathbf{b} = \mathbf{A}^{-1}\sigma^{\mathbf{p}\mathbf{v}}. \quad (\text{A.2})$$

The convention-adapted version of (5) is given by

$$\sigma^\Delta = \mathbf{B}'\mathbf{b}. \quad (\text{A.3})$$

To transform an array of volatilities σ^Δ into an array of option prices $\mathbf{c}(\mathbf{X}, \tau)$, we insert the convention-adapted Garman and Kohlhagen (1983) formula for Δ into (A.3). Given an array of exercise prices \mathbf{X} , we numerically solve for $\sigma^\mathbf{X}$. Again applying the Garman and Kohlhagen formula, we finally obtain $\mathbf{c}(\mathbf{X}, \tau)$.

B The Vanna-Volga Method

Given a set of option quotes σ^{xDP} , σ^{DN} and σ^{xDC} , we compute strike prices along the lines of Reiswich and Wystup (2010). Next, we compute volatility greeks and assemble the vector \mathbf{y} and matrix \mathbf{A} ,

$$\mathbf{y} = \begin{pmatrix} \frac{\partial c(X, \tau, S_t, \sigma_t^X)}{\partial \sigma_t} \\ \frac{\partial^2 c(X, \tau, S_t, \sigma_t^X)}{\partial S_t \partial \sigma_t} \\ \frac{\partial^2 c(X, \tau, S_t, \sigma_t^X)}{\partial \sigma_t^2} \end{pmatrix}, \quad (\text{B.1})$$

and

$$\mathbf{A} = \begin{pmatrix} \frac{\partial c(X_{xDP}, \tau, S_t, \sigma_t^{xDP})}{\partial \sigma_t} & \frac{\partial c(X_{DN}, \tau, S_t, \sigma_t^{DN})}{\partial \sigma_t} & \frac{\partial c(X_{xDC}, \tau, S_t, \sigma_t^{xDC})}{\partial \sigma_t} \\ \frac{\partial^2 c(X_{xDP}, \tau, S_t, \sigma_t^{xDP})}{\partial S_t \partial \sigma_t} & \frac{\partial^2 c(X_{DN}, \tau, S_t, \sigma_t^{DN})}{\partial S_t \partial \sigma_t} & \frac{\partial^2 c(X_{xDC}, \tau, S_t, \sigma_t^{xDC})}{\partial S_t \partial \sigma_t} \\ \frac{\partial^2 c(X_{xDP}, \tau, S_t, \sigma_t^{xDP})}{\partial \sigma_t^2} & \frac{\partial^2 c(X_{DN}, \tau, S_t, \sigma_t^{DN})}{\partial \sigma_t^2} & \frac{\partial^2 c(X_{xDC}, \tau, S_t, \sigma_t^{xDC})}{\partial \sigma_t^2} \end{pmatrix}. \quad (\text{B.2})$$

The vega and volga of an option are given by

$$\frac{\partial c(\cdot)}{\partial \sigma_t} = e^{-r_f \tau} S_t \frac{e^{-d_1^2/2}}{\sqrt{2\pi}} \sqrt{\tau} \quad \frac{\partial^2 c(\cdot)}{\partial \sigma_t^2} = \frac{\partial c(\cdot)}{\partial \sigma_t} d_1 \frac{d_2}{\sigma_t}, \quad (\text{B.3})$$

where $d_1 = \frac{\log(S_t/X) + (r_d - r_f + \sigma_t^2/2)\tau}{\sigma_t \sqrt{\tau}}$ and $d_2 = d_1 - \sigma_t \sqrt{\tau}$.

The vanna depends on the premium convention. The regular (reg.) and premium-adjusted (p.a.) vanna are given by

$$\frac{\partial^2 c(\cdot)}{\partial \sigma_t \partial S_t} \stackrel{(reg.)}{=} -e^{-r_f \tau} N'(d_1) \frac{d_2}{\sigma_t} \quad \frac{\partial^2 c(\cdot)}{\partial \sigma_t \partial S_t} \stackrel{(p.a.)}{=} \frac{\partial^2 c(\cdot)}{\partial \sigma_t \partial S_t} \stackrel{(reg.)}{=} \frac{K}{F_t}, \quad (\text{B.4})$$

where $N'(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$ and $F_t = S_t e^{(r_d - r_f)\tau}$.

The market excess prices used in (8) are given by

$$\mathbf{c}^{\text{me}} = \begin{pmatrix} c(X_{xDP}, \tau, S_t, \sigma_t^{xDP}) - c(X_{xDP}, \tau, S_t, \sigma_t^{BS}) \\ c(X_{DN}, \tau, S_t, \sigma_t^{DN}) - c(X_{DN}, \tau, S_t, \sigma_t^{BS}) \\ c(X_{xDC}, \tau, S_t, \sigma_t^{xDC}) - c(X_{xDC}, \tau, S_t, \sigma_t^{BS}) \end{pmatrix}. \quad (\text{B.5})$$

Inserting (B.1), (B.2) and (B.5) into (7) and (8), we numerically solve for $\sigma^{\mathbf{X}}$ given an array of exercise prices X . A numerical procedure is required since \mathbf{y} and hence the weights \mathbf{x} depend on $\sigma^{\mathbf{X}}$. To map $\sigma^{\mathbf{X}}$ to $\mathbf{c}(\mathbf{X}, \tau)$, we apply the Garman and Kohlhagen (1983) function.

C Maximum Likelihood Estimation of the Kalman Filter

We refer to Hamilton (1994) for a detailed exposition of the Kalman filter. Under the assumption that ξ_1 and $\{\varepsilon_t, \omega_t\}_{t=1}^T$ are multivariate Gaussian, the conditional distribution of \mathbf{c}_t^{me} is given by

$$\mathbf{c}_t^{\text{me},X} \mid \mathbf{Y}_t, \mathcal{I}_{t-1} \sim \mathbf{N}(\mathbf{Y}_t' \hat{\xi}_{t|t-1} + \mathbf{Y}_t' \bar{\alpha}, \mathbf{Y}_t' \mathbf{P}_{t|t-1} \mathbf{Y}_t + \mathbf{R}), \quad (\text{C.1})$$

i.e.

$$f_{\mathbf{c}_t^{\text{me},X} \mid \mathbf{Y}_t, \mathcal{I}_{t-1}}(\mathbf{c}_t^{\text{me},X} \mid \mathbf{Y}_t, \mathcal{I}_{t-1}) = (2\pi)^{-\frac{T}{2}} \mid Y_t' P_{t|t-1} Y_t + R \mid^{-\frac{1}{2}} e^{\{-\frac{1}{2}(\mathbf{c}_t^{\text{me},X} - Y_t' \hat{\xi}_{t|t-1} - Y_t' \bar{\alpha})' (Y_t' P_{t|t-1} Y_t + R)^{-1} (\mathbf{c}_t^{\text{me},X} - Y_t' \hat{\xi}_{t|t-1} - Y_t' \bar{\alpha})\}}. \quad (\text{C.2})$$

The sample log likelihood function is therefore

$$\begin{aligned} \sum_{t=1}^T \log f(\mathbf{c}_t^{\text{me},X} \mid \mathbf{Y}_t, \mathcal{I}_{t-1}) &= -\left(\frac{T}{2}\right) \log(2\pi) - \\ &\frac{1}{2} \sum_{t=1}^T \log \mid Y_t' P_{t|t-1} Y_t + R \mid - \frac{1}{2} \sum_{t=1}^T (\mathbf{c}_t^{\text{me},X} - Y_t' \hat{\xi}_{t|t-1} - Y_t' \bar{\alpha})' \\ &(Y_t' P_{t|t-1} Y_t + R)^{-1} (\mathbf{c}_t^{\text{me},X} - Y_t' \hat{\xi}_{t|t-1} - Y_t' \bar{\alpha}). \end{aligned} \quad (\text{C.3})$$

The parameters are found by maximizing (C.3).

References

- Ait-Sahalia, Yacine and Duarte, Jefferson, 2003. "Nonparametric option pricing under shape restrictions." *Journal of Econometrics* 116, 9-47.
- Ait-Sahalia, Yacine and Lo, Andrew W., 1998. "Nonparametric estimation of state-price densities implicit in financial asset prices." *The Journal of Finance* 53, 499-547.
- Ait-Sahalia, Yacine and Lo, Andrew W., 2000. "Nonparametric risk management and implied risk aversion." *Journal of Econometrics* 94, 9-51.
- Ait-Sahalia, Yacine, Wang, Yubo and Yared, Francis, 2001. "Do option markets correctly price the probabilities of movement of the underlying asset?" *Journal of Econometrics* 102, 67-110.
- Arrow, Kenneth J., 1964. "The role of securities in the optimal allocation of risk bearing." *Review of Economic Studies* 31, 91-96.
- Bahra, Bhupinder, 1996. "Probability distributions of future asset prices implied by option prices." *Bank of England Quarterly Bulletin*.
- Bakshi, Gurdip and Kapadia, Nikunj, 2003. "Delta-hedged gains and the negative market volatility risk premium." *The Review of Financial Studies* 16, 527-566.
- Bakshi Gurdip, Kapadia, Nikunj and Madan, Dilip, 2003. "Stock return characteristics, skew laws, and the differential pricing of individual options." *The Review of Financial Studies* 16, 101-143.
- Bedendo, Mascia and Hodges, Stewart D., 2009. "The dynamics of the volatility skew: A Kalman filter approach." *Journal of Banking & Finance* 33, 1156-1165.
- Berkowitz, Jeremy, 2001. "Testing density forecasts, with applications to risk management." *Journal of Business & Economic Statistics* 19, 465-474.
- Black, Fischer and Scholes, Myron, 1973. "The pricing of options and corporate liabilities." *Journal of Political Economy* 81, 637-654.
- Bollerslev, Tim, 1986. "Generalized autoregressive conditional heteroskedasticity." *Journal of Econometrics* 31, 307-327.
- Breeden, Douglas T. and Litzenberger, Robert H., 1978. "Prices of state-contingent claims implicit in option prices." *The Journal of Business* 51, 621-651.

- Carr, Peter and Madan, Dilip, 1998. "Towards a theory of volatility trading." In Robert Jarrow (ed.), *Risk Book on Volatility*. New York: Risk Publications, 417-427.
- Carr, Peter and Wu, Liuren, 2009. "Variance risk premiums." *The Review of Financial Studies* 22, 1311-1341.
- Castagna, Antonio and Mercurio, Fabio, 2007. "The vanna-volga method for implied volatilities." *Risk*, 106-111.
- Christoffersen, Peter and Mazzotta, Stefano, 2005. "The accuracy of density forecasts from foreign exchange options." *Journal of Financial Econometrics* 3, 578-605.
- Debreu, Gerard, 1959. "Theory of Value." Wiley, New York.
- Demeterfi Kresimir, Derman, Emanuel, Kamal, Michael and Zou, Joseph, 1999. "A guide to volatility and variance swaps." *The Journal of Derivatives* 7, 9-32.
- Garman, Mark B. and Kohlhagen, Steven W., 1983. "Foreign currency option values." *Journal of International Money and Finance* 2, 231-237.
- Hamilton, James D., 1994. "Time series analysis." Princeton University Press, Princeton, New Jersey.
- Hansen, Lars Peter, 1982. "Large sample properties of generalized method of moments estimators." *Econometrica* 50, 1029-1054.
- Jackwerth, Jens C. and Rubinstein, Mark, 1996. "Recovering probability distributions from option prices." *The Journal of Finance* 51, 1611-1631.
- Jackwerth, Jens C., 2000. "Recovering risk aversion from option prices and realized returns." *The Review of Financial Studies* 13, 433-451.
- Lamoureux Christopher G. and Lastrapes, William D., 1993. "Forecasting stock-return variance: toward an understanding of stochastic implied volatilities." *The Review of Financial Studies* 6, 293-326.
- Malz, Allan M., 1997. "Estimating the probability distribution of the future exchange rate from option prices." *Journal of Derivatives* 5, 18-36.
- Melick, William R. and Thomas, Charles P., 1997. "Recovering an asset's implied pdf from option prices: An application to crude oil during the gulf crisis." *Journal of Financial and Quantitative Analysis* 32, 91-115.

Newey, Whitney K. and West, Kenneth D., 1987. "A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix." *Econometrica* 55, 703-708.

Pan, Jun, 2002. "The jump-risk premia implicit in options: evidence from an integrated time-series study." *Journal of Financial Economics* 63, 3-50.

Reiswich, Dimitri and Wystup, Uwe, 2010. "A guide to FX options quoting conventions." *The Journal of Derivatives* 18, 58-68.

Shkolnikov, Yuriy, 2009. "Generalized vanna-volga method and its applications." Working paper.

Part II

Variance Risk Premiums in Foreign Exchange Markets

Abstract

Based on the theory of static replication of variance swaps we assess the sign and magnitude of variance risk premiums in foreign exchange markets. We find significantly negative risk premiums when realized variance is computed from intraday data with low frequency. As a likely consequence of microstructure effects however, the evidence is ambiguous when realized variance is based on high-frequency data. Common to all estimates, variance risk premiums are highly time-varying and inversely related to the risk-neutral expectation of future variance.

When we test whether variance risk premiums can be attributed to classic risk factors or fear of jump risk, we find that conditional premiums remain significantly negative. However, we observe a strong relationship between the size of log variance risk premiums and the VIX, the TED spread and the general shape of the implied volatility function of the corresponding currency pair. Overall, we conclude that there is a separately priced variance risk factor which commands a highly time-varying premium.

1 Introduction

The increase of traded volumes in foreign exchange derivative markets over the past decades suggests that it becomes ever more important to understand risk factors and their potential premiums in currency markets. The aim of this paper is a careful examination of the variance risk premiums in foreign exchange markets. Specifically, we investigate the sign, size and evolution of variance risk premiums by means of the model-free approach of Carr and Wu (2009). Based on the theoretical work from Carr and Madan (1998), Demeterfi, Derman, Kamal and Zou (1999) and Britten-Jones and Neuberger (2000), we synthesize zero cost variance swaps, which is equivalent to constructing risk-neutral forecasts of future variance. Using a model-free variance estimator has the advantage that we can be agnostic about the volatility process of the underlying exchange rate. More precisely, our estimator produces robust forecasts under an arbitrary volatility process and thus avoids an important source of potential error. When variance swaps are compared with a measure of ex-post realized variance, the sign and magnitude of the average variance risk premiums can be directly inferred.

To the best of our knowledge, we are the first to apply the model-free methodology proposed by Carr and Wu (2009) to study variance risk premiums in foreign exchange markets. Our analysis is based on OTC options, for which only a handful of standard quotes in the strike domain are available. For this reason, we compute variance risk premiums for different interpolation methods and settings. We also propose a novel interpolation technique that extends the information set to all option quotes in the strike domain. Finally, we contribute to the literature by carefully examining the relationship between variance risk premiums and the sampling frequency of the spot rate on which realized variance estimates are based.

A number of studies on currency markets document biases for option-implied volatility in predicting future realized volatility. Early work on the subject includes Scott (1992), who introduces the notion of a volatility risk premium, Jorion (1995) and Bates (1996a). Covrig and Low (2003), Christoffersen and Mazzotta (2005) and Charoenwong, Jenwittayaroje and Low (2009) use OTC options to study the accuracy of implied volatility forecasts. While their evidence is ambiguous as to whether or not implied volatility is a biased predictor of future realized volatility, they agree that implied volatility subsumes the information contained in competing time-series models. Martens and Zein (2004) and Pong, Shackleton, Taylor and Xu (2004) compare implied volatility with forecasts from high-frequency historical data. They conclude that the latter provide accurate forecasts of future realized volatility. Contrary to this

body of research, we do not benchmark implied against historical volatility forecasts. Instead, our attention is devoted to a thorough analysis of the variance risk premiums in currency markets.

Variance risk premiums in equity markets are relatively well studied. For example, Coval and Shumway (2001) and Bakshi and Kapadia (2003) conduct analyses with a focus on the performance of hedged option positions. The former construct so-called zero-beta index straddles, while the latter examine returns to delta-neutral call option strategies. Both report significant negative returns and attribute these to negatively priced variance risk. Carr and Wu (2009) quantify variance risk premiums for both index options and individual stocks. Although they report some cross-sectional differences for the latter, the overall evidence is strongly indicative of negative variance risk premiums. Currency markets in contrast have so far received little attention, in spite of their very distinct nature compared to equities. Guo (1998) investigates variance risk premiums in the context of the Heston (1993) stochastic volatility model, whereas we assess them in a model-free manner. Low and Zhang (2005) adapt the approach of Bakshi and Kapadia (2003). Compared to their analysis, our approach has two distinct advantages: First, we can directly quantify the magnitude of variance risk premiums. Second, we account for the information in the cross-section of option prices, whereas their evidence rests entirely on the at-the-money quotes.

Currency markets are fundamentally different from equities in that one of the key explanations for negative variance risk premiums does not necessarily apply. The classical argument goes as follows: Since equity investors are primarily concerned with a decrease in share prices, and since negative returns tend to coincide with an increase in volatility, instruments with a positive exposure to volatility pay out in bad states of the economy. As such, risk-averse investors should be willing to pay a premium for holding such instruments. In foreign exchange, the relationship between the level of volatility and the direction of the underlying currency pair is not as clear-cut. Evidently, the so-called leverage effect first pointed out by Black (1976) is absent. More importantly, there are likewise domestic and foreign investors and firms with opposite interest in the valuation of one currency against another. As a result, a currency depreciation need not be a bad thing. Thinking in the context of the mean-variance framework, an increase in volatility is however likely to adversely affect the opportunity set of an international investor. Furthermore, it impedes the budgeting and planning process of an internationally operating firm. Provided market participants are risk-averse, this reasoning suggests that potentially negative variance risk premiums can be attributed to a separately priced variance risk factor.

In line with the economic argument, we find significantly negative variance risk premiums when realized variance is computed from intraday data with low frequencies. However, we report a considerable difference in average variance risk premiums when spot data with daily sampling frequency as opposed to high-frequency data is used. Our results suggest that the observed discrepancies are owed to microstructure effects that come into play as the sampling frequency is increased. We can further assert that variance risk premiums are highly time-varying and inversely related to the risk-neutral expectation of future variance. Finally, our results are robust to whether or not we include data covering the financial crisis of 2008.

In an attempt to better comprehend the nature of variance risk premiums, we interpret our results in the context of classic risk factors. Specifically, we regress log variance risk premiums on excess returns in the S&P 500, returns on the VIX and first differences in the TED spread. While the latter two share a significant relationship with the magnitude of log variance risk premiums, the conditional premiums remain significantly negative. We also examine whether variance risk premiums subsume fear of jump risk. Assuming that jump risk is well proxied by the prices for risk reversal and butterfly strategies, we can conclude that jump risk aversion cannot account for the observed variance risk premiums in currency markets. Overall, our results point towards an independent variance risk factor which commands a time-varying premium.

The remainder of the paper is organized as follows. Section 2 provides an outline of the general estimation methodology applied in this paper. In section 3, we present the data set and elaborate on the details of replicating variance swaps in foreign exchange markets. Section 4 provides evidence on average variance risk premiums as well as their time series characteristics. Section 5 investigates variance risk premiums during the financial crisis of 2008. In section 6, we assess the variance risk premiums within the framework of classical factor models. Section 7 concludes.

2 Static Hedging and Model-Free Variance Forecasting

Throughout this paper, we work with the risk-neutral variance forecast developed in Carr and Madan (1998), Demeterfi et al. (1999) and Britten-Jones and Neuberger (2000). We follow the common foreign exchange quotation convention and assume that the evolution of a currency pair under the risk-neutral measure \mathbb{Q} is governed by

a stochastic differential equation (SDE) of the form

$$dS(t) = (r_d - r_f)S(t)dt + \sigma(t)S(t)dB(t), \quad (1)$$

where $S(t)$ is the price of a foreign currency measured in domestic units. As is readily seen from (1), the domestic and foreign interest rates r_d and r_f are assumed to be constant. In contrast, no assumption is made with regard to the stochastic dynamics of the volatility $\sigma(t)$. In fact, $\sigma(t)$ represents an arbitrary stochastic process that we leave unspecified throughout the paper. Given the SDE in (1), it follows that the risk-neutral expectation of future variance is given by¹

$$\begin{aligned} RNV_{[t,T]} &= \frac{1}{T-t} E_t^{\mathbb{Q}} \left(\int_t^T \sigma_s^2 ds \right) \\ &= \frac{2}{T-t} e^{r_d(T-t)} \left[\int_0^{F_t} \frac{1}{K^2} p(K, T) dK + \int_{F_t}^{\infty} \frac{1}{K^2} c(K, T) dK \right]. \end{aligned} \quad (2)$$

Equation (2) is also known as synthetic variance swap rate. Specifically, it is the rate that makes a variance swap which pays the difference between the future realized variance and the swap rate zero cost. The estimator in (2) has the advantage that it does not hinge on a particular option pricing model. Under the premise of absence of risk-aversion, the risk-neutral expectation is the best estimator of future realized variance. When risk aversion is present, $RNV_{[t,T]}$ encompasses both the expectation of realized variance under the physical measure and a risk premium. As proposed by Carr and Wu (2009), a simple means to quantify the variance risk premium is thus to compare the ex post realized variance with its risk-neutral forecast. In particular, the variance risk premium over the period from t to T is given by

$$RP_{[t,T]} = RV_{[t,T]} - RNV_{[t,T]}, \quad (3)$$

where $RV_{[t,T]}$ is the realized variance of the underlying spot rate. We use equation (3) to investigate the evolution of variance risk premiums in foreign exchange markets. To do so, both a sample estimator of the actual realized variance and a suitable discretization of the static replication of a variance swap are needed. Regarding realized variance, we consider two distinct estimators. In their seminal work on variance risk premiums in equity markets, both Bakshi and Kapadia (2003) and Carr and Wu (2009) employ daily data. Similarly, Low and Zhang (2005) employ daily spot data in the analysis of foreign exchange markets. This choice is reasonable also in the current context, since the majority of variance swaps settle against a daily fixing schedule. On the other hand, when spot prices exhibit a lot of intraday variation but tend to close

¹We refer to Appendix A for a derivation.

around the opening price, daily estimates of realized variance will be rather imprecise (Andersen and Bollerslev, 1998). Also, option traders can re hedge several times a day, which potentially materializes in higher implied volatilities and hence variance swap rates. Arguably, this intraday hedging opportunity should be accounted for when estimating realized variance. More appropriate estimates are then obtained from intraday data. Since foreign exchange markets are characterized by round-the-clock trading, a simple estimator based on 5-minute intraday mid-quotes is given by

$$\hat{RV}_{[t,T]} = \frac{260}{D} \sum_{i=1}^N \log \left(\frac{S_{i+1}}{S_i} \right)^2, \quad (4)$$

where N is the total number of observations over the interval from t to T and $D = N/288$ is the number of active trading days. We scale by 260 business days to obtain an annualized measure of realized variance.²

The use of 5-minute intraday data appears to have established as a standard for the analysis of currency markets.³ In fact, Pong et al. (2004) and Charoenwong et al. (2009) employ an identical estimator as in (4). As proposed in Andersen and Bollerslev (1997), we use log returns, and we dismiss any data between Friday and Sunday 2100 GMT. Moreover, we compute realized variance over a window from 2pm Eastern Time (ET) on trade date to 10am ET on expiry date. The first instance corresponds to the time stamp on the option quotes, whereas the latter is identical to New York cut, i.e. the time when the options expire.

Regarding a suitable estimator for the risk-neutral variance, we choose

$$R\hat{NV}_{[t,T]} = \frac{2}{(T-t)} e^{ra(T-t)} \sum_{i=1}^m \frac{\Delta K}{K_i^2} (p(K_i, T) I_{K_i \leq F_0} + c(K_i, T) I_{K_i > F_0}), \quad (5)$$

where I denotes an indicator function and m is the number of option quotes in the cross-section of strike prices. Specifically, we define a ± 6 standard deviation interval around the current forward price $F(t, T)$, where the standard deviation is based on the implied volatility of the delta-neutral strike quote. Furthermore, we set $m = 80$ with equally-sized subintervals.⁴

²The same formula is used to compute realized variance for daily spot data.

³See Andersen and Bollerslev (1997), Andersen and Bollerslev (1998), Pong, Shackleton, Taylor and Xu (2004) and Charoenwong, Jenwittayaroje and Low (2009) to name a few.

⁴The results are virtually indifferent to alternative choices of 4 respectively 8 standard deviations. Jiang and Tian (2005) provide an excellent account on truncation and other approximation errors in the context of model-free forecasting. We also tried out a number of different partitions. As it turns out, 80 subintervals is sufficient to obtain a fine enough grid. Note that since this equals a step size of 0.15 standard deviations, our findings are consistent with figure 2 from Jiang and Tian (2005).

The literature on foreign exchange markets provides some evidence that risk-neutral price processes exhibit jumps.⁵ Ultimately, this renders the SDE in (1) an inadequate description of the true return generating process. However, Jiang and Tian (2005) show that (2) still provides an unbiased estimator in a jump-diffusion setting where the jumps and the diffusion part are assumed to be orthogonal. In the most general case, Carr and Wu (2009) explicitly derive the approximation error due to jumps. Their numerical analysis suggests that the error is small. Finally, Todorov (2010) relies on the estimator in (2) even though he explicitly assumes a role for jumps.

3 Data and Methodology

We investigate the variance risk premiums for EURUSD (U.S. dollar per 1 euro), GBPUSD (U.S. dollar per 1 British pound), USDJPY (Japanese yen per 1 U.S. dollar) and EURGBP (British pound per 1 euro). Since the currencies involved form part of the group of majors, undesirable liquidity effects can largely be ruled out. Still, the different nature of the currency pairs and the fact that we include a so-called cross biproduct suggest that we deal with a fairly representative set of volatility shapes commonly observed in foreign exchange markets.⁶ We consider variance risk premiums over the 1-month and 3-month horizon. Our data set covers the period from January 2003 to August 2009. To avoid that any evidence is confound by the extreme market conditions of the financial crisis in autumn 2008, we first look at a pre-crisis subperiod from January 2003 to August 2008. The cut-off in August 2008 pre-dates the collapse of Lehman Brothers and the ensuing market turbulence by 2 weeks.

We obtain intraday spot data from Olsen & Associates, a currency trader and provider of high-frequency data. Interest rates are from Bloomberg. Volatility quotes on delta-neutral straddles (DN) and 5-delta, 10-delta and 25-delta call (DC) and put (DP) options have been provided by UBS, a major investment bank and market maker in foreign exchange. Care must be taken when these quotes are mapped to their respective delta. For EURUSD and GBPUSD, the domestic currency, i.e. USD, is the premium currency. As a result, a 25-delta quote refers to a regular spot delta of (-)0.25. In contrast, USDJPY and EURGBP are quoted with a foreign currency premium. Accordingly, a so-called premium-adjusted delta convention applies. For a detailed discussion, we refer to Reischich and Wystup (2010).

The difficulty with the estimator in (5) is that we only have a handful of delta quotes

⁵See for example Bates (1996), Daal and Madan (2005), Carr and Wu (2007) and references therein.

⁶We use the term biproduct synonym for currency pair.

available, which essentially precludes the use of a non-parametric interpolation method. We therefore apply two different techniques that are tailored to the specifics of foreign exchange markets. Malz (1997) proposes a parabolic interpolation method that rests on so-called risk reversal and butterfly strategies. Unfortunately, his approach relates to a forward delta and is valid only under a regular delta convention. In appendix B, we propose a generalized version that produces market consistent implied volatility functions under any convention. For brief, set $B = (\sigma^{DN}, \sigma^{xrr}(\Delta - \Delta_{DN}), \sigma^{xbf}(\Delta - \Delta_{DN})^2)'$, where σ^{xrr} and σ^{xbf} refer to a risk reversal and butterfly constructed from our volatility quotes,

$$\sigma^{xrr} := \sigma^{xDC} - \sigma^{xDP}, \quad \sigma^{xbf} := \frac{\sigma^{xDC} + \sigma^{xDP}}{2} - \sigma^{DN}. \quad (6)$$

\mathbf{x} denotes a particular choice of pivot options, e.g. the 25-delta calls and puts. The volatility of an arbitrary option with delta Δ follows from

$$\hat{\sigma}^\Delta = \mathbf{B}'\mathbf{a}, \quad (7)$$

where \mathbf{a} is a 3×1 vector of parameters that depends on the pivot choice and the delta convention. From $\hat{\sigma}^\Delta$, call and put prices to be used in (5) are obtained through the Garman and Kohlhagen (1983) function for currency options. Again, the mapping from the delta to the strike space requires that attention is paid to the quoting convention.

As an alternative to the Malz approach, we employ the vanna-volga method presented by Castagna and Mercurio (2007). The vanna-volga method maintains that any option in the strike domain can be replicated by a delta-hedge, a money market position and a suitable portfolio of three options. As such, the vanna-volga method rests on the same volatility quotes as the Malz approach. The reason for holding a portfolio of options is that under the premise of stochastic volatility, hedging an option with second-order accuracy requires offsetting the greeks $\partial c(\cdot)/\partial\sigma$, $\partial^2 c(\cdot)/\partial S \partial\sigma$ and $\partial^2 c(\cdot)/\partial\sigma^2$, which are commonly referred to as the vega, the vanna and the volga of an option.

For an arbitrary option with strike price K , suppose these so-called volatility greeks are stacked into a 3×1 vector \mathbf{y} . Furthermore, presume that the 3×3 matrix \mathbf{A} concatenates the volatility greeks for the 3 pivot options. Finally, define a 3×1 vector \mathbf{c}^{me} of market excess prices. More precisely, \mathbf{c}^{me} subsumes the differences between the observable market prices and the theoretical Black-Scholes prices of the options in the hedging portfolio. Then, by the usual replication arguments, it follows that the price of the option with exercise price K is given by

$$c(K, T, \sigma^K) = c(K, T, \sigma^{BS}) + (\mathbf{c}^{\text{me}})'\mathbf{x}, \quad (8)$$

where $\mathbf{x} = \mathbf{A}^{-1}\mathbf{y}$ is the 3×1 vector of portfolio weights. An array of option prices generated by the vanna-volga method can directly be employed in (5).⁷

The Malz and the vanna-volga method essentially differ in how they extrapolate the implied volatility function beyond the out-of-the-money (OTM) quotes. When the Malz approach is used, the volatility smile flattens out, whereas this is not the case for the vanna-volga method. Potentially, this has a non-trivial impact on the magnitude of the risk-neutral variance forecast, with the more aggressive estimates coming from the vanna-volga method. To deal with this problem, we consider a number of different estimates. Specifically, we produce four base estimates combining both methods with either a 10DP-DN-10DC or 25DP-DN-25DC pivot set. We leave out the 5-delta quotes to dispel any doubts regarding liquidity. Next, we employ a mixed estimation scheme where on each day, the approach is chosen which best fits the option quotes not used for interpolation. The 5-delta quotes thus serve as a benchmark in assessing tail-modelling accuracy. Figure 1 illustrates this principle.




Figure 1: Implied Volatility Functions

Implied volatility functions from the Malz (light grey) and the vanna-volga method (dark grey) generated alternatively from 25-delta (solid line) and 10-delta (dashed line) option quotes. The black crosses mark the observed 1-month option quotes for EURUSD on 2nd January 2003. The EURUSD spot rate on 2nd January 2003 was 1.0351.

The dark grey solid line marks the implied volatility function obtained when the vanna-volga method is fitted to the 25-delta option quotes. In contrast, the dashed dark line is obtained when the method is applied to the 10-delta options. The light grey lines mark the corresponding curves from the Malz approach. While the methods more or

⁷Note that the vanna of an option is again dependent on the delta convention.

less agree on the strike domain inside the pivot options used, they substantially differ in their prediction of deep OTM quotes. The black crosses represent the market quotes for 1-month EURUSD options on 2nd January 2003. In this instance, our mixed approach would have chosen to synthesize variance swaps from the 10-delta Malz approach.

To complement the methodological toolbox, we consider a sixth alternative that takes the information content of all quotes into account. In equation (8), substitute $\mathbf{x} = \mathbf{A}^{-1}\mathbf{y}$ and rearrange to obtain

$$c(K, T, \sigma^K) - c(K, T, \sigma^{BS}) = \mathbf{y}'\alpha, \quad (9)$$

where $\alpha = (\mathbf{A}^{-1})'\mathbf{c}^{\text{me}}$ is a 3×1 price vector associated with the volatility greeks. Seen from this perspective, the vanna-volga method purports that an option with exercise price K has a market price in excess of the Black-Scholes value that is determined by the option's vega, vanna and volga times the respective hedging costs. Using (9), we can estimate α from all quotes in the strike domain. Specifically, stack all observed excess prices into a (7×1) vector $\mathbf{c}^{\text{me},\forall}$ to obtain a system of equations

$$\mathbf{c}^{\text{me},\forall} = \mathbf{Y}'\alpha + \varepsilon, \quad (10)$$

where Y is a 3×7 matrix of volatility greeks with each column corresponding to one of the option quotes observed in the market. We dub the estimates resulting from (10) as least-square vanna-volga (LSVV) estimates.

4 Time Series Dynamics of Variance Risk Premiums

4.1 Evidence on the Sign and Size of Variance Risk Premiums

We conjectured that on average, variance risk premiums in foreign exchange markets ought to be negative. To evaluate this conjecture, consider an investor that maintains a variance swap investment with daily spot rate fixing. We ask for the annualized average amount earned over the period from January 2003 to August 2008, for a 1-month and 3-month investment horizon. This procedure is equivalent to investigating the sign and size of the average variance risk premium. Table 1 presents the results, along with t-statistics from a test on the null hypothesis that variance risk premiums are zero. Since the overlapping estimation procedure introduces serial correlation, Newey and West (1987) standard errors are employed. To address the possibility that variance

		EURUSD		GBPUSD		USDJPY		EURGBP	
		1m	3m	1m	3m	1m	3m	1m	3m
Malz									
25D	$\bar{R}P$	-0.0008	-0.0009	-0.0004	-0.0004	-0.0008	0.0000	-0.0007	-0.0006
	$tstat$	-2.882***	-2.109**	-1.261	-1.103	-1.790*	0.055	-4.757***	-2.270**
10D	$\bar{R}P$	-0.0010	-0.0011	-0.0006	-0.0006	-0.0011	-0.0004	-0.0008	-0.0007
	$tstat$	-3.612***	-2.689***	-1.722*	-1.537	-2.423**	-0.520	-5.179***	-2.678***
Vanna-Volga									
25D	$\bar{R}P$	-0.0013	-0.0015	-0.0009	-0.0010	-0.0023	-0.0017	-0.0010	-0.0010
	$tstat$	-4.917***	-3.953***	-2.469**	-2.300**	-4.192***	-2.685***	-5.776***	-3.673***
10D	$\bar{R}P$	-0.0014	-0.0015	-0.0009	-0.0010	-0.0019	-0.0013	-0.0010	-0.0010
	$tstat$	-4.875***	-3.901***	-2.518**	-2.318**	-3.647***	-1.809*	-5.809***	-3.474***
Mixed									
	$\bar{R}P$	-0.0013	-0.0015	-0.0009	-0.0009	-0.0013	-0.0006	-0.0010	-0.0009
	$tstat$	-4.788***	-3.776***	-2.512**	-2.269**	-3.166***	-0.948	-5.755***	-3.216***
LSVV									
	$\bar{R}P$	-0.0013	-0.0015	-0.0009	-0.0010	-0.0018	-0.0011	-0.0010	-0.0009
	$tstat$	-4.710***	-3.697***	-2.426**	-2.220**	-3.549***	-1.603	-5.722***	-3.344***

Table 1: Average Variance Risk Premiums from Daily Data

Annualized average variance risk premiums at the 1-month and 3-month horizon. Synthetic variance swaps are alternatively constructed using the Malz and vanna-volga method with 25-delta and 10-delta options. The mixed approach refers to the optimal choice from these base interpolation methods, while the LSVV method comprises the full information content of option quotes in the strike domain. The t-statistics are based on Newey and West (1987) standard errors with lags up to one year. *, ** and *** indicate significance at the 10%, 5% and 1% confidence level.

risk premiums are autocorrelated even for non-overlapping windows, we include lags up to one year.

In line with our expectations, variance risk premiums are uniformly negative. For example, an investment of USD 100 in a 1-month EURUSD variance swap would return an average annualized loss of about 13 cents.⁸ As hypothesized in the previous section, evidence of negative variance risk premiums is strongest for the vanna-volga method: With the exception of GBPUSD and 3-month USDJPY, the t-statistics are always significant on the 1% level. The mixed and LSVV methods, which arguably provide the most balanced assessment, suggest that except for the 3-month USDJPY estimates, variance risk premiums are significantly negative at least on the 5% level. Concerning USDJPY, we observe a material difference between the 1-month and 3-month estimates, irrespective of the chosen method. This finding is in line with a decreasing term structure of volatility risk premiums as reported in Low and Zhang (2005).

⁸This number is based on the estimates from the mixed or LSVV approach.

The continuously compounded returns to an investor going long the variance swap at rate $RNV_{[t,T]}$ are given by $\log(RV_{[t,T]}/RNV_{[t,T]})$. Following Carr and Wu (2009), we define this ratio as the log variance risk premium. A significantly negative log variance risk premium suggests that investors are willing to lose money on average to protect against a rise in variance. We test this hypothesis over the period from January 2003 to August 2008, again for all biproducts.

		EURUSD		GBPUSD		USDJPY		EURGBP	
		1m	3m	1m	3m	1m	3m	1m	3m
Malz									
25D	$L\bar{R}P$	-15.7%	-14.2%	-12.5%	-9.7%	-17.4%	-5.1%	-24.2%	-20.4%
	$tstat$	-3.839***	-2.386***	-3.335***	-2.339**	-3.490***	-0.664	-7.492***	-2.908***
10D	$L\bar{R}P$	-18.1%	-16.8%	-14.8%	-12.2%	-20.4%	-9.0%	-26.9%	-23.2%
	$tstat$	-4.397***	-2.873***	-3.850***	-2.885***	-3.893***	-1.179	-8.318***	-3.293***
Vanna-Volga									
25D	$L\bar{R}P$	-22.2%	-21.6%	-19.3%	-17.2%	-30.3%	-20.8%	-32.1%	-28.8%
	$tstat$	-5.401***	-3.858***	-4.511***	-3.728***	-5.604***	-2.967***	-10.09***	-4.242***
10D	$L\bar{R}P$	-22.3%	-21.7%	-19.3%	-17.2%	-27.4%	-17.3%	-32.0%	-28.7%
	$tstat$	-5.390***	-3.853***	-4.610***	-3.766**	-4.874***	-2.260***	-9.944***	-4.085***
Mixed									
	$L\bar{R}P$	-22.0%	-21.3%	-18.7%	-16.3%	-22.3%	-11.6%	-31.3%	-27.9%
	$tstat$	-5.318***	-3.745***	-4.695***	-3.790***	-4.838***	-1.649*	-9.654***	-3.852***
LSVV									
	$L\bar{R}P$	-21.7%	-20.9%	-18.6%	-16.5%	-26.1%	-15.8%	-31.2%	-27.8%
	$tstat$	-5.260***	-3.692***	-4.541***	-3.669**	-4.862***	-2.113**	-9.688***	-3.951***

Table 2: Average Log Variance Risk Premiums from Daily Data

Annualized average log variance risk premiums at the 1-month and 3-month horizon. Synthetic variance swaps are alternatively constructed using the Malz and vanna-volga method with 25-delta and 10-delta options. The mixed approach refers to the optimal choice from these base interpolation methods, while the LSVV method comprises the full information content of option quotes in the strike domain. The t-statistics are based on Newey and West (1987) standard errors with lags up to one year.

*, ** and *** indicate significance at the 10%, 5% and 1% confidence level.

Table 2 shows annualized average log variance risk premiums. The numbers are considerably larger than in Table 1, which is a result of using the variance swap rate as the investment basis, rather than a fixed USD amount. The absolute magnitudes of the premiums are comparable to those reported in Carr and Wu (2009) for individual stocks, but with -10% to -30% generally smaller than for stock indices. The results in Table 2 broadly confirm our findings from Table 1. Except for the 3-month USD-JPY estimates, the t-statistics suggest that the negative log variance risk premiums are highly significant across different methods. To account for potentially long-lasting

autocorrelations, the t-statistics are again based on Newey and West (1987) robust standard errors with a one-year lag. For USDJPY, 3-month estimates are highly significant when the vanna-volga method is employed, and still moderately significant when either the LSVV or the mixed approach is used. The difference of about 10% between 1-month and 3-month risk premiums is striking however. In contrast, this difference amounts to just 1% to 3% for EURUSD, GBPUSD and EURGBP.

4.2 Variance Risk Premiums from Intraday Data

Alternative estimates of variance risk premiums are obtained by computing realized variance from 5-minute intraday data. When currency pairs exhibit a lot of intraday variation, we can expect variance risk premiums to become more positive. Thus, some of the previously reported negative risk premiums may become insignificant. Regarding the computation of variance swap rates, we stick to the common set of methodologies outlined in section 3. Table 3 shows the results.

		EURUSD		GBPUSD		USDJPY		EURGBP	
		1m	3m	1m	3m	1m	3m	1m	3m
Malz									
25D	$\bar{R}P$	0.0001	-0.0001	0.0002	0.0001	0.0006	0.0013	0.0017	0.0017
	$tstat$	0.293	-0.404	0.949	0.182	1.840*	2.329**	4.176***	4.384***
10D	$\bar{R}P$	-0.0001	-0.0004	0.0000	-0.0001	0.0002	0.0009	0.0016	0.0016
	$tstat$	-0.460	-1.064	0.208	-0.400	0.696	1.647*	3.928***	4.121***
Vanna-Volga									
25D	$\bar{R}P$	-0.0005	-0.0008	-0.0003	-0.0005	-0.0009	-0.0004	0.0013	0.0013
	$tstat$	-1.702*	-2.372**	-1.135	-1.506	-1.750*	-0.905	3.218***	3.461***
10D	$\bar{R}P$	-0.0005	-0.0008	-0.0003	-0.0005	-0.0005	0.0000	0.0014	0.0013
	$tstat$	-1.708*	-2.336**	-1.126	-1.500	-1.183	-0.045	3.348***	3.523***
Mixed									
	$\bar{R}P$	-0.0005	-0.0007	-0.0002	-0.0004	0.0000	0.0006	0.0014	0.0014
	$tstat$	-1.622	-2.221**	-0.977	-1.323	0.143	1.218	3.522***	3.675***
LSVV									
	$\bar{R}P$	-0.0004	-0.0007	-0.0002	-0.0004	-0.0004	0.0001	0.0014	0.0014
	$tstat$	-1.555	-2.142**	-0.957	-1.347	-0.901	0.280	3.445***	3.626***

Table 3: Average Variance Risk Premiums from Intraday Data

Annualized average variance risk premiums at the 1-month and 3-month horizon. Synthetic variance swaps are alternatively constructed using the Malz and vanna-volga method with 25-delta and 10-delta options. The mixed approach refers to the optimal choice from these base interpolation methods, while the LSVV method comprises the full information content of option quotes in the strike domain. The t-statistics are based on Newey and West (1987) standard errors with lags up to one year. *, ** and *** indicate significance at the 10%, 5% and 1% confidence level.

The variance risk premiums computed from intraday data are considerably higher

than those reported for daily data. For EURUSD and GBPUSD, variance risk premiums have increased by about 9 and 6 cents per 100 USD. For USDJPY and EURGBP, the increase is even more pronounced with 14 and 24 cents respectively. For all currency pairs, both the 1-month and 3-month premiums have increased by roughly the same amount. Variance risk premiums are still negative for EURUSD and GBPUSD. They are positive for EURGBP and 3-month USDJPY and either positive or negative, depending on the chosen interpolation method, for 1-month USDJPY. The previously strong evidence for negative risk premiums in foreign exchange markets has mostly evaporated. 3-month EURUSD estimates are significantly negative on a 5% level for all except the Malz approach. Apart from this outlier, the 5-minute data provides evidence against the presence of a variance risk premium. In fact, we now report significantly positive risk premiums for EURGBP, both at the 1-month and 3-month horizon.

		EURUSD		GBPUSD		USDJPY		EURGBP	
		1m	3m	1m	3m	1m	3m	1m	3m
Malz									
25D	$L\bar{R}P$	1.2%	-1.1%	2.7%	1.0%	5.9%	12.9%	37.8%	36.5%
	$tstat$	0.417	-0.284	0.978	0.237	1.906*	2.471**	5.629***	5.944***
10D	$L\bar{R}P$	-1.1%	-3.7%	0.4%	-1.6%	2.9%	9.0%	35.1%	33.6%
	$tstat$	-0.375	-0.961	0.142	-0.393	0.859	1.717*	5.283***	5.445***
Vanna-Volga									
25D	$L\bar{R}P$	-5.2%	-8.5%	-4.1%	-6.6%	-7.1%	-2.8%	29.9%	28.1%
	$tstat$	-1.746*	-2.376**	-1.375	-1.625	-1.842*	-0.617	4.424***	4.543***
10D	$L\bar{R}P$	-5.4%	-8.6%	-4.1%	-6.6%	-4.1%	0.6%	30.0%	28.2%
	$tstat$	-1.768*	-2.388**	-1.384	-1.629	-1.094	0.121	4.544***	4.522***
Mixed									
	$L\bar{R}P$	-5.1%	-8.2%	-3.4%	-5.7%	0.9%	6.4%	30.7%	29.0%
	$tstat$	-1.668*	-2.242**	-1.233	-1.423	0.315	1.322	4.750***	4.673***
LSVV									
	$L\bar{R}P$	-4.8%	-7.8%	-3.4%	-5.9%	-2.8%	2.2%	30.8%	29.0%
	$tstat$	-1.588	-2.150**	-1.186	-1.450	-0.784	0.426	4.660***	4.664***

Table 4: Average Log Variance Risk Premiums from Intraday Data

Annualized average log variance risk premiums at the 1-month and 3-month horizon. Synthetic variance swaps are alternatively constructed using the Malz and vanna-volga method with 25-delta and 10-delta options. The mixed approach refers to the optimal choice from these base interpolation methods, while the LSVV method comprises the full information content of option quotes in the strike domain. The t-statistics are based on Newey and West (1987) standard errors with lags up to one year.

*, ** and *** indicate significance at the 10%, 5% and 1% confidence level.

Similar conclusions can be drawn from the results shown in Table 4. For EURUSD

and GBPUSD, log variance risk premiums are roughly 11% to 17% higher compared to the estimates based on daily data. For USDJPY, the increase is more pronounced with approximately 23% at the 1-month and 18% at the 3-month horizon. The discrepancy, which is of the order of 60%, is again most dramatic for EURGBP. For all except the Malz approach, EURUSD 1-month and 3-month estimates are significantly negative on the 10% or 5% level. EURGBP premiums are positive and highly significant across all methods. The evidence on GBPUSD and USDJPY is mixed. The former estimates are mostly negative but insignificant, while USDJPY results tend to be positive, in particular at the 3-month horizon.

4.3 Realized Variance at Alternative Frequencies

The theoretical arguments developed in Andersen, Bollerslev, Diebold and Labys (2001, 2003) suggest that the precision of realized variance estimates increases with the frequency of the underlying data. With an ever increasing frequency however, market microstructure effects such as bid-ask bounce, price discreteness, spread positioning or strategic order flow come into play, which effectively results in a bias-efficiency trade off (Andersen and Benzoni, 2008). The optimal sampling frequency is unknown and depends on the asset under scrutiny. Andersen and Benzoni (2008) suggest that for liquid stocks, realized volatility estimates stabilize at frequencies between 5 and 40 minutes. Given the high activity in currency markets, one is tempted to conclude that estimates based on 5-minute intervals should be unbiased. To shed light on this issue, Table 5 shows realized volatilities at 6 different sampling frequencies between 5 minutes and 1 day. The numbers in Table 5 represent averages of annualized estimates that are taken over intervals corresponding to the variance swap maturities.

Two key findings emerge. First, realized volatilities computed from 5-minute intraday data are materially higher than the corresponding estimates from daily spot data. For example, the average EURUSD realized volatility at the 1-month horizon is 9.1% for 5-minute intraday data and gradually decreases to 8.5% for daily data. For EURUSD and GBPUSD, the difference between 5-minute and 1-day estimates is between 0.3% and 0.6%. For USDJPY, it amounts to 0.8% to 0.9%. Not surprisingly, the largest differences are observed for EURGBP. Second, realized volatilities appear to stabilize at 30-minute or 1-hour frequencies for EURUSD, GBPUSD and USDJPY, which is in line with the findings in Andersen et al. (2003). In contrast, EURGBP estimates continue to fall as the frequency is decreased.

To gain additional insight on the discrepancy between the various estimates, we

	EURUSD		GBPUSD		USDJPY		EURGBP	
	1m	3m	1m	3m	1m	3m	1m	3m
5min	9.1%	9.1%	8.4%	8.4%	10.0%	10.1%	7.6%	7.6%
10min	8.8%	8.9%	8.2%	8.3%	9.7%	9.8%	7.0%	7.1%
30min	8.6%	8.6%	8.1%	8.1%	9.4%	9.5%	6.5%	6.6%
1hour	8.6%	8.6%	8.1%	8.1%	9.2%	9.4%	6.3%	6.4%
2hours	8.6%	8.6%	8.0%	8.1%	9.1%	9.3%	6.2%	6.2%
1day	8.5%	8.6%	7.9%	8.1%	9.1%	9.3%	5.7%	5.9%

Table 5: Average Realized Volatility at Different Sampling Frequencies

Average realized volatility for sampling frequencies between 5 minutes and 1 day. The numbers represent averages of annualized volatilities that are computed over the tenor of the corresponding 1-month and 3-month variance swaps. The horizon under consideration ranges from January 2003 to August 2008.

compute autocorrelation functions.⁹ The first lag autocorrelation for EURUSD and GBPUSD 5-minute data is of the order of -0.03 and highly significant. For USDJPY and EURGBP, we obtain -0.06 and -0.11. These numbers gradually decrease and become insignificant at the 30-minute interval, except for EURGBP. Although serial correlation is somewhat closer to zero if we only consider day time trading between 8am and 5pm GMT, the general findings still apply. For EURGBP, our analysis suggests that a 2-hour frequency is an appropriate choice. Interestingly, the first lag autocorrelation spikes to 0.09 for daily data, which explains why the realized volatility estimates from Table 5 continue to fall as the sampling frequency decreases.

Figure 2 shows the evolution of variance risk premiums as the sampling frequency decreases. Variance swap rates have been constructed from the mixed approach. The upper-left panel reveals 1-month risk premiums for EURUSD (black solid line), GBPUSD (grey solid), USDJPY (black dashed) and EURGBP (grey dashed). For sampling frequencies lower than or equal to 30 minutes, all currency pairs exhibit a negative risk premium. The same finding applies for the 3-month horizon (upper-right panel). In general, the most negative premiums are obtained for EURUSD, followed by USDJPY for the 1-month and GBPUSD for the 3-month horizon. In accordance with Table 5, variance risk premiums tend to stabilize for sampling frequencies of 30 minutes or lower. The lower panels plot the corresponding t-statistics. 1-month EURUSD risk premiums become highly significant for any sampling frequency larger than 5 minutes. Except for EURGBP, highly negative risk premiums are obtained for frequencies equal

⁹To conserve space, we refrain from revealing the full-blown analysis. However, autocorrelation plots and result-tables are available upon request.

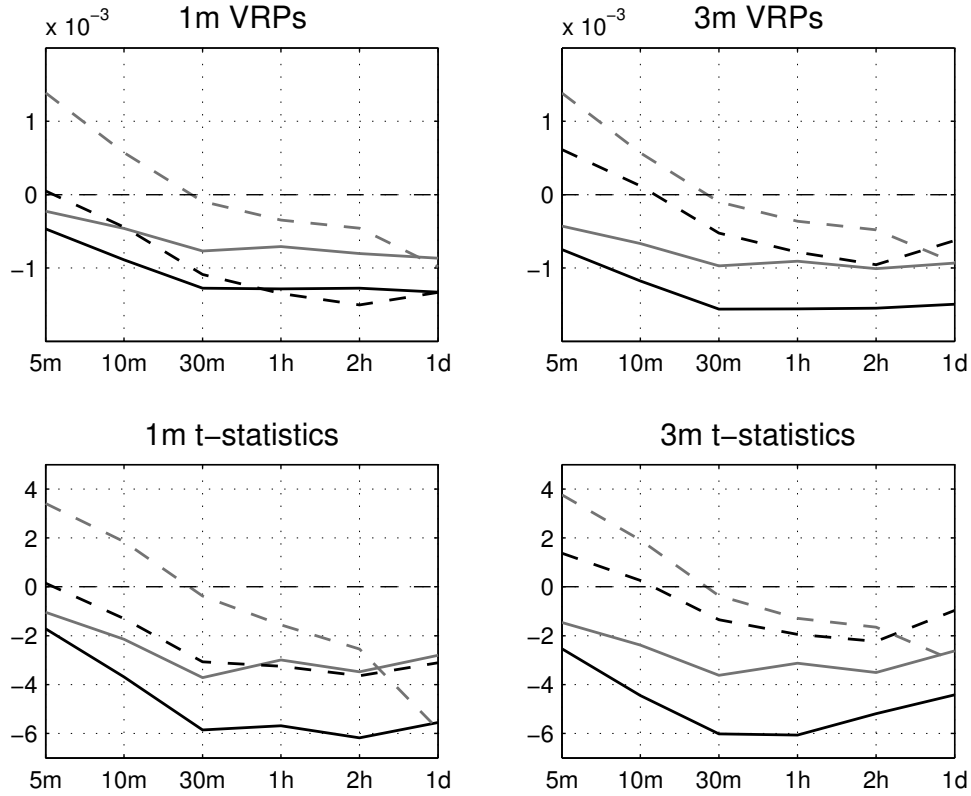


Figure 2: Variance Risk Premiums at Alternative Frequencies

Variance risk premiums (VRPs) with realized variance computed at alternative sampling frequencies. EURUSD (black solid), GBPUSD (grey solid), USDJPY (black dashed) and EURGBP (grey dashed) premiums are computed using the mixed approach. The lower panels show the t-statistics as the sampling frequency decreases.

to or lower than 30 minutes. EURGBP estimates are significantly negative on the 5% level for the 2-hour interval data. For the 3-month horizon (lower-right panel), we report highly significant premiums for EURUSD and GBPUSD. The estimates for USDJPY are significant on a 10% level for a 30-minute frequency, 5% for the 1-hour and 2-hour frequency and insignificant when daily data is used. The latter finding can be attributed to a negative, albeit insignificant autocorrelation of -0.03. For EURGBP, a moderately significant risk-premium is obtained for the 2-hour interval data.

In sum, we conclude that the intraday estimates reported in Table 3 and Table 4 may be confounded by microstructure effects. The results in this section suggest that variance risk premiums in foreign exchange markets are significantly negative, provided realized variance is computed using data with an appropriate sampling frequency. One

upshot is that some evidence from the previous literature against the presence of a volatility risk premium needs to be revisited, since the use of 5-minute intraday data has been widespread. For the remainder of this paper, we work with variance risk premiums computed from 30-minute intraday data for EURUSD, GBPUSD and USDJPY, and 2-hour data for EURGBP.

4.4 Are Variance Risk Premiums Time-varying?

In this section, we examine whether variance risk premiums in currency markets are time-varying. First, we plot variance risk premiums over the time horizon from January 2003 to August 2008. Second and in analogy to Carr and Wu (2009), we run so-called expectation hypothesis regressions. These regressions not only reveal whether variance risk premiums are time-varying, but also whether they are systematically linked to the risk-neutral expectation of future variance.

Figure 3 shows the variance risk premiums at the 1-month (dark grey crosses) and 3-month horizon (light grey crosses). Variance swaps have been synthesized using the mixed approach. The realized variance is based on 30-minute data for EURUSD, GBPUSD and USDJPY and 2-hour data for EURGBP. To facilitate interpretation, we have superimposed two-sided Gaussian kernel estimates with a standard deviation of 1/4 year. The solid and dashed line correspond to smoothed estimates of the 1-month and 3-month variance risk premiums. For all currency pairs, variance risk premiums are highly time-varying and oscillate around or just below the zero line. Despite the negative average premiums reported in the previous sections, a considerable amount of the realizations lie in the positive domain, as is seen from the scattering of the crosses. The largest swings in variance risk premiums are observed for USDJPY (note that it plots on a different scale). The smoothed estimates are negative throughout the entire period for both EURUSD and GBPUSD. They reveal a pronounced slump in the second half of 2004 which is not shared by USDJPY and EURGBP. The 3-month smoothed curves tend to lie below the 1-month estimates, although this relationship is inverted as the financial crisis is approached. In general, the outset of the crisis marks a shift in investor sentiment which is reflected in increasingly negative variance risk premiums across the different biproducts.

To formally confirm the notion of time-varying variance risk premiums, we run the following expectation hypothesis regression

$$RV_t = \alpha + \beta RNV_t + \varepsilon_t. \tag{11}$$

The time-varying component of the variance risk premium, if present, is captured by

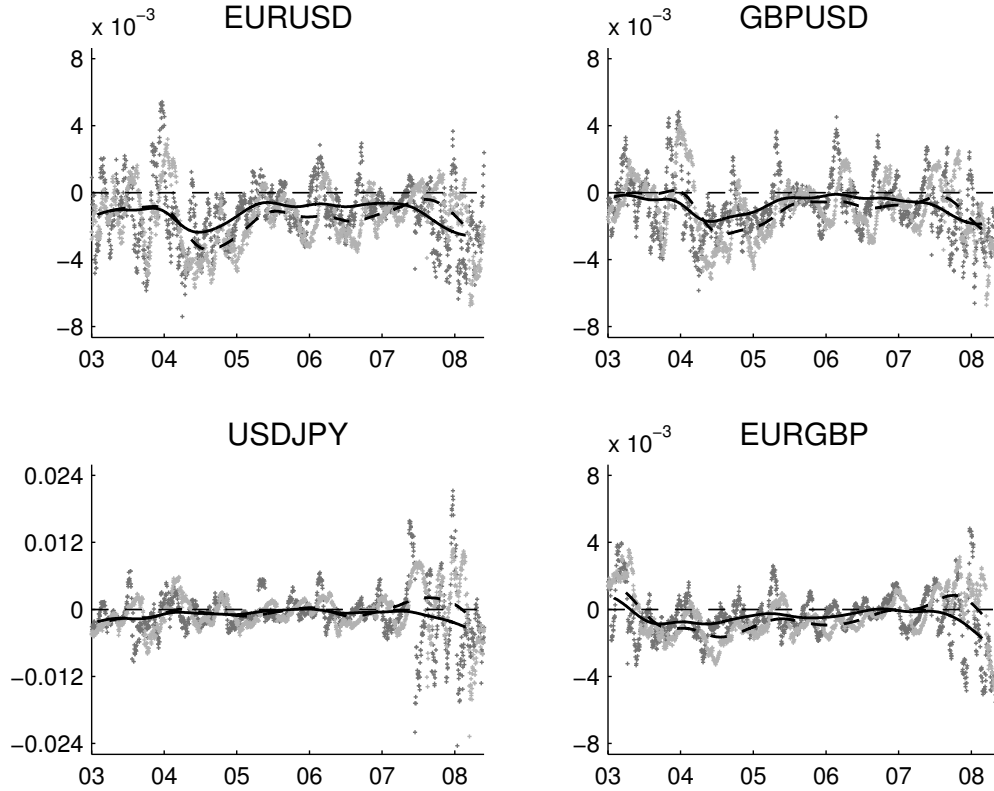


Figure 3: Time-Variation in Variance Risk Premiums

Variance risk premiums generated from the mixed approach and realized variances that are based on sampling frequencies of 30 minutes for EURUSD, GBPUSD, USDJPY and 2 hours for EURGBP. The dark (light) grey crosses show the premiums at the 1-month (3-month) horizon. The solid (dashed) black line depicts a Gaussian kernel estimate with standard deviation of 1/4 year.

the error term ε . Under the null hypothesis of a zero variance risk premium, $\alpha = 0$ and $\beta = 1$. If a variance risk premium is present but constant, we can expect $\alpha \neq 0$ but β still equal to 1. Finally, if the variance risk-premium is time-varying and systematically linked to the variance swap rate RNV_t , it generally holds that $\beta \neq 1$.

We run (11) for both the 1-month and 3-month horizons. The left-hand side is computed using a 2-hour sampling frequency for EURGBP and a 30-minute frequency for the remainder. Variance swaps have been synthesized from the mixed approach. Estimates are obtained using Hansen's (1982) GMM. To account for serial dependence, we use Newey-West standard errors with a lag of one year. The results are shown in Table 6.

The slope parameters in Table 6 are considerably smaller than one, irrespective of the currency pair and time horizon under examination. In all instances, the reported

	EURUSD		GBPUSD		USDJPY		EURGBP	
	1m	3m	1m	3m	1m	3m	1m	3m
α	0.0003	0.0004	0.0011	0.0020	0.0031	0.0040	0.0005	0.0004
$tstat$	0.683	0.529	1.625	2.255**	4.494***	4.807***	1.478	1.010
β	0.820	0.793	0.755	0.613	0.594	0.536	0.801	0.801
$tstat$	-2.495**	-2.269**	-2.411**	-2.981***	-7.449***	-6.869***	-2.269**	-2.102**

Table 6: Expectation Hypothesis Regression

Expectation hypothesis regressions, where risk-neutral variance is computed from the mixed approach. Realized variances are based on sampling frequencies of 30 minutes for EURUSD, GBPUSD, USDJPY and 2 hours for EURGBP. The t-statistics are based on Newey and West (1987) standard errors with lags up to one year. *, ** and *** indicate significance at the 10%, 5% and 1% confidence level.

t-statistics indicate significance at least on the 5% level, which suggests a strong relationship between the premiums charged for loading variance risk and the risk-neutral expectation of future variance. Specifically, market participants are willing to pay a larger premium when they expect a higher variance for the nearby future. On the other hand, the larger the negative correlation between the time-varying component of the variance risk premium and the future expected variance, the larger the constant part of the risk premium. This finding is deduced from the significantly positive α s for USDJPY and 3-month GBPUSD and the simultaneously small slope parameters.

5 How Does the Financial Crisis of 2008 Affect the Results?

So far we have disregarded data after August 2008. Common sense suggests that with the rise of the financial crisis, the average market participant became more risk averse and hence put a larger premium on accepting variance risk. To see whether this conjecture holds true, we consider variance risk premiums for a subperiod from January 2008 to August 2009. Since this period is relatively short, we confine ourselves to a graphical analysis. Figure 4 reveals analogous plots to Figure 3, with the light and dark grey marks corresponding to realizations of variance risk premiums for the 1-month and 3-month horizon. The solid and dashed black lines depict Gaussian kernel estimates with a 1-month standard deviation.

The absolute magnitude of variance risk premiums during the financial crisis has substantially increased. Measured on the larger scale of Figure 4, the previously reported dips into negative territory in Summer 2008 are hardly visible. At the onset

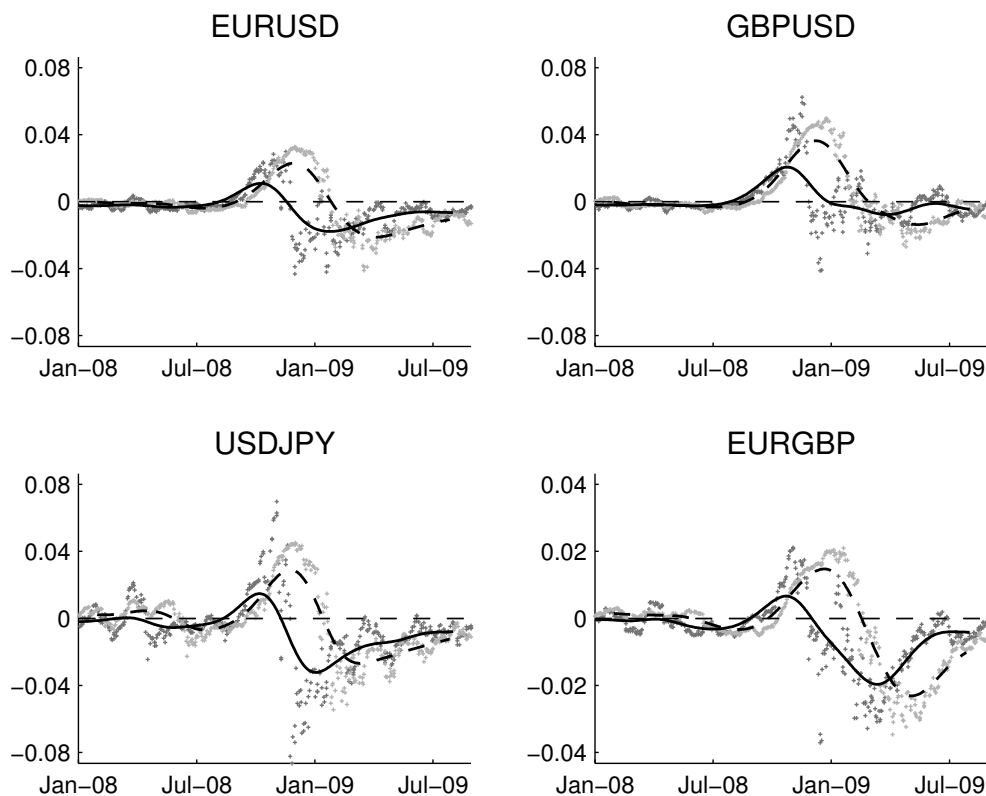


Figure 4: Variance Risk Premiums During the Financial Crisis

Variance risk premiums generated from the mixed approach and realized variances that are based on sampling frequencies of 30 minutes for EURUSD, GBPUSD, USDJPY and 2 hours for EURGBP. The dark (light) grey crosses show the premiums at the 1-month (3-month) horizon. The solid (dashed) black line depicts a Gaussian kernel estimate with standard deviation of 1 month.

of the crisis, realized variances have exceeded the swap rates by far, resulting in large positive premiums from about August 2008 to November 2008 for the 1-month and February 2009 for the 3-month horizon. This observation can likely be attributed to an option market which did not fully anticipate the increase in variance that followed the bankruptcy of Lehman Brothers. Since the computation of variance risk premiums is based on swap rates that are fixed for the tenor of the option contracts, the 1-month risk premiums turn negative considerably before the 3-month premiums. Eventually, both premiums reside in the negative domain, which suggests that the market reacted strongly to the crisis by demanding a large premium for compensating variance risk.

To assess the robustness of the result from Section 4 in light of the financial crisis, we compute variance risk premiums over the full sample period from January 2003 to

August 2009. To conserve space, we report variance and log variance risk premiums only for the mixed approach. Realized variance is based on data from 3 alternative sampling frequencies. Table 7 presents the results.

		EURUSD		GBPUSD		USDJPY		EURGBP	
		1m	3m	1m	3m	1m	3m	1m	3m
VRP									
1day	$\bar{R}P$	-0.0023	-0.0019	-0.0009	-0.0001	-0.0035	-0.0020	-0.0018	-0.0015
	$tstat$	-2.468**	-1.924*	-1.852*	-0.104	-1.806*	-1.208	-2.257**	-1.699*
2h/30min	$\bar{R}P$	-0.0020	-0.0018	-0.0004	0.0001	-0.0025	-0.0012	-0.0013	-0.0011
	$tstat$	-2.526**	-1.880*	-0.672	0.093	-1.759*	-0.932	-1.530	-1.176
5min	$\bar{R}P$	-0.0013	-0.0011	0.0002	0.0007	-0.0010	0.0003	0.0009	0.0011
	$tstat$	-1.492	-1.158	0.404	0.577	-0.838	0.214	1.196	1.260
LRP									
1day	$L\bar{R}P$	-23.1%	-19.0%	-16.5%	-10.4%	-26.0%	-13.3%	-31.2%	-25.9%
	$tstat$	-5.518***	-2.977***	-4.495***	-1.746*	-4.682***	-1.816*	-9.767***	-3.465***
2h/30min	$L\bar{R}P$	-16.5%	-17.0%	-9.3%	-7.7%	-15.0%	-7.2%	-13.4%	-12.6%
	$tstat$	-6.930***	-3.631***	-3.115***	-1.314	-3.853***	-1.318	-4.190***	-1.890*
5min	$L\bar{R}P$	-6.5%	-7.3%	-1.4%	-0.2%	-2.1%	5.1%	25.8%	25.9%
	$tstat$	-2.057**	-1.711*	-0.536	-0.036	-0.492	0.916	3.531***	3.881***

Table 7: Variance Risk Premiums for the Full Sample

Annualized average variance and log variance risk premiums for the period from January 2003 to August 2009. Synthetic variance swaps are constructed using the mixed approach. The sampling frequency for realized variance is indicated in the very left column. The t-statistics are based on Newey and West (1987) standard errors with lags up to one year. *, ** and *** indicate significance at the 10%, 5% and 1% confidence level.

In general, variance risk premiums are more negative than the previously reported premiums for the truncated period, in particular for USDJPY. The exception to this finding are the premiums for GBPUSD. Figure 4 suggests that in 2009, variance risk premiums for GBPUSD did not drop as much as for the other pairs. The results in Table 7 tend to be less significant than before, owing to a material increase in the variation of variance risk premiums since 2008. For the daily sampling frequency, evidence of negative risk premiums is still strong, with 6 out of 8 reported premiums being significant at least on the 10% confidence level. For the 30-minute data, the EURUSD and 1-month USDJPY premiums remain significantly negative. None of the premiums are significant when realized variance is computed using 5-minute interval data. The log variance risk premiums reported in the lower half of Table 7 are of the same order as previously for the truncated sample. On the basis of 30-minute (2-hour for EURGBP) sampled data, log variance risk premiums are significantly negative except for 3-month GBPUSD and USDJPY. In sum, extending the data set to include the financial crisis does not fundamentally alter the results. To avoid any issues with

possibly non-stationary time series, we continue to work with the truncated sample.

6 What Drives Variance Risk Premiums in Foreign Exchange Markets?

6.1 S&P 500, the VIX and the TED Spread

Variance risk premiums in currency markets can result from co-movement between exchange rate variance and classic risk factors. As such, they might simply be a reflection of other well-documented risk premiums. Alternatively, they may be attributable to an independent variance risk factor which commands a premium. In the introduction, we emphasized that the first explanation for variance risk premiums in currency markets is less compelling. In here, we test the two competing arguments by hypothesizing a relationship between log variance risk premiums and classic risk factors. If the risk factors can account for the returns on the variance swaps, the notion of an independent variance risk factor can be rejected.

We start by looking at the excess return on the stock market, proxied by the S&P 500, over the T-bill rate. We augment this CAPM-like regression with the return on the VIX and the first difference of the 3-month TED spread series. The VIX reveals whether variance risk premiums in currency markets are attributable to the overall level of market volatility. On the other hand, the TED spread, i.e. the difference between the U.S. Libor and the T-Bill rate, provides information with regard to the extent to which credit risk is reflected in variance risk premiums. Also, the TED spread is a robust measure of overall risk sentiment in financial markets.

We estimate the regression

$$LRP_{j,t} = \alpha_j + \beta_j^1 RE_t^{SPX} + \beta_j^2 R_t^{VIX} + \beta_j^3 \Delta_t^{TED} + \varepsilon_{j,t}. \quad (12)$$

by means of GMM. The log variance risk premiums are derived using the mixed approach. Realized variance is computed on the basis of 30-minute interval data for EURUSD, GBPUSD and USDJPY and 2-hour data for EURGBP. T-statistics for the coefficients are based on Newey and West (1987) robust standard errors with a one-year lag. Table 8 reveals the results.

The first row in Table 8 shows the unconditional log variance risk premiums. The second row provides the risk-adjusted premiums. The two are approximately of the same magnitude. Table 8 provides little evidence that the S&P 500 excess returns are linked to variance risk premiums in currency markets. β^1 is modestly significant only

	EURUSD		GBPUSD		USDJPY		EURGBP	
	1m	3m	1m	3m	1m	3m	1m	3m
$L\bar{R}P$	-0.165	-0.194	-0.122	-0.141	-0.127	-0.065	-0.123	-0.131
α	-0.171	-0.197	-0.129	-0.150	-0.134	-0.064	-0.127	-0.137
$tstat$	-8.177***	-6.128***	-5.184***	-4.218***	-5.079***	-2.307**	-4.455***	-1.992**
β^1	1.474	-0.350	1.799	0.411	0.837	-1.457	1.649	-0.423
$tstat$	1.800*	-0.414	1.698*	0.379	0.537	-1.266	1.191	-0.278
β^2	0.475	0.007	0.470	0.182	0.742	0.156	0.380	-0.156
$tstat$	3.366***	0.040	2.911***	0.751	3.493***	0.793	1.775*	-0.625
β^3	0.154	0.173	0.158	0.167	0.340	0.441	0.000	0.204
$tstat$	7.817***	3.975***	2.597***	2.10**	3.658***	3.697***	-0.007	1.675*
R^2	12.9%	6.9%	9.0%	6.9%	21.4%	24.5%	2.1%	3.0%

Table 8: Classic Risk Factor Regression

Risk-adjusted average log variance risk premiums (α) and factor loadings for S&P 500 returns, returns on the VIX and changes in the TED spread. The t-statistics are based on Newey and West (1987) standard errors with lags up to one year. *, ** and *** indicate significance at the 10%, 5% and 1% confidence level.

for 1-month EURUSD and GBPUSD. In contrast, the results suggest that the VIX returns are strongly related to all of the 1-month risk premiums. An increase in the VIX tends to be accompanied by an increase in realized variance. Since variance swap rates are fixed on the trade date of the underlying option contracts, the sign on the VIX is positive, except for 3-month EURGBP. The TED spread is also correlated with the returns on the variance swaps. An increase in the spread is generally associated with higher realized variance, leading to smaller variance risk premiums. Significant coefficients on the TED spread are obtained in all cases except for 1-month EURGBP. With an R^2 above 20%, the strongest relationship between variance risk premiums and the risk factors is observed for USDJPY. In general however, the results suggest that variance risk premiums cannot be accounted for by classic risk factors. In particular, the conditional variance risk premiums remain significantly negative across all currency pairs and tenors. Table 8 thus provides evidence in favor of an independent variance risk factor for currency markets.

6.2 Do Variance Risk Premiums Subsume Fear of Jump Risk?

In the literature on option markets, it has been widely documented that apart from variance risk, jump risk commands a premium too (see e.g. Jackwerth, 2000 or Pan, 2002). Bakshi and Kapadia (2003) conjecture that variance risk premiums may subsume the fear of jump risk. To test this hypothesis, they regress gains from delta-hedged option positions on the model-free metrics of Bakshi et al. (2003). Implicitly, they assume that

jump risk can be surrogated by the skew and kurtosis of the risk-neutral distribution. In our case, constructing implied skew and kurtosis along their lines might result in noisy regressors, since the volatility data is delimited by an approximate strike range of ± 1.65 standard deviations.¹⁰ For this reason, we proxy the mean jump size and jump intensity directly by the observed mid-quotes for 25-delta risk reversal and butterfly strategies. Specifically, we set up the regression equation

$$LRP_{j,t} = \alpha_j + \beta_j^1 \Delta_t^{25RR} + \beta_j^2 \Delta_t^{25BF} + \varepsilon_{j,t}. \quad (13)$$

Although risk reversals and butterflies are bounded economically, they are non-stationary in a statistical sense. We thus consider differences over the tenor of the variance swaps rather than levels. As before, log variance risk premiums are constructed from the mixed approach. Realized variance estimates are based on a 30-minute sampling frequency for EURUSD, GBPUSD and USDJPY and a 2-hour frequency for EURGBP. T-statistics are computed from Newey and West (1987) robust standard errors with a one-year lag. Table 9 shows the results.¹¹

	EURUSD		GBPUSD		USDJPY		EURGBP	
	1m	3m	1m	3m	1m	3m	1m	3m
$L\bar{R}P$	-0.165	-0.194	-0.122	-0.141	-0.127	-0.065	-0.123	-0.131
α	-0.162	-0.198	-0.124	-0.148	-0.131	-0.081	-0.125	-0.144
$tstat$	-7.021***	-5.914***	-4.313***	-4.093***	-4.675***	-1.898*	-4.147***	-3.056***
β^1	-0.030	-0.072	-0.081	-0.045	-0.158	-0.133	0.134	0.037
$tstat$	-0.789	-1.110	-1.150	-0.449	-5.907***	-2.538**	1.354	0.239
β^2	2.833	1.294	3.587	2.519	1.314	0.905	7.268	5.036
$tstat$	3.552***	2.794***	2.390**	5.192***	2.898***	1.915*	5.476***	4.510***
R^2	7.3%	7.6%	10.6%	18.2%	40.1%	33.3%	12.7%	21.0%

Table 9: Jump Risk Factor Regression

Risk-adjusted average log variance risk premiums (α) and factor loadings for changes in 25-delta risk reversals and butterflies. The t-statistics are based on Newey and West (1987) standard errors with lags up to one year. *, ** and *** indicate significance at the 10%, 5% and 1% confidence level.

The conditional variance risk premiums from Table 9 are of the same size as the unconditional premiums, and with the exception of 3-month USDJPY, they are highly significant despite the inclusion of the jump fear proxies. Provided fear of jump risk is well reflected in risk reversal and butterfly strategies, we can conclude that variance risk premiums do not subsume jump risk aversion, which is consistent with the findings in

¹⁰We thank an anonymous referee for pointing this out.

¹¹Alternatively, we have estimated (13) using 10-delta and 5-delta option quotes. The results are literally the same and therefore omitted.

Low and Zhang (2005). As in the previous section, the results thus provide evidence in favor of an independent, possibly currency-specific variance risk factor which commands a premium. However, variance risk premiums are not independent of the jump risk factors. For USDJPY, the size of the log variance risk premium is significantly related to changes in prices for risk reversals. Specifically, owing to the fact that for the period under consideration, the average USDJPY implied volatility function has been skewed to the downside, a normalization of the skew tends to coincide with lower returns for a variance swap strategy. We thus observe what we would expect to hold in equity markets. Increases in prices for butterfly strategies have a significantly positive impact on log variance risk premiums for all currency pairs. Intuitively, an increase in the curvature of the implied volatility function reflects heightened fear of jump risk, which in turn is positively correlated with the realized variance over a given period. With an R^2 of 40.1% and 33.3%, the relationship in (13) is material for USDJPY. In conclusion however, the results from Table 9 suggest that fear of jump risk cannot account for the variance risk premiums observed in currency markets.

7 Conclusion

Based on the theory of static replication of variance swaps, we have constructed risk-neutral estimates of future variance and assessed the sign and magnitude of variance risk premiums in foreign exchange markets. When realized variance is computed from data with a low sampling frequency, there is robust evidence of negative average variance risk premiums for both a pre-crisis sample and the full sample from January 2003 to August 2009. Evidence is ambiguous when realized variance is obtained from 5-minute interval data, which can likely be attributed to microstructure effects. Common to all estimates and currency pairs examined, variance risk premiums are highly time-varying and inversely linked to the risk-neutral expectation of future variance. Thus, whenever market participants are concerned about a rise in variance, they are willing to pay an extra premium to hedge away the corresponding risk.

In an effort to enhance our understanding of variance risk premiums in currency markets, we have assessed the roles of commonly priced risk factors and fear of jump risk. We report a robust link between variance risk premiums and the VIX, the TED spread and the general shape of the implied volatility function of the corresponding currency pair. For some biproducts, the latter explains a considerable fraction of the time variation in variance risk premiums. However, the conditional premiums remain significantly negative. Therefore, premiums associated with variance swaps are likely to

be the result of a distinct variance risk factor. This conjecture stands in contrast with the evidence from equity markets, where the fear of downside risk plays a significant role.

8 Appendix

A Derivation of Risk-Neutral Variance Forecasts

For the reader not familiar with the topic, we derive the variance estimates used throughout the paper. The exposition largely follows Carr and Madan (1998, 1999). Consider a two-period setting where investments are made at time t and payoffs are received at a future time T . No intermediary trading takes place. Assume that there is a forward market on currency pairs for delivery at time T .¹² Furthermore, presume the existence of a market for European options on currency forwards for a continuum of strike prices.¹³ Under these assumptions, any smooth, twice differentiable payoff function $h(F(T, T))$, where $F(T, T)$ denotes the terminal forward price for delivery in future time T , can be replicated by an appropriate static position in option contracts traded at initial time t . Carr and Madan (1998) establish the following formal result.¹⁴

$$h(F_T) = h(x) + h'(x)[(F_T - x)^+ - (x - F_T)^+] + \int_0^x h''(K)(K - F_T)^+ dK + \int_x^\infty h''(K)(F_T - K)^+ dK, \quad (\text{A.1})$$

where, for notational convenience, we write F_T instead of $F(T, T)$. If arbitrage is to be ruled out, the relation in the previous equation has to hold at the outset, i.e.

$$V_t = E_t^{\mathbb{Q}}(h(F_T)) = e^{-r_d T} h(x) + h'(x)[c(x, T) - p(x, T)] + \int_0^x h''(K)p(K, T)dK + \int_x^\infty h''(K)c(K, T)dK, \quad (\text{A.2})$$

where $E_t^{\mathbb{Q}}(\cdot)$ denotes the expectation operator at time t under the risk-neutral measure \mathbb{Q} and $c(\cdot)$ and $p(\cdot)$ are the time t call and put prices. Given the stochastic differential equation for currency pairs under measure \mathbb{Q} ,

$$dS(s) = (r_d - r_f)S(s)ds + \sigma(s)S(s)dB(s), \quad (\text{A.3})$$

it follows from Itô's Lemma that the forward price evolution is given by

$$dF(s, T) = \sigma(s)F(s, T)dB(s). \quad (\text{A.4})$$

¹²Currency options usually trade on the spot rather than the forward rate. Note though that as long as the forward and the option expiry coincide, the spot and the forward price eventually converge. Hence, options written on either underlying must have the same price.

¹³Clearly, this is a very idealistic assumption. However, since FX options are traded over-the-counter, a very narrow partition of the strike domain is imaginable.

¹⁴For a derivation, see Appendix A of their online version, available on <http://www.math.nyu.edu/research/carrp/research.html>. Last verified in December 2011.

Taking expectations, we deduce that the forward price is a martingale under measure \mathbb{Q} .

Reconsider the payoff function $h(F_T)$. Again using Itô's Lemma, we have

$$h(F_T) = h(F_t) + \int_t^T h'(F_s) dF_s + \frac{1}{2} \int_t^T h''(F_s) \sigma_s^2 F_s^2 ds, \quad (\text{A.5})$$

where $F_s = F(s, T)$. Consider now the twice-differentiable function

$$g(F_s) = 2 \left[\frac{F_s - F_t}{F_t} - \log \left(\frac{F_s}{F_t} \right) \right]. \quad (\text{A.6})$$

It is straightforward to verify that $g(F_t) = g'(F_t) = 0$ and $g''(F_s) = 2/F_s^2$. Replacing $h(\cdot)$ with $g(\cdot)$ in (A.5) and taking expectations under measure \mathbb{Q} yields

$$E_t^{\mathbb{Q}}(g(F_T)) = E_t^{\mathbb{Q}} \left(\int_t^T \sigma_s^2 ds \right), \quad (\text{A.7})$$

where the second term in (A.5) vanishes as a result of the martingale property of the forward rate. (A.7) is the risk-neutral expectation of future realized variance over the interval from t to T . Similarly, if we substitute for the function $g(\cdot)$ in (A.2), set $x = F_t$, and compute the future value of the expectation taken at time t , we obtain

$$\begin{aligned} V_T &= e^{rd(T-t)} E_t^{\mathbb{Q}}(g(F_T)) \\ &= \int_0^{F_t} e^{rd(T-t)} \frac{2}{K^2} p(K, T) dK + \int_{F_t}^{\infty} e^{rd(T-t)} \frac{2}{K^2} c(K, T) dK. \end{aligned} \quad (\text{A.8})$$

Equating (A.7) and (A.8), the risk-neutral expectation of future variance is given by

$$\begin{aligned} RNV_{[t,T]} &= \frac{1}{T-t} E_t^{\mathbb{Q}} \left(\int_t^T \sigma_s^2 ds \right) \\ &= \frac{2}{T-t} e^{rd(T-t)} \left[\int_0^{F_t} \frac{1}{K^2} p(K, T) dK + \int_{F_t}^{\infty} \frac{1}{K^2} c(K, T) dK \right], \end{aligned} \quad (\text{A.9})$$

where the division by $(T - t)$ is for annualization.

B Generalization of the Malz Interpolation Method

We propose a general version of the Malz interpolation method that is consistent with the regular and premium-adjusted delta convention. Given a set of option quotes σ^{DP} , σ^{DN} , σ^{DC} , we back out strike prices and compute spot deltas Δ_P , Δ_{DN} , Δ_C along the lines of Reiswich and Wystup (2010). Next, we set up the matrix

$$\mathbf{C} = \begin{pmatrix} \sigma^{DN} & 0 & 0 \\ 0 & \sigma^{RR}(\Delta_C - \Delta_P) & \sigma^{BF}((\Delta_C - \Delta_{DN})^2 - (\Delta_P - \Delta_{DN})^2) \\ 0 & 0.5\sigma^{RR}(\Delta_C + \Delta_P - 2\Delta_{DN}) & 0.5\sigma^{BF}((\Delta_C - \Delta_{DN})^2 + (\Delta_P - \Delta_{DN})^2) \end{pmatrix} \quad (\text{B.1})$$

Define a vector of parameters $\mathbf{a} = (a_0, a_1, a_2)'$. Furthermore, assemble a vector $\sigma = (\sigma^{DN}, \sigma^{RR}, \sigma^{BF})'$, where σ^{RR} and σ^{BF} follow from (6). The matrix \mathbf{C} is simply an assemblage of the basic Malz equation (7) for the delta-neutral quote, the risk reversal and the butterfly. Since (7) has to hold for any quote in the delta space, it follows that the parameters are given by

$$\mathbf{a} = \mathbf{C}^{-1}\sigma. \quad (\text{B.2})$$

An implied volatility function is obtained from (7).

References

- Andersen Torben G. and Benzoni, Luca, 2008. "Realized volatility." Working paper.
- Andersen Torben G. and Bollerslev, Tim, 1997. "Heterogeneous information arrivals and return volatility dynamics: uncovering the long-run in high frequency returns." *Journal of Finance* 52, 975-1005.
- Andersen Torben G. and Bollerslev, Tim, 1998. "Deutsche mark-dollar volatility: Intra-day activity patterns, macroeconomic announcements, and longer run dependencies." *Journal of Finance* 53, 219-265.
- Andersen, Torben G., Bollerslev, Tim, Diebold, Francis X. and Labys, Paul, 2001. "The distribution of realized exchange rate volatility?" *Journal of the American Statistical Association* 96, 42-55.
- Andersen, Torben G., Bollerslev, Tim, Diebold, Francis X. and Labys, Paul, 2003. "Modeling and forecasting realized volatility." *Econometrica* 71, 579-625.
- Bakshi Gurdip and Kapadia, Nikunj, 2003. "Delta-hedged gains and the negative market volatility risk premium." *The Review of Financial Studies* 16, 527-566.
- Bakshi Gurdip, Kapadia, Nikunj and Madan, Dilip, 2003. "Stock return characteristics, skew laws, and the differential pricing of individual options." *The Review of Financial Studies* 16, 101-143.
- Bates, David S., 1996a. "Dollar jump fears, 1984-1992: distributional abnormalities implicit in currency future options." *Journal of International Money and Finance* 15, 65-93.
- Bates, David S., 1996b. "Jumps and stochastic volatility: exchange rate processes implicit in deutsche mark options." *The Review of Financial Studies* 9, 69-107.
- Black, Fischer, 1976. "Studies of stock price volatility changes." *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistics Section*, 177-181.
- Black, Fischer and Scholes, Myron, 1973. "The pricing of options and corporate liabilities." *Journal of Political Economy* 81, 637-654.
- Britten-Jones Mark and Neuberger, Anthony, 2000. "Option prices, implied price processes and stochastic volatility." *The Journal of Finance* 55, 839-866.

- Carr, Peter and Madan, Dilip, 1998. "Towards a theory of volatility trading." In Robert Jarrow (ed.), *Risk Book on Volatility*. New York: Risk Publications, 417-427.
- Carr, Peter and Madan, Dilip, 1999. "Introducing the covariance swap." *Risk*, 47-52.
- Carr, Peter and Wu, Liuren, 2007. "Stochastic skew in currency options." *Journal of Financial Economics* 86, 213-247.
- Carr, Peter and Wu, Liuren, 2009. "Variance risk premiums." *The Review of Financial Studies* 22, 1311-1341.
- Castagna, Antonio and Mercurio, Fabio, 2007. "The vanna-volga method for implied volatilities." *Risk*, 106-111.
- Charoenwong, Charlie, Jenwittayaroje, Nattawut and Low, Buen Sin, 2009. "Who knows more about future currency volatility?" *Journal of Futures Markets* 29, 270-295.
- Christoffersen, Peter and Mazzotta, Stefano, 2005. "The accuracy of density forecasts from foreign exchange options." *Journal of Financial Econometrics* 3, 578-605.
- Coval, Joshua D. and Shumway, Tyler, 2001. "Expected option returns." *The Journal of Finance* 56, 983-1009.
- Covrig, Vicentiu and Low, Buen Sin, 2003. "The quality of volatility traded on the over-the-counter currency market: a multiple horizons study." *The Journal of Futures Markets* 23, 261-285.
- Daal, Elton A. and Madan, Dilip, 2005. "An empirical examination of the variance-gamma model for foreign currency options." *Journal of Business* 78, 2121-2152.
- Demeterfi Kresimir, Derman, Emanuel, Kamal, Michael and Zou, Joseph, 1999. "A guide to volatility and variance swaps." *The Journal of Derivatives* 7, 9-32.
- Garman, Mark B. and Kohlhagen, Steven W., 1983. "Foreign currency option values." *Journal of International Money and Finance* 2, 231-237.
- Guo, Dajiang, 1998. "The risk premium of volatility implicit in currency options." *Journal of Business & Economic Statistics* 16, 498-507.
- Hansen, Lars Peter, 1982. "Large sample properties of generalized method of moments estimators." *Econometrica* 50, 1029-1054.

- Heston, Steve L., 1993. "A closed-form solution for options with stochastic volatility with applications to bond and currency options." *The Review of Financial Studies* 6, 327-343.
- Jackwerth, Jens Carsten, 2000. "Recovering risk aversion from option prices and realized returns." *The Review of Financial Studies* 13, 433-451.
- Jiang, George J. and Tian, Yisong S., 2005. "The model-free implied volatility and its information content." *The Review of Financial Studies* 18, 1305-1342.
- Jorion, Philippe, 1995. "Predicting volatility in the foreign exchange market." *The Journal of Finance* 50, 507-528.
- Low, Buen Sin and Zhang, Shaojun, 2005. "The volatility risk premium embedded in currency options." *Journal of Financial and Quantitative Analysis* 40, 803-832.
- Malz, Allan M., 1997. "Estimating the probability distribution of the future exchange rate from option prices." *Journal of Derivatives* 5, 18-36.
- Martens, Martin and Zein, Jason, 2004. "Predicting financial volatility: High-frequency time-series forecasts vis-a-vis implied volatility." *Journal of Futures Markets* 24, 1005-1028.
- Newey, Whitney K. and West, Kenneth D., 1987. "A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix." *Econometrica* 55, 703-708.
- Pan, Jun, 2002. "The jump-risk premia implicit in options: evidence from an integrated time-series study." *Journal of Financial Economics* 63, 3-50.
- Pong, Shiuyan, Shackleton, Mark B., Taylor, Stephen T. and Xu, Xinzhong, 2004. "Forecasting currency volatility: a comparison of implied volatilities and AR(FI)MA models." *Journal of Banking and Finance* 28, 2541-2563.
- Reiswich, Dimitri and Wystup, Uwe, 2010. "A guide to FX options quoting conventions." *The Journal of Derivatives* 18, 58-68.
- Scott, Louis O., 1992. "The information content of prices in derivative security markets." *IMF Staff Papers* 39, 596-625.
- Todorov, Viktor, 2010. "Variance risk premium dynamics: The role of jumps." *The Review of Financial Studies* 23, 345-383.

Part III

The Pricing of Foreign Exchange One-Touch Options

Abstract

This article examines the structure of risk-neutral currency returns as implied in one-touch options. For this purpose, I specify models comprising pure or time-changed diffusion risk, pure or time-changed jumps, or both. The models are calibrated to vanilla options and subsequently applied to the one-touch market. Since one-touches are unspanned by a complete set of vanilla options, they lend themselves to a rigorous out-of-sample test.

The results suggest that vanilla and one-touch option markets do not generally agree on the risk-neutral dynamics of currency returns: Evidence from the vanilla market favors a complex model with stochastic volatility and jumps, whereas one-touch options imply purely diffusive currency dynamics. This latter finding gives rise to two interpretations. Either, the high activity in currency markets is best reflected by the infinite variation of a diffusive risk factor. Alternatively, the result is an artefact of market makers who anchor their quotes to what the pure diffusion Black-Scholes model implies.

1 Introduction

In this article, I examine the pricing of foreign exchange one-touch options. Foreign exchange one-touches pay out a fixed amount of cash at expiry if at any time during the term of the contract, the spot rate trades at or beyond a predefined barrier level. One-touch options are among the most liquid exotic contracts and serve as building blocks for more complex derivative securities. Hence, gaining a solid understanding of the currency dynamics implied by the one-touch market is important. In comparing different building blocks for a prospective model, one-touches have the advantage that they put stronger restrictions on the admissible pricing kernel. Specifically, one-touches depend on the whole trajectory of spot and hence its transition probabilities. In contrast, knowledge of the spot distribution at maturity is sufficient for the valuation of vanilla options. Therefore, one-touches provide valuable information on risk-neutral currency dynamics that are not revealed by the vanilla market.

Since the emergence of Black and Scholes (1973), or more precisely, the currency equivalent of Garman and Kohlhagen (1983), numerous models have been proposed to better capture the stylized facts of currency option markets. For example, Heston (1993) suggests a model that accounts for the stochastic volatility implied in option prices. Bates (1996) applies a combination of the Heston (1993) and the Merton (1976) jump-diffusion model to Deutsche mark options. The finite activity jump structure in Bates maintains that on a given time interval, currency pair dynamics exhibit a countable number of (large) jumps. More recently, Madan, Carr and Chang (1998) devise an option pricing model where an infinity of jumps is generated. To account for the high activity in currency markets, the assumption of finitely many jumps often necessitates the inclusion of a diffusion part which governs the resolution of the more subtle pieces of information. In contrast, the so-called variance gamma (VG) process of Madan et al. subsumes high-frequent small jumps and the rare occasions of large jumps in a single risk factor. In an empirical investigation, Daal and Madan (2005) find the model to cope well with the intricacies of currency markets. Finally, Carr and Wu (2007) propose a stochastic skew model where tail-specific risk factors are evaluated on a separate stochastic clock. Among other things, they report superior performance versus the Bates model.

To assess the goodness-of-fit of a particular model, all of these studies have in common that they exclusively focus on the market for vanilla options. Given the size of the market for exotic currency derivatives, this perspective is possibly too narrow. The contribution of this article is to study the risk-neutral currency dynamics using two

different information sets, namely the prices of vanilla and one-touch options. With this approach, a potential complication arises in the form of an identification problem, since any specification analysis on the basis of two unspanned information sets is a joint test of a certain model structure and a particular equivalent martingale measure. For example, a particular model that accommodates the stylized facts of the vanilla market may provide a poor fit to the one-touch data either because its structure is inappropriate or because the wrong martingale measure has been selected.

To resolve this issue, it is important to recognize the practical link between the two markets. Vanilla options are often used to hedge specific risks associated with one-touches. In fact, large banks risk-manage the two contracts in the same books. A demand for consistency therefore requires that vanilla and one-touch options are priced using the same model structure and equivalent martingale measure. This requirement allows for the subsequent analysis to be designed as an out-of-sample test. First and in accordance with previous studies, the models are calibrated to prices for vanilla options. While the resulting in-sample evidence is interesting in itself, the calibration primarily serves the purpose of selecting the equivalent martingale measure. Second, the models are applied to quoted prices for one-touch options. This approach allows for a number of interesting questions to be addressed. First, the structure of currency returns is inferred from a richer and more restrictive information set. Second, it is examined whether the vanilla and one-touch markets agree on the key ingredients of risk-neutral currency dynamics. Third and most relevant for practitioners, it is assessed whether a martingale measure identified in the vanilla market is accurate for the out-of-sample valuation of one-touch options.

The specification of currency returns follows a bottom-up approach. Specifically, I start from the basic pure diffusion model (Garman and Kohlhagen, 1983) and progressively enhance it by considering additional, possibly time-changed risk sources. This approach allows me to gauge the incremental benefit from hypothesizing more complex currency dynamics. Since countless option pricing models can be constructed by considering different constituent risk factors, I am guided by both practical and theoretical considerations to limit the scope of this article. First, option pricing models are selected that can be cast into the general framework of Carr and Wu (2004). Second, the focus lies on models that allow for an unbiased valuation of continuously monitored barrier contracts. This restriction precludes models where the joint extrema of two or more risk factors are unknown. Third, the modelling of the activity rate in currency markets is limited to the popular square root diffusion process of Cox, Ingersoll and Ross (1985). Finally, when modelling the distribution of jumps, I consider special cases

of the Lévy density proposed in Carr, Geman, Madan and Yor (2002) and Wu (2006). This jump structure is nested for example in the stochastic skew model of Carr and Wu (2007).

The results of this article suggest that the markets for vanilla and one-touch options convey different information on the dynamics of currency returns. For vanilla options, the goodness-of-fit is most accurate for a model comprising a time-changed diffusion and jumps with stochastic intensity. In contrast, one-touch option prices appear to reflect little jump risk: The best performance is achieved by the ordinary Heston (1993) model. Also, while the ranking of model performance implied by vanilla options is the same for both currency pairs examined, this is not the case for the one-touch market. Overall, the parsimony and yet strong performance of the Heston model make it an attractive choice. Two explanations for its supremacy are plausible: Either, the high activity in currency markets is best captured by diffusive price risk. Alternatively, the results are due to market makers' quoting behavior. As the one-touch option data suggests, they tend to anchor their quotes to what the Black-Scholes model, itself a pure diffusion model, implies.

The remainder of this paper is organized as follows. Section 2 develops the various currency rate dynamics examined in this study. Section 3 presents the vanilla market data, elaborates on the calibration approach and presents in-sample evidence on model performance. Section 4 deals with the market for one-touch options. It provides detailed summary statistics of the data set, explains the valuation of one-touch options and shows the out-of-sample results. Section 5 concludes.

2 The Risk-Neutral Dynamics of Currency Pairs

2.1 Pure Diffusion Model

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{Q})$ denote a filtered probability space with risk-neutral measure \mathbb{Q} . Under measure \mathbb{Q} , the Black and Scholes (1973) evolution of a currency pair, quoted as units of domestic currency (DOM) per one unit of foreign currency (FOR), is specified as

$$S_t = S_0 e^{(r_d - r_f)t + (\sigma W_t - \eta t)}, \quad (1)$$

where r_f and r_d denote the foreign and domestic interest rates, W_t is a Brownian motion, σ a volatility parameter and η the corresponding convexity adjustment.

To compute prices for vanilla options, I make use of the Fourier cosine expansion method (COS) of Fang and Oosterlee (2008). This approach requires that the risk-neutral currency price dynamics are expressed in terms of the Fourier transform of the log returns $s = \log(S_t/S_0)$,

$$\phi_s(u) = E^{\mathbb{Q}}[e^{ius_t}] = e^{iu(r_d-r_f)t} E^{\mathbb{Q}}[e^{iu(\sigma W_t - \eta t)}]. \quad (2)$$

By the Lévy-Khintchine Theorem,¹ the Fourier transform of a generic Lévy process Y_t with characteristic triplet (μ, σ^2, ν) is given by

$$\phi_Y(u) = E^{\mathbb{Q}}[e^{iuY_t}] = e^{-\psi_Y(u)t}, \quad (3)$$

where $\psi_Y(u)$ is the characteristic exponent

$$\psi_Y(u) = -iu\mu + \frac{1}{2}u^2\sigma^2 + \int_{\mathbb{R}_0} (1 - e^{iux} + iux1_{|x|<1})\nu(x)dx. \quad (4)$$

For the Black and Scholes (1973) model with triplet $(\eta, \sigma^2, 0)$, it follows that

$$\phi_s^D(u) = e^{iu(r_d-r_f)t} e^{-\psi_D(u)t}, \quad (5)$$

where²

$$\psi_D(u) = \frac{1}{2}\sigma^2(u^2 + iu). \quad (6)$$

2.2 Pure Jump and Jump-Diffusion Model

For the purpose of modelling a pure jump model, the risk-neutral currency dynamics are specified as

$$S_t = S_0 e^{(r_d-r_f)t + (X_t - \zeta t)}, \quad (7)$$

where X_t is a Lévy process with characteristic triplet $(0, 0, \nu)$ and ζ denotes the corresponding convexity adjustment. In this article, I consider jump structures that are encompassed by the process proposed in Carr, Geman, Madan and Yor (2002). The CGMY process has a Lévy density of the form

$$\nu(x) = \begin{cases} C \frac{e^{-Mx}}{x^{1+Y}}, & x > 0 \\ C \frac{e^{-G|x|}}{|x|^{1+Y}}, & x < 0 \end{cases}, \quad (8)$$

¹See e.g. Bertoin (1996).

²Note that for currency pairs to obey the martingale property, it must hold that $\eta = \frac{1}{2}\sigma^2$. In general, the convexity-adjustment is obtained by evaluating the characteristic exponent of the Lévy process without drift at the negative imaginary unit $-i$. For details, see e.g. Wu (2008).

with $C, M, G > 0$ and $Y \in [-1, 2]$. C determines the level of jump intensity, while M and G are tail-specific dampening factors. The parameter Y classifies the jump structure into finite activity ($Y < 0$), infinite activity/finite variation and infinite variation ($Y \geq 1$).³ Two popular special cases are the finite activity model of Kou (2002) and the infinite activity variance gamma (VG) model of Madan and Seneta (1990) and Madan, Carr and Chang (1998). These models are characterized by restricting the activity parameter Y to -1 and 0 respectively. By the Lévy-Khintchine Theorem, the characteristic function of the log return s is

$$\phi_s^J(u) = E^{\mathbb{Q}}[e^{ius_t}] = e^{iu(r_d-r_f)t} e^{-\psi_J(u)t}, \quad (9)$$

with characteristic exponents

$$\begin{aligned} \psi_J^{Kou}(u) = & iuC[(M-1)^{-1} - M^{-1} + (G+1)^{-1} - G^{-1}] - \\ & C[(M-iu)^{-1} - M^{-1} + (G+iu)^{-1} - G^{-1}] \end{aligned} \quad (10)$$

for the Kou and

$$\begin{aligned} \psi_J^{VG}(u) = & -iuC[\log(M-1) - \log M + \log(G+1) - \log G] + \\ & C[\log(M-iu) - \log M + \log(G+iu) - \log G] \end{aligned} \quad (11)$$

for the VG model.

Daal and Madan (2005) provide strong evidence that currency returns exhibit infinite activity. Therefore, I only consider the VG model as a representative for the pure jump models. In contrast, when jumps are nested in a more complex setting, I choose the finite activity model of Kou (2002). As will be explained in Section 4, this restriction is made in lieu of valuing the continuously monitored one-touch options.

Given the two building blocks for the pure diffusion and jump models, it is straightforward to devise a jump-diffusion model with currency pair dynamics

$$S_t = S_0 e^{(r_d-r_f)t + (\sigma W_t - \eta t) + (X_t - \zeta t)}. \quad (12)$$

This model is analogous to Merton (1976) but with the distinct jump structure of Kou (2002). The characteristic function of the log return follows from combining the characteristic functions of the diffusion and jump components, i.e.

$$\phi_s^{JD}(u) = e^{iu(r_d-r_f)t} e^{-\psi_{JD}(u)t}, \quad (13)$$

where

$$\psi_{JD}(u) = \psi_D(u) + \psi_J^{Kou}(u). \quad (14)$$

³Technically speaking, a finite activity model is defined by $\int_{\mathbb{R}_0} v(x)dx < \infty$, while finite variation refers to the condition $\int_{\mathbb{R}_0} (1 \wedge |x|)v(x)dx < \infty$.

2.3 Stochastic Volatility

Stochastic volatility can stem from either time variation in the activity rate of the diffusion process, time variation in the jump intensity, or both. Following Carr and Wu (2004), I define a stochastic clock

$$\mathcal{T}_t = \int_0^t v_s ds, \quad (15)$$

where v_t denotes the activity rate in currency markets. Throughout this article, I make the popular assumption that the activity rate follows a Cox et al. (1985) square-root diffusion process

$$dv_t = \kappa(1 - v_t)dt + \sigma_v \sqrt{v} dW_t^v. \quad (16)$$

κ specifies the mean reversion speed and σ_v governs the volatility of the activity rate. The long-run activity parameter is normalized to one since the long-run variation of the underlying is identified through the volatility or jump intensity parameter in case of a diffusion or jump model.

First, I introduce stochastic volatility by running a pure diffusion component on the stochastic clock \mathcal{T}_t . The currency dynamics are specified as

$$S_t = S_0 e^{(r_d - r_f)t + (\sigma W_{\mathcal{T}_t} - \eta \mathcal{T}_t)}, \quad (17)$$

with $E(dW_t dW_t^v) = \rho dt$. Equations (16) and (17) constitute the renowned Heston (1993) stochastic volatility model. The Fourier transform of the log return s is given by

$$\phi_s^{SV(D)}(u) = E^{\mathbb{Q}}[e^{ius_t}] = e^{iu(r_d - r_f)t} E^{\mathbb{Q}}[e^{iu(\sigma W_{\mathcal{T}_t} - \eta \mathcal{T}_t)}]. \quad (18)$$

Carr and Wu (2004) show that under a suitable complex measure change, the problem of obtaining the Fourier transform of a time-changed Lévy process becomes one of solving the Laplace transform of the stochastic time \mathcal{T}_t . Formally,

$$\phi_s^{SV(D)}(u) = e^{iu(r_d - r_f)t} \mathcal{L}_{\mathcal{T}_t}^{\mathbb{M}}(\psi_D(u)), \quad (19)$$

where $\mathcal{L}_{\mathcal{T}_t}^{\mathbb{M}}$ refers to the Laplace transform of the stochastic time \mathcal{T}_t , \mathbb{M} is a new complex measure and $\psi_D(u)$ is the characteristic exponent in (6).⁴

Alternatively, stochastic volatility is obtained by subordinating a VG process to the random clock in (15). This model, with currency dynamics

$$S_t = S_0 e^{(r_d - r_f)t + (X\mathcal{T}_t - \zeta \mathcal{T}_t)}, \quad (20)$$

⁴See Appendix A for some details.

has been proposed by Carr, Geman, Madan and Yor (2003). Following Carr and Wu (2004), the Fourier transform of the log return s is expressed as

$$\phi_s^{SV(J)}(u) = e^{iu(r_d-r_f)t} \mathcal{L}_{\mathcal{T}_t}^{\mathbb{M}}(\psi_J^{VG}(u)). \quad (21)$$

I consider two further models that combine the diffusion and jump components. The first may be viewed as an equivalent to Bates (1996). Its currency dynamics are given by,

$$S_t = S_0 e^{(r_d-r_f)t+(\sigma W_{\mathcal{T}_t}-\eta\mathcal{T}_t)+(X_t-\zeta t)}. \quad (22)$$

The Fourier transform of the log return follows from combining the characteristic functions of the stochastic volatility and jump components, i.e.

$$\phi_s^{SV(D)J}(u) = e^{iu(r_d-r_f)t} \mathcal{L}_{\mathcal{T}_t}^{\mathbb{M}}(\psi_D(u)) e^{-\psi_J^{Kou}(u)t}. \quad (23)$$

For the second model, the time change applies to both the diffusion and jump component, i.e.

$$S_t = S_0 e^{(r_d-r_f)t+(\sigma W_{\mathcal{T}_t}-\eta\mathcal{T}_t)+(X_{\mathcal{T}_t}-\zeta\mathcal{T}_t)}. \quad (24)$$

Bates (2000) suggests a related model but with Gaussian jump structure. To obtain the characteristic function, the Laplace transform is applied to the characteristic exponent in (14),

$$\phi_s^{SV(JD)}(u) = e^{iu(r_d-r_f)t} \mathcal{L}_{\mathcal{T}_t}^{\mathbb{M}}(\psi_{JD}(u)). \quad (25)$$

Under the acronyms SV1 and SV3, Huang and Wu (2004) consider variations of (22) and (24) for equity index options.

Figure 1 provides an overview of the models examined in this study. They are classified by whether they include a (time-changed) diffusion, (time-changed) jumps, or both. In general, the models increase in complexity as we move from the lower-left to the upper-right corner. The number of parameters including the initial state variable is indicated in brackets. In total, 7 different models are considered.




Figure 1: Model Overview

The figure shows the models examined in this study. The models are classified by whether they comprise a (time-changed) diffusion, (time-changed) jumps, or both. In brackets is the number of parameters including, if applicable, the initial state variable.

3 Calibrating the Models to the Vanilla Market

3.1 Data

Daily option quotes for EURUSD and USDJPY, quoted as FORDOM, are obtained from UBS, a major investment bank and market-maker in foreign exchange. The data comprises 7 quotes in the cross section for each the 1-month, 3-month, 6-month and 12-month maturity. This yields a total of 28 option quotes per currency pair and day. In the cross-section, the data set comprises quotes on delta-neutral straddles and 5-delta, 10-delta and 25-delta out-of-the-money (OTM) put and call options. The fact that 5-delta quotes are available should enhance our ability to identify the market-consistent martingale measure. Intuitively, 5-delta options reveal more refined information on the tail behavior than the more at-the-money (ATM) options. For example, Carr and Wu (2007) report difficulties in classifying the jump structure on the basis of a strike range delimited by 10-delta options. Because one-touch options are essentially contracts on threshold probabilities, the inclusion of 5-delta options is potentially important.

Since the data entails two currency pairs and 4 different maturities, the models can be challenged with regard to different implied volatility functions and term structures. For example, risk-neutral returns for USDJPY are highly skewed and do not to degenerate to normality with an increasing time horizon (see Figure 2). The stochastic volatility models of section 2.3 are well equipped to capture this slow decay. In contrast, models comprising jumps are at a comparable advantage in replicating the skew and kurtosis at the short horizon.

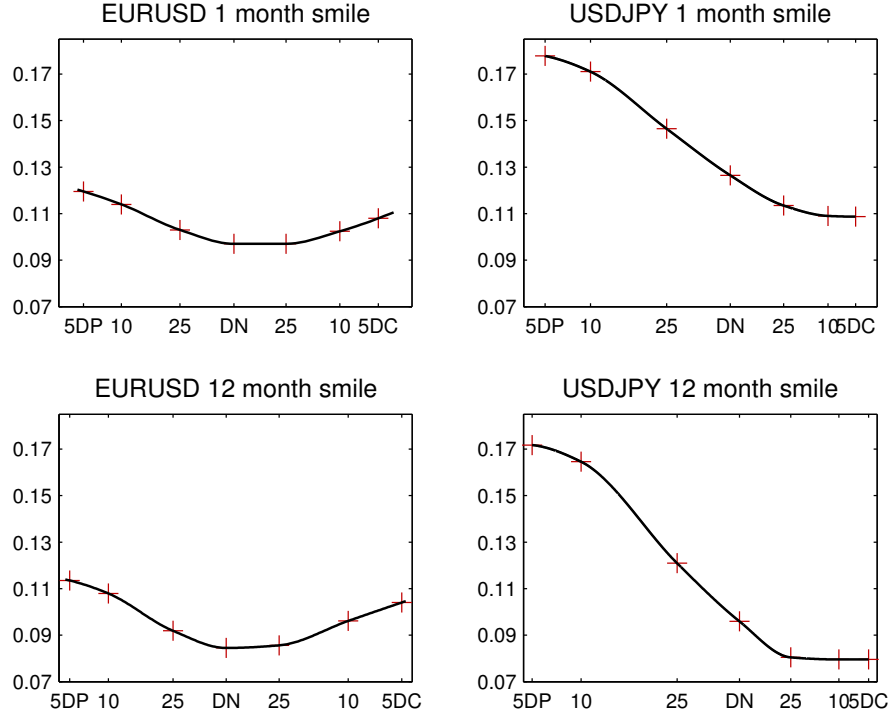


Figure 2: Exemplary Implied Volatility Functions

Exemplary implied volatility functions for 1-month and 12-month EURUSD and USDJPY, as recorded on January 2nd, 2008. The crosses mark the volatility quotes against their moneyness. The quotes have been connected by cubic spline interpolation.

The data set covers the period from January 2007 to August 2008. I have selected this period because liquidity has dried-up in the wake of the market turbulence in autumn 2008 and in particular the collapse of Lehman Brothers in mid-September. Currency spot rates for the same period are from Bloomberg. Since foreign exchange options are traded over the counter, daily interbank rates as published by the British Banker’s Association (BBA) are used. These rates are obtained from Datastream.

3.2 Calibration

The vector of parameters Θ_t is estimated by means of a search algorithm that minimizes the weighted sum of squared pricing errors

$$\Theta_t \equiv \arg \min_{\Theta_t} \sum_{i=1}^N w_i (O_{i,t} - \hat{O}_{i,t}(\Theta_t))^2. \quad (26)$$

N is the total number of quotes across the maturity and strike domain and $O_{i,t}$ and $\hat{O}_{i,t}(\Theta_t)$ denote the market and model option prices normalized by the spot rate

at time t . I follow here the standard set by Carr and Wu (2003) and Huang and Wu (2004) and apply the minimization criterion to normalized option prices rather than implied volatilities. To avoid that the calibration is biased towards long-dated ATM options, Huang and Wu (2004) suggest a weighting matrix where the elements w_i are set equal to the reciprocal of the variance of the corresponding option prices over the sample period.⁵

Concerning the market option prices for each moneyness level and maturity, I first apply the Garman and Kohlhagen (1983) formula for currency options to map volatility quotes from delta to strike space. An important subtlety is the fact that option premiums for USDJPY are quoted in USD terms. For this reason, a so-called premium-adjusted delta convention applies (see Reiswich and Wystup, 2010). To obtain option prices, implied volatilities are transformed again by use of the Garman and Kohlhagen formula.

Model prices are obtained by means of the Fourier-cosine expansion method (COS) of Fang and Oosterlee (2008). Under this method, the price of a call option is expressed as

$$c(s, t) = e^{-r_d(T-t)} \sum_{k=0}^{N-1} {}'Re \left[\phi \left(\frac{k\pi}{b-a} \right) e^{ik\pi \frac{s-a}{b-a}} \right] V_k, \quad (27)$$

where N is the number of expansion terms, ϕ any of the characteristic functions from the previous section and V_k a series of payoff-dependent coefficients. a and b denote truncation parameters, while \sum' indicates that the first summation term only counts half. Details have been relegated to Appendix B. I choose the COS method over the more common Fast Fourier transform (FFT) approach of Carr and Madan (1999) for two reasons. First, the COS method is more efficient, i.e. it achieves the same pricing accuracy with considerably fewer summation terms (Fang and Oosterlee, 2008). Second, rather than interpolating from a discrete log-strike grid as is the case for the FFT method, option prices for arbitrary strikes can immediately be computed.

When a model features at least two state variables, it can be cast into a dynamic estimation framework which enforces model consistency through time.⁶ However, in the current context, this approach has a number of drawbacks: First, not all models from the previous section satisfy the two-state-variable condition. Second, the approach naturally favours models with a larger number of state variables. Third, frequent parameter re-estimation appears to have established as today's industry standard. Therefore,

⁵See Huang and Wu (2004) for more details on this weighting scheme and the conditions under which it is optimal.

⁶See e.g. the estimation approach in Carr and Wu (2007).

results from dynamic estimation may be of limited interest to practitioners. In line with the estimation procedure in Carr and Wu (2003) and Huang and Wu (2004),⁷ I calibrate the models on a daily basis. Implicitly, a model is asked to consistently price the volatility surface on any given date, but not necessarily to be parameter-stable through time. With this approach, the initial state variable v_0 can be used as an additional free parameter. At the initial estimation stage on January 2nd 2007, I support the calibration by manually fine-tuning the parameters based on a graphical inspection of the market and model implied volatility surfaces. On subsequent days, the previous-day parameter estimates are selected as an initial guess for the search.

3.3 In-Sample Evidence on the Structure of Currency Returns

Table 1 presents the EURUSD parameter estimates for the different models under consideration. The parameters reported are averages computed from the daily estimates for the period from January 2007 to August 2008. With regard to model performance, two statistics are considered. The mean absolute price error (MAPE) is computed as

$$MAPE = \frac{1}{T} \frac{1}{N} \sum_{t=1}^T \sum_{i=1}^N \frac{|O_{i,t} - \hat{O}_{i,t}|}{O_{i,t}}, \quad (28)$$

where the first summation indicates that an average is taken over all daily observations. I further report a statistic that is based on a comparison of the model implied volatility surface and the volatility quotes observed in the market. Specifically, I transform the model prices $\hat{O}_{i,t}$ into $\hat{\sigma}_{i,t}$ and compute

$$MAVPD = \frac{1}{T} \frac{1}{N} \sum_{t=1}^T \sum_{i=1}^N |\sigma_{i,t} - \hat{\sigma}_{i,t}|, \quad (29)$$

where MAVPD denotes 'mean absolute vol point deviation'.

Among the different specifications, the pure jump VG model attains the worst performance. The model has an average MAPE and MAVPD of 17.2% and 0.52.⁸ Contrary to the findings of Daal and Madan (2005), the jump-diffusion model with finite jump activity considerably improves upon this performance. Both the MAPE and MAVPD drop by about 60%. A still better fit is obtained for the pure diffusion stochastic volatility model SV(D). Its MAPE and MAVPD are 3.5% and 0.09. The initial activity rate v_0 is below 1, which implies that the average volatility term structure

⁷See also the seminal paper of Bakshi, Cao and Chen (1997).

⁸In line with the common volatility quoting convention, I omit the percentage sign for the MAVPD statistic.

	<i>VG</i>	<i>JD</i>	<i>SV(D)</i>	<i>SV(J)</i>	<i>SV(D)J</i>	<i>SV(JD)</i>
σ	-	7.3%	9.1%	-	4.4%	6.3%
κ	-	-	2.50	1.67	0.78	1.57
σ_v	-	-	3.24	2.52	8.47	2.54
ρ	-	-	-0.04	-	-0.02	-0.09
v_0	-	-	0.83	0.87	34.4	0.83
C	230.0	2.3	-	101.4	291.4	3170.4
G	243.5	10.6	-	132.8	101.8	318.2
M	264.4	12.6	-	177.5	79.3	147.9
MAPE	17.2%	6.7%	3.5%	4.7%	2.9%	2.0%
MAVPD	0.52	0.22	0.09	0.13	0.09	0.06

Table 1: EURUSD Parameter Estimates and Model Performance

EURUSD parameter estimates and summary statistics for the different models. The parameter estimates refer to averages over the sample period from January 2007 to August 2008. MAPE and MAVPD denote the 'mean absolute price error' and 'mean absolute vol point deviation'.

is upward sloping. The correlation parameter of -0.04 suggests that the EURUSD implied volatility function tends to exhibit little skew. When stochastic volatility is introduced via jumps, the MAPE and MAVPD are 4.7% and 0.13. In line with the findings of Huang and Wu (2004), the pricing errors for the SV(J) are larger than for the SV(D) model. A comparatively strong performance is achieved when the pure diffusion stochastic volatility model is combined with Kou jumps. However, the SV(D)J model exhibits a disturbingly large initial activity rate. With a MAPE and MAVPD of only 2.0% and 0.06, the best performance is achieved by the SV(JD) model. This finding should not surprise, since it offers the most flexible specification. In accordance with the presumably high activity in currency markets, its jump intensity parameter C is far larger compared to the JD or SV(D)J model.

Table 2 presents the results for USDJPY. In general, the average pricing errors suggest that it is much harder to capture the volatility surface of a skewed currency pair. As for EURUSD, the VG model performs worst, with a MAPE and MAVPD of 23.1% and 1.06. Undoubtedly, the results of this section recommend against the use of the VG model as a self-contained pricing paradigm. Contrary to the findings for EURUSD, the performance is not materially different for the JD model. On the other hand, the SV(D) model is again very successful in reducing average pricing errors. Its MAPE and MAVPD are 6.5% and 0.24 or about 75% lower compared to the VG model. The parameter estimates suggest that the USDJPY implied volatility surface is highly skewed to the downside ($\rho = -0.57$) and negatively sloped in the time-dimension ($v_0 = 1.37$). The SV(J) model fails to adequately reflect these stylized facts. The initial activity rate seems unnaturally high while the jump intensity C is very low.

	<i>VG</i>	<i>JD</i>	<i>SV(D)</i>	<i>SV(J)</i>	<i>SV(D)J</i>	<i>SV(JD)</i>
σ	-	5.9%	10.0%	-	9.8%	8.7%
κ	-	-	2.36	0.65	1.56	1.42
σ_v	-	-	4.20	9.33	4.26	3.23
ρ	-	-	-0.57	-	-0.66	-0.73
v_0	-	-	1.37	2970.1	1.41	1.23
C	13.1	4.7	-	0.6	525.2	578.1
G	20.2	7.0	-	21.8	250.0	186.6
M	41.0	37.5	-	45.2	651.2	133.4
MAPE	23.1%	21.8%	6.5%	19.6%	5.9%	5.8%
MAVPD	1.06	1.03	0.24	0.85	0.23	0.23

Table 2: USDJPY Parameter Estimates and Model Performance

USDJPY parameter estimates and summary statistics for the different models. The parameter estimates refer to averages over the sample period from January 2007 to August 2008. MAPE and MAVPD denote the 'mean absolute price error' and 'mean absolute vol point deviation'.

Of course, the former makes up for the latter, but apparently without success. Both the SV(D)J and SV(JD) models perform well, this time with little difference between the two. Two major findings emerge: First, when the implied volatility skew is large, allowing for correlation between the spot and variance risk factor is important. For the JD or SV(J) model where this feature is ruled out, a large discrepancy between the EURUSD and USDJPY pricing performance is observed. As a consequence, diffusive price risk is a key component of risk-neutral currency dynamics. Second, while offering a strong performance, the SV(D) model is still enhanced when jumps are included. The impact of a time-varying jump intensity is thereby larger for a moderately skewed currency pair.

4 Evidence from One-Touch Options

In this section, the out-of-sample tests on the one-touch option market are presented. To fix matters, let $H(S_u, B, T)$ denote the event of a barrier crossing,

$$H(S_u, B, T) := 1_{\tau_B < T}, \quad \tau_B := \inf\{u \geq 0 : \phi S_u \leq \phi B, \quad 0 < u < T\} \quad (30)$$

where B denotes the barrier level and $\phi \in \{-1, 1\}$ is an indicator variable corresponding to an upper and lower barrier. Measured in domestic currency, the value of a one-touch option paying at expiry one unit of domestic currency is given by

$$G(S_u, B, T) = e^{-r_d T} E^{\mathbb{Q}}[H(S_u, B, T)], \quad 0 < u < T]. \quad (31)$$

One-touch options in foreign exchange are quoted in terms of (31), i.e. as a percentage discounted hit probability. For EURUSD and USDJPY, the convention is to

quote one-touches in foreign units, i.e. EUR in the first and USD in the second case. This requires the modification

$$G(S_u, B, T) = e^{-r_f T} E^{\mathbb{Q}}[H(S_u, B, T)] B E^{\mathbb{Q}}[e^{(r_d - r_f)(T - \tau_B)}] / S, \quad (32)$$

since the value of one unit of foreign currency at the barrier level is different from today's spot price. The second expectation in (32) is the forward value of the foreign currency unit contingent upon a barrier event. Under the Black-Scholes model assumptions, (31) and (32) can be solved analytically.⁹ I refer to the corresponding values as the unskewed theoretical value (TV). In case of the more complex models, the expectations need to be evaluated using numerical techniques.

The interbank market for one-touch options is organized as a broker market. A bank with an open interest asks a broker to obtain quotes from other market makers. For this purpose, the bank and the broker, and subsequently any of the market makers willing to show a quote, seek mutual agreement on the interest rates, the reference spot price and the reference volatility, and thus the unskewed TV. This mutual agreement ensures that the parties involved talk about the same contract. When a quote is shown, it remains binding unless spot and volatility change considerably.¹⁰ The parties involved are kept updated about the most recent bid-offer spread, and they may choose to improve their quote. They remain anonymous throughout the process, and only those parties whose bid and offer finally match get informed about the counterparty they are dealing with.

4.1 Data and Biases

Quotes on EURUSD and USDJPY one-touches are received from the same source as the vanilla data. The period under consideration ranges from January 2007 to August 2008. The quotes have been observed in the interbank market, where access is granted only to sophisticated market makers, primarily the major investment banks. The interbank market for one-touch options is highly liquid. To further enhance the quality of the data, several filters are applied: First, to match the maturities of the vanilla options, one-touches with a tenor shorter than three weeks and longer than one year are dismissed. Second, obviously erroneous records, e.g. those where the bid is above the ask price, are excluded. Third, while the reference spot price and volatility are given, the mutually-agreed interest rates are not observed. In general, these rates will be somewhat different from the BBA interbank rates used for this study. To avoid

⁹A brief exposition along with the explicit pricing formulae is given in Appendix C.

¹⁰One-touch options are traded simultaneously with a delta and vega hedge. Therefore, large deviations from the reference spot price and volatility must occur to warrant a cancellation of a quote.

that the results are affected by this subtlety, I compare the unskewed TV based on the BBA interest rates with the reference TV recorded in the data set. One-touches where this difference is larger than 1% are dismissed. After applying all filters, 587 quotes for EURUSD and 519 quotes for USDJPY remain.

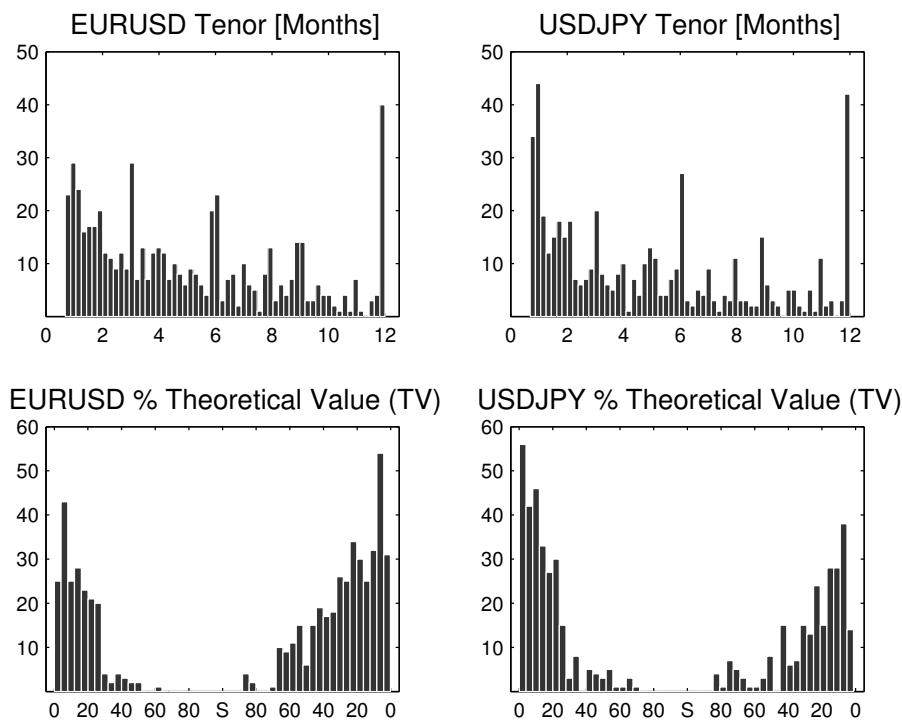


Figure 3: One-Touch Option Data

Summary statistics for the one-touch option data. For both currency pairs, the figure reveals the distribution of the one-touch option quotes in the time and moneyness dimension, where moneyness is proxied by the unskewed theoretical value.

Figure 3 reveals summary statistics for the one-touch data under examination. The upper panels show the distribution of contract terms over the admitted range from three weeks to one year. For both EURUSD and USDJPY, there are peaks at the 1-month, 3-month, 6-month and 12-month tenor. Although one-touches trade over-the-counter, these maturities are standard. The lower panels show the distribution of one-touch quotes in terms of unskewed TV: At the center of each plot is the spot rate. Upside and downside one-touches are distributed to the right and left, where upside (downside) refers to the case when the barrier lies above (below) the spot rate at trade date. Since the unskewed TV roughly corresponds to the hit probability, it may be interpreted as a measure of moneyness. For EURUSD, 65.4% (34.6%) of the quotes correspond to upside (downside) one-touches. Figure 3 shows that most of them

concentrate in the tails: Roughly 53% have a TV below 20%. Out of the 587 quotes, 169 or about 29% are traded prices. For the remainder, the average (maximum) bid-ask spread is 1.00% (4.00%). For USDJPY, the distribution between upside (45.5%) and downside (54.5%) is more even. 62% of the quotes have a TV below 20%. About 21% or 109 quotes are traded prices. For the others, the average (maximum) bid-ask spread is 1.59% (5.25%).




Figure 4: Black-Scholes Model Biases in One-Touch Options

EURUSD and USDJPY valuation biases for the Black-Scholes model across maturities and in the moneyness dimension. Biases are computed as market minus model-implied prices.

The horizontal axis corresponds to the unskewed (Black-Scholes) TV.

Figure 4 provides an overview of the relation between market and Black-Scholes model prices, where the latter are computed using the reference volatility from the data set. For each panel, the vertical axis shows the valuation bias, measured as the mid-market quote minus the unskewed TV. Each cross corresponds to a single one-touch option. The upper panels show the biases across maturities, while the lower panels plot them against the unskewed TV.

The lower-left panel shows that for EURUSD, low TV one-touches, whether upside or downside, trade considerably above the unskewed TV. This finding suggests that for

barriers far away from spot, traders assign a higher hit probability than is implied by the Black-Scholes model. In contrast, one-touches with an unskewed TV above 20% tend to trade at a "discount". For USDJPY, the lower-right panel shows that low TV downside one-touches trade up to 8% above the unskewed TV. On the other hand, the Black-Scholes model overestimates the market price of upside one-touches. These facts imply an above normal concentration of risk-neutral probability mass in the lower tail. For both currency pairs, the bias pattern observed in the moneyness dimension has been remarkably stable for the period from January 2007 to August 2008. One reason for this finding may be that traders anchor their quotes to the Black-Scholes TV: To obtain market prices, they seemingly adjust the unskewed TV by a TV-dependent mark-up. Contrary to the moneyness dimension, the valuation biases are mostly unrelated to the maturity of the one-touches (upper panels).

4.2 Calibration, Pricing and Numerical Issues

The various models are applied to the one-touch data set in a manner that is consistent with industry practice. On each date where one or more one-touch quotes are available, the parameter estimates from calibrating the models to the vanilla surface are recalled. One problem is that while the vanilla data is recorded at the end of a trading day, i.e. 2pm eastern time, the one-touch quotes may be from any time during the day. To account for this intra-day mismatch, I linearly interpolate the delta-neutral vanilla quotes and select the quote that corresponds to the tenor of the one-touch option. Next, a parallel-shift to the implied volatility surface is applied such that the interpolated vanilla quote coincides with the one-touch reference volatility. The models are recalibrated to this adjusted curve, using the same-day vanilla parameters as the starting values for the search algorithm.

With the exception of the VG model, all the models feature multiple sources of randomness. Given the resultant complexity, Monte Carlo simulation is particularly suitable to evaluate the expectation in (32). Trajectories of spot prices are directly obtained from (7) for the VG, (12) for the jump-diffusion, (17) for the Heston, (20) for the stochastic jump intensity, (22) for the jump-enhanced stochastic volatility and (24) for the stochastic jump-diffusion model. I discretize the time line using an equidistant calendar-time grid with one spot realization per day. To price a single one-touch option, I simulate 150k sample paths. This number ensures that the sample standard deviation for a one-touch price is in the vicinity of 0.10% or slightly below.

To simulate trajectories for the activity rate, I discretize the square-root diffusion process in (16) using the Euler scheme. This choice is justified by Higham and Mao

(2005), who show that the scheme converges strongly as the time step goes to zero. One problem though with the Euler discretization is the chance of negative activity rates, which would imply that business time moves backwards. To circumvent this problem, Lord, Koekkoek and van Dijk (2010) suggest a full truncation scheme which they show to be superior relative to other commonly used discretization approaches. While I have alternatively considered the unbiased algorithm of Broadie and Kaya (2006) and the quadratic-exponential scheme of Andersen (2008), it turns out that for the step size used in this study, the full truncation scheme of Lord et al. offers the best accuracy-speed trade-off.

When simulating spot trajectories, the chosen grid step size is too coarse to provide accurate prices for the continuously monitored one-touch options. For example, comparing simulated and analytical Black-Scholes TVs, the average simulation bias over all EURUSD one-touches is approximately -1.9% (the average analytical TV is 22.3%). The reason for this bias is that the true extrema almost surely lie between the grid points and are thus ignored. For a Brownian motion, Beaglehole, Dybvig and Zhou (1997) propose to simulate a Brownian bridge from the known distribution of extrema to connect adjacent grid points. Ribeiro and Webber (2006) extend this result to the class of models which can be represented as time-changed Brownian motion. This includes the VG jump model, which may be expressed as a Brownian motion subordinated to a gamma clock. To the best of my knowledge, no comparable result is available for the joint maximum or minimum distribution of two or more general Lévy processes. This fact precludes, for example, the simulation of unbiased prices for the stochastic skew model of Carr and Wu (2007). Furthermore, it restricts jumps to the finite activity class when they are combined with a diffusion term.¹¹ Since the models from Section 2 have been selected on these grounds, the results of Ribeiro and Webber apply for all models. Safe for the approximation error mentioned in their paper, simulated one-touch prices are unbiased.

Simulation of the Brownian risk factors is straightforward. I refer to Glasserman (2004) for a general reference. In case of the pure jump VG model, I simulate X_t using the well-known subordinator representation. The simulation of Kou jumps is cast into a compound Poisson framework. Some details are revealed in Appendix D.

¹¹When jumps exhibit finite activity, the equidistant grid can be augmented by simulated jump times. Brownian bridge sampling can then be applied to this extended grid.

4.3 Results

Analogous to Figure 4, Figure 5 shows the EURUSD valuation biases for the various models under examination. Since the pricing errors for the Black-Scholes model were unrelated to maturity, the focus lies on the moneyness dimension. As before, biases are calculated as the mid-market minus the model-implied prices. The horizontal axis refers to the unskewed (Black-Scholes) TV. To facilitate interpretation, I have augmented the charts with a Gaussian kernel estimate of the pricing bias using a bandwidth of 10% (black line).

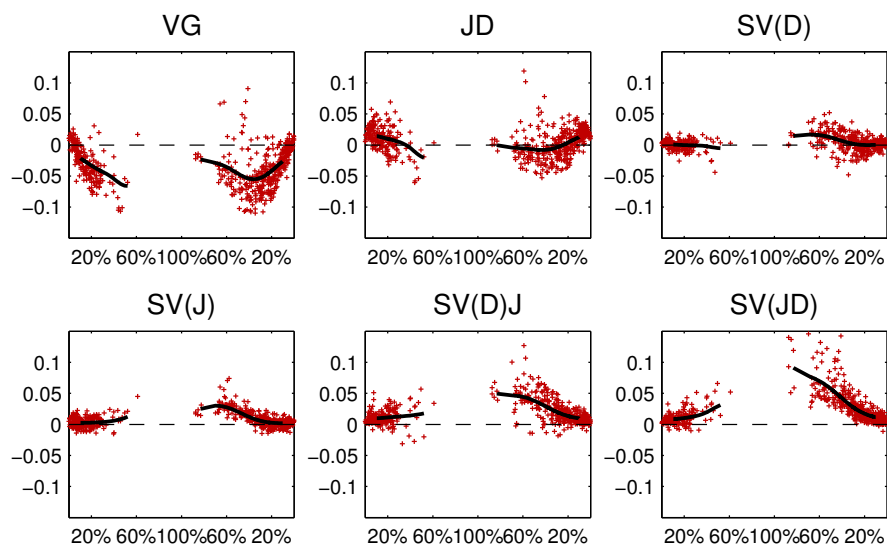


Figure 5: EURUSD Pricing Errors

EURUSD pricing performance for the various models. The pricing bias (vertical axis) is depicted against the unskewed TV (horizontal axis), with downside and upside one-touch options (crosses) being plotted left and right of the center. A kernel estimate of the pricing biases is given by the black line.

The VG model (upper-left panel) provides fairly accurate prices for low TV upside and downside one-touches. However, when the unskewed TV is above 10%, the model severely overestimates the market prices. Moreover, the dispersion of pricing errors is large. For example, for upside one-touch options with a TV of approximately 30%, pricing errors range between plus and minus 10%, meaning that the model-implied TV is anywhere between 20% and 40%. For practitioners, this degree of precision is not acceptable. The jump-diffusion model (upper-middle panel) improves on both accounts: The valuation bias (black line) is closer to zero and the dispersion of pricing errors is smaller. However, the model tends to underestimate market prices in the extreme tails. This bias is corrected by the Heston model (upper-right panel). No

evident biases are observed for downside and low TV upside one-touch options. A similarly good performance is observed for the stochastic jump intensity model (lower-left panel). When the Heston model is augmented with Kou jumps (lower-middle panel), it tends to underestimate market prices for moderate and high TV one-touches. This pattern is aggravated when the jumps are sampled from a model with stochastic jump intensity (lower-right panel). Therefore, the sophisticated models perform considerably worse out-of-sample than their more parsimonious counterparts.

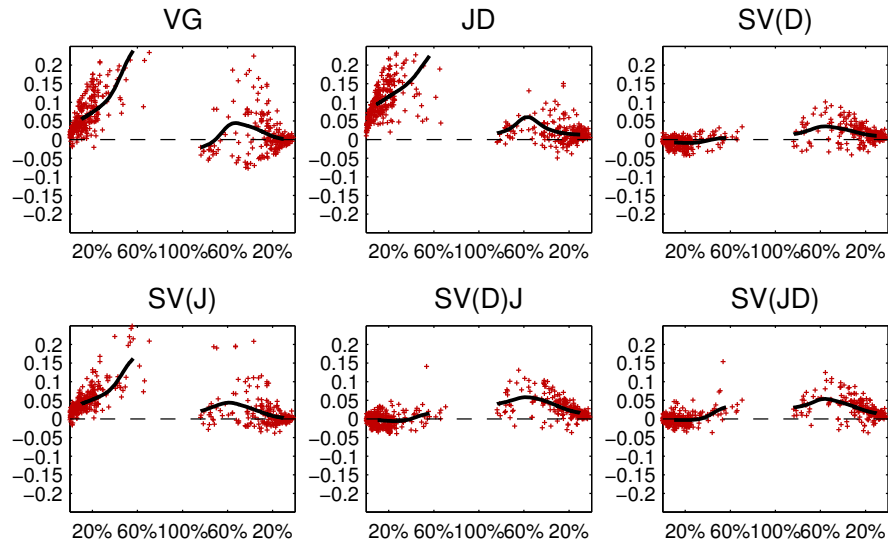


Figure 6: USDJPY Pricing Errors

USDJPY pricing performance for the various models. The pricing bias (vertical axis) is depicted against the unskewed TV (horizontal axis), with downside and upside one-touch options (crosses) being plotted left and right of the center. A kernel estimate of the pricing biases is given by the black line.

Figure 6 shows the valuation biases for USDJPY. The VG model (upper-left panel) exhibits large biases for downside one-touch options and large pricing error dispersion for contracts where the barrier lies above the spot rate.¹² Compared to EURUSD, the average bias has a positive sign, i.e. the model underestimates the probability of a barrier event. The jump-diffusion model (upper-middle panel) somewhat reduces the pricing error dispersion for upside one-touches. However, the downside biases are as pronounced as for the VG model. The Heston model (upper-right panel) removes most of these biases and is very accurate in general. Contrary to the findings for EURUSD, the performance of the stochastic jump intensity model (lower-left panel) closely resembles the VG model performance. In particular, the model shares a large bias for

¹²Note that the scale is different compared to Figure 4 and Figure 5.

moderate and high TV one-touch options in the lower tail. The sophisticated $SV(D)J$ and $SV(JD)$ models (lower-middle and lower-right panels) reveal a bias pattern that is similar to the Heston model. They provide approximately unbiased prices for downside one-touches but tend to underestimate market prices for moderate TV upside barrier contracts.

To allow for a more precise assessment of model performance, Table 3 sorts one-touch options into different moneyness buckets. Specifically, it is distinguished between upside and downside one-touches with an unskewed TV below 10%, between 10% and 20% and above 20%. The very-left column reveals the number of quotes n in each bucket. The remaining columns show the 'mean absolute price error' (MAPE) in percentage terms, now defined as

$$MAPE = \frac{1}{n} \sum_{i=1}^n |O_{i,t} - \hat{O}_{i,t}|. \quad (33)$$

For EURUSD, the Heston model has the lowest overall MAPE (0.87%). This performance is closely matched by the stochastic jump intensity model (0.94%), while the next best model ($SV(D)J$) performs considerably worse. Compared to Black-Scholes, the $SV(D)$ and $SV(J)$ models reduce the MAPE by about 54% and 50%. If the various moneyness buckets are viewed in isolation, it holds that in all except one instance, the Heston model comes first, followed by the $SV(J)$ model. The overall improvement of the valuation performance relative to Black-Scholes is very modest for the jump-diffusion (0.5%) and the jump-augmented stochastic volatility models (3%). Most surprisingly, the sophisticated $SV(JD)$ model exhibits an average absolute pricing error in excess of Black-Scholes. The overall performance of the $SV(JD)$ model may be driven by its inability to mirror prices for upside one-touches with a TV between 20% and 100%. In contrast, all models perform superior versus Black-Scholes for one-touch options in the 10% upside and downside tail buckets.

For USDJPY, the best valuation performance is again achieved by the Heston model. It has a MAPE of 1.63%, which amounts to an error reduction of approximately 61% relative to Black-Scholes. Contrary to the findings for EURUSD, the $SV(JD)$ model provides a good fit. In particular, the model exhibits the smallest MAPE for downside one-touches with a TV below 20%. Thus, the model is well suited to capture a fat-tailed left distribution that is usually characteristic of equity markets. The $SV(J)$ model is accurate for pricing low TV upside one-touches, but fails to properly reflect the remainder. Both the VG and the jump-diffusion models attain an MAPE that is larger compared to Black-Scholes, although they provide fairly accurate prices for upside one-touches with a TV below 20%.

	<i>BS</i>	<i>VG</i>	<i>JD</i>	<i>SV(D)</i>	<i>SV(J)</i>	<i>SV(D)J</i>	<i>SV(JD)</i>
EURUSD							
10%D (n=82)	2.25	1.29	2.04	0.46	0.70	0.94	0.72
10%U (n=104)	2.18	1.25	2.06	0.42	0.44	0.74	0.79
20%D (n=56)	0.76	3.43	1.40	0.48	0.62	0.96	0.94
20%U (n=67)	0.83	3.91	1.50	0.68	0.50	1.08	1.31
100%D (n=65)	1.80	5.32	1.95	0.81	0.92	1.91	1.86
100%U (n=213)	2.25	5.74	1.93	1.42	1.50	3.11	4.51
Overall	1.88	3.85	1.87	0.87	0.94	1.82	2.32
USDJPY							
10%D (n=119)	5.13	3.57	7.02	0.85	2.80	0.94	0.79
10%U (n=56)	1.83	0.64	1.27	0.76	0.53	1.03	0.94
20%D (n=84)	3.67	6.52	11.2	1.28	4.90	1.57	1.14
20%U (n=63)	5.12	1.53	1.58	1.37	1.03	1.93	1.83
100%D (n=80)	2.37	14.3	16.7	1.42	9.51	1.81	1.77
100%U (n=117)	5.20	5.80	3.71	3.37	3.72	4.64	4.33
Overall	4.13	5.64	7.16	1.63	3.92	2.14	1.94

Table 3: Pricing Errors by Moneyness Bucket

Pricing errors for the different models grouped by moneyness buckets. n in the far-left column reveals the number of quotes per bucket. The letters U and D refer to upside and downside option quotes. The pricing errors are computed as the percentage absolute pricing bias.

To draw robust conclusions on the fine structure of currency returns, the models are formally compared on the basis of the 'mean absolute price error' difference. For models j and k, t-statistics on the MAPE are computed as

$$tstat = \frac{MAPE_j - MAPE_k}{\hat{\sigma}_{(APE_{j,i} - APE_{k,i})}}. \quad (34)$$

A non-significant t-statistic suggests that model j and k provide a similar fit to the market for one-touch options. Since the series $(APE_{j,i} - APE_{k,i})$ is likely to be autocorrelated, $\hat{\sigma}_{(APE_{j,i} - APE_{k,i})}$ is computed using the Bartlett estimate as propose by Newey and West (1987). To determine the number of lags, I plot autocorrelation functions and set the number of lags equal to the first non-significant autocorrelation number. Table 4 reads as follows: The entry in row j and column k show the difference in MAPE for a performance comparison of model j and k. A positive number suggests that model k, i.e. the column model, is superior. Asterisks indicate significance. The lower triangular matrix is empty because it simply mirrors the statistics shown in the upper triangle.

The first row of Table 4 shows that for EURUSD, the Black-Scholes model outperforms the VG and SV(JD) model on a 1% and 5% confidence level, respectively. In contrast, the model's performance is indistinguishable from the *JD* and *SV(D)J*,

	<i>BS</i>	<i>VG</i>	<i>JD</i>	<i>SV(D)</i>	<i>SV(J)</i>	<i>SV(D)J</i>	<i>SV(JD)</i>
EURUSD							
<i>BS</i>	-	-1.96%***	0.01%	1.02%***	0.95%***	0.07%	-0.44%**
<i>VG</i>	-	-	1.97%***	2.98%***	2.91%***	2.03%***	1.53%***
<i>JD</i>	-	-	-	1.00%***	0.93%***	0.06%	-0.45%**
<i>SV(D)</i>	-	-	-	-	-0.07%*	-0.95%***	-1.45%***
<i>SV(J)</i>	-	-	-	-	-	-0.88%***	-1.38%***
<i>SV(D)J</i>	-	-	-	-	-	-	-0.50%***
USDJPY							
<i>BS</i>	-	-1.51%***	-3.03%***	2.50%***	0.21%	1.99%***	2.19%***
<i>VG</i>	-	-	-1.52%***	4.01%***	1.72%***	3.50%***	3.70%***
<i>JD</i>	-	-	-	5.53%***	3.23%***	5.01%***	5.22%***
<i>SV(D)</i>	-	-	-	-	-2.29%***	-0.51%***	-0.31%***
<i>SV(J)</i>	-	-	-	-	-	1.78%***	1.98%***
<i>SV(D)J</i>	-	-	-	-	-	-	0.20%***

Table 4: Performance Comparison

Performance comparison of the various models based on the mean absolute price error (MAPE). The reported MAPE differences for row j and column k reveal whether model j is superior to model k . T-statistics are computed using Newey-West (1987) robust standard errors. *, **, and *** indicate significance at the 10%, 5% and 1% confidence level.

and is significantly worse than the performance of the $SV(J)$ and $SV(D)$ models. The jump-diffusion model is on par with the more sophisticated $SV(D)J$, but significantly worse than the stochastic volatility and jump intensity models. The supremacy of Heston is significant on a 10% level versus the stochastic jump intensity model and highly significant versus the remainder. For USDJPY, the Black-Scholes model outperforms the VG and the jump diffusion models on a 1% confidence level, but is outperformed by all models comprising a time-changed diffusion. The USDJPY one-touch market is the only place where I find similar evidence as Daal and Madan (2005), who report superior performance for the VG versus the jump-diffusion model. One reason for this seemingly contradictory finding may be that I employ a Kou rather than a Gaussian jump structure. For USDJPY, the superior performance of the Heston model is significant on a 1% level against all models. In contrast, the $SV(J)$ model is just on par with Black-Scholes and significantly worse than the jump models comprising a time-changed diffusion.

Several key findings emerge. First, the unique characteristics of currency pairs may require a distinct modelling of risk-neutral currency dynamics. For USDJPY, both the in- and out-of-sample results emphasize the importance of including a time-changed diffusion: By allowing currency returns and variance to be negatively correlated, the USDJPY skew is captured well. For EURUSD, where strong performance is also reported for the stochastic jump intensity model, this feature seems less important. Second and

irrespective of the $SV(J)$'s performance, I conclude that a time-changed diffusion is an indispensable ingredient of risk-neutral currency dynamics, while models without a time-changed risk factor should be discouraged. Most notably, the $SV(D)$ model provides the best fit for both EURUSD and USDJPY one-touches, even though it came only third in replicating the vanilla surface. Third, while the same ranking of model performance was obtained in-sample for EURUSD and USDJPY, this is no longer true for the one-touch option market. For example, the $SV(JD)$ model does well in capturing the stylized facts of USDJPY, but provides a relatively poor fit for EURUSD. One problem in this context is certainly the danger of overfitting. Still, I conjecture that the vanilla and one-touch market need not necessarily agree on the structure of currency returns. In particular, the prices for one-touch options appear to reflect little jump risk, while the same can not be said for the vanilla market. For practitioners, this finding has potentially far-reaching implications. Either they include exotic derivatives when calibrating a model so as to enforce consistency or they keep currency dynamics simple. Given the solid performance of the Heston model also for the vanilla market, modelling risk-neutral currency dynamics as a time-changed diffusion is a reasonable choice.

The source of the strong relative performance of the $SV(D)$ model remains subject to speculation, but two hypotheses can be put forward. Either it may be that risk-neutral currency dynamics, as implied by the one-touch option data, exhibit infinite variation. While a Brownian motion is characterized by infinite variation, the jump processes considered in this study are not. If we imagine an overall risk budget that is allocated, through the process of calibration, to the various risk factors, then the models comprising jumps will partially fail to reproduce the stylized facts. This interpretation is supported by the average volatility parameter estimates from Table 1 and Table 2. Since the parameters tend to be smaller for the models comprising jumps, the contribution of the diffusive factor to overall model-implied risk is diminished. Alternatively, the superior performance of the Heston (1993) model is the result of the aforementioned use of the Black-Scholes value as a reference price. Such behavior by traders is likely to favor a diffusive risk specification, since the Black-Scholes model itself is a diffusion model.

5 Conclusion

Within the framework of Carr and Wu (2004), I have specified a range of option pricing models to determine the key ingredients of risk-neutral currency dynamics. Currency

dynamics are specified as pure or time-changed diffusion, pure or time-changed jumps, or both. The models have been calibrated to a set of vanilla option quotes and subsequently applied to obtain prices for one-touch options.

In-sample evidence from calibrating the models to vanilla options favors complex specifications including a time-changed diffusion and jumps. This finding is consistent with the previous literature. In contrast, the comparably simple Heston (1993) model provides the best fit for the one-touch options. Both markets agree on the minimum requirement that risk-neutral currency dynamics should comprise a time-changed diffusive risk factor. However, they do not agree on whether jumps should be included. More generally, a ranking of model performance yields different results for the two markets. Therefore, caution is warranted when drawing inference from the vanilla market and transferring it to the valuation of one-touch options.

Augmenting risk-neutral currency dynamics with jumps may still be sensible if the focus lies on pricing very short-dated options. The shortest maturity examined in this study is one month for vanilla and three weeks for one-touch options. Hence, no conclusions can be drawn for the very front-end of the volatility term structure. However, given the robust performance of the parsimonious Heston model in both markets, the practitioner has good reasons to prefer it over more complex specifications.

6 Appendix

A Measure Change and Laplace Transform

This Appendix draws on Carr and Wu (2004). Consider, by example, the Fourier transform in (18). Define the Radon-Nikodym derivative for a measure change from \mathbb{Q} to \mathbb{M} as

$$\frac{d\mathbb{M}}{d\mathbb{Q}} = e^{iu\sigma W_{\mathcal{T}_t} - iu\eta\mathcal{T}_t + \psi_D(u)\mathcal{T}_t}. \quad (\text{A.1})$$

From $E^{\mathbb{Q}}[x] = E^{\mathbb{M}}[x d\mathbb{Q}/d\mathbb{M}]$ it follows that

$$\begin{aligned} E^{\mathbb{Q}}[e^{iu(\sigma W_{\mathcal{T}_t} - \eta\mathcal{T}_t)}] &= E^{\mathbb{M}}[e^{iu(\sigma W_{\mathcal{T}_t} - \eta\mathcal{T}_t)} \frac{d\mathbb{Q}}{d\mathbb{M}}] = E^{\mathbb{M}}[e^{-\psi_D(u)\mathcal{T}_t}] \\ &= \mathcal{L}_{\mathcal{T}_t}^{\mathbb{M}}(\psi_D(u)). \end{aligned} \quad (\text{A.2})$$

The same argumentation applies for the models in (20), (22) and (24). Except for the stochastic jump intensity model, the measure change has implications for the drift of the activity rate process in (16). Wu (2008) shows that under \mathbb{M} , the activity rate process becomes

$$dv_t = (\kappa - \kappa^{\mathbb{M}}v_t)dt + \sigma_v\sqrt{v_t}dW_t^v, \quad (\text{A.3})$$

with $\kappa^{\mathbb{M}} = \kappa - iu\sigma\sigma_v\rho$.

To solve the Laplace transform in (A.2), Carr and Wu (2004) substitute for the stochastic clock, i.e.

$$\mathcal{L}_{\mathcal{T}_t}^{\mathbb{M}}(\psi_D(u)) = E^{\mathbb{M}}[e^{-\psi_D(u)\int_0^t v(s)ds}]. \quad (\text{A.4})$$

What is inside the expectation operator bears much resemblance with the pricing formula for a zero coupon bond. Since the activity rate is modeled as a square-root diffusion process, the solution approach from Cox et al. (1985) applies.¹³ Formally,

$$\mathcal{L}_{\mathcal{T}_t}^{\mathbb{M}}(\psi_D(u)) = e^{-a(t) - b(t)v_0}, \quad (\text{A.5})$$

where

$$b(t) = \frac{2\psi_D(u)(1 - e^{-\gamma t})}{2\gamma - (\gamma - \kappa^{\mathbb{M}})(1 - e^{-\gamma t})}, \quad (\text{A.6})$$

$$a(t) = \frac{\kappa}{\sigma_v^2} \left[2\log \left(1 - \frac{\gamma - \kappa^{\mathbb{M}}}{2\gamma}(1 - e^{-\gamma t}) \right) + (\gamma - \kappa^{\mathbb{M}})t \right], \quad (\text{A.7})$$

and

$$\gamma = \sqrt{(\kappa^{\mathbb{M}})^2 + 2\sigma_v^2\psi_D(u)}. \quad (\text{A.8})$$

¹³See also Ingersoll (1987).

B Option Pricing via Fourier-cosine Expansion

This appendix follows Fang and Oosterlee (2008). On a finite interval $[a, b]$, the cosine expansion of the density function of the log return s reads as

$$f(s) = \sum_{k=0}^{\infty} 'A_k \cos \left(k\pi \frac{s-a}{b-a} \right), \quad (\text{B.1})$$

where the coefficients A_k are given by

$$A_k = \frac{2}{b-a} \int_a^b f(s) \cos \left(k\pi \frac{s-a}{b-a} \right) ds. \quad (\text{B.2})$$

From the Fourier transform $\phi(u) = \int_{\mathbb{R}} e^{ius} f(s) ds \approx \int_a^b e^{ius} f(s) ds$ and the fact that $\text{Re} [e^{i\theta}] = \cos(\theta)$, it follows that

$$A_k \approx \frac{2}{b-a} \text{Re} \left[\phi \left(\frac{k\pi}{b-a}, s \right) e^{-i \frac{k\pi a}{b-a}} \right]. \quad (\text{B.3})$$

The risk-neutral price of a call option may be represented in terms of the density function, i.e.

$$c(s, t) = e^{-r_d(T-t)} \int_{\mathbb{R}} c(y, T) f(y | s) dy \approx e^{-r_d(T-t)} \int_a^b c(y, T) f(y | s) dy, \quad (\text{B.4})$$

where $c(\cdot)$ is the option payoff function, y is the log return at expiry and $f(y | s)$ is the conditional density of y given s . Substituting for the density, interchanging summation and integration, and truncating the summation at N yields

$$c(s, t) = e^{-r_d(T-t)} \sum_{k=0}^{N-1} ' \text{Re} \left[\phi \left(\frac{k\pi}{b-a} \right) e^{ik\pi \frac{s-a}{b-a}} \right] V_k, \quad (\text{B.5})$$

where $V_k = \frac{2}{b-a} \int_a^b c(y, T) \cos \left(k\pi \frac{y-a}{b-a} \right) dy$ are the cosine series coefficients of the terminal payoff $c(y, T)$. For a call option with payoff $c(y, T) = [K(e^y - 1)]^+$, they are given by

$$V_k = \frac{2}{b-a} K [\chi_k(0, b) - \psi_k(0, b)], \quad (\text{B.6})$$

where

$$\chi_k(c, d) = \frac{1}{1 + \left(\frac{k\pi}{b-a}\right)^2} \left[\cos \left(k\pi \frac{d-a}{b-a} \right) e^d - \cos \left(k\pi \frac{c-a}{b-a} \right) e^c + \frac{k\pi}{b-a} \sin \left(k\pi \frac{d-a}{b-a} \right) e^d - \frac{k\pi}{b-a} \sin \left(k\pi \frac{c-a}{b-a} \right) e^c \right] \quad (\text{B.7})$$

$$\psi_k(c, d) = \begin{cases} \frac{b-a}{k\pi} \left[\sin \left(k\pi \frac{d-a}{b-a} \right) - \sin \left(k\pi \frac{c-a}{b-a} \right) \right] & k \neq 0 \\ d - c & k = 0 \end{cases} \quad (\text{B.8})$$

C One-Touch Option Prices when Currency Pairs Follow a Geometric Brownian Motion

Consider the model in (1) with convexity adjustment $\eta = \psi_x^{BS}(-i) = -\frac{1}{2}\sigma^2$, i.e.

$$S_t = S_0 e^{(r_d - r_f - \frac{1}{2}\sigma^2)t + \sigma W_t}. \quad (\text{C.1})$$

Define a drift variable $\theta = \frac{1}{\sigma}(r_d - r_f - \frac{1}{2}\sigma^2)$ and set

$$S_t = S_0 e^{\sigma \hat{W}_t}, \quad (\text{C.2})$$

where $\hat{W}_t = \theta t + W_t$ is a Brownian motion with drift. Consider a one-touch with upper barrier level B. The first passage time to B is defined as

$$\tau_B = \inf\{u \geq 0 : S_u = B\} = \inf\{t \geq 0 : W_t + \theta t = \frac{1}{\sigma} \log\left(\frac{B}{S_0}\right)\}. \quad (\text{C.3})$$

The first passage time τ_B has a density (see Wystup, 2010 and the reference therein)

$$P[\tau_B \in dt] = \frac{\frac{1}{\sigma} \log\left(\frac{B}{S_0}\right)}{t\sqrt{2\pi t}} \exp\left(-\frac{\left(\frac{1}{\sigma} \log\left(\frac{B}{S_0}\right) - \theta t\right)^2}{2t}\right) dt. \quad (\text{C.4})$$

From (31), the domestic price of a one-touch is given by

$$G(S_u, B, T) = e^{-r_d T} \int_0^T \frac{\frac{1}{\sigma} \log\left(\frac{B}{S_0}\right)}{t\sqrt{2\pi t}} \exp\left(-\frac{\left(\frac{1}{\sigma} \log\left(\frac{B}{S_0}\right) - \theta t\right)^2}{2t}\right) dt. \quad (\text{C.5})$$

Weber and Wystup (2010) evaluate this integral and derive a general formula,

$$G(S, B, T) = e^{-r_d T} \left((B/S)^{\frac{\theta+|\theta|}{\sigma}} \mathcal{N}(-\phi d_+) + (B/S)^{\frac{\theta-|\theta|}{\sigma}} \mathcal{N}(\phi d_-) \right), \quad (\text{C.6})$$

where $d_{\pm} = \frac{1}{\sigma\sqrt{t}}(\pm \log(S/B) - \sigma|\theta|t)$. For quotes under the foreign risk-neutral measure, C.6. is applied with reciprocal values for B and S, interest rates exchanged and ϕ replaced by $-\phi$ (See Weber and Wystup (2010) and Wystup (2010)).

D Simulating Jumps

A general reference for simulating jump trajectories is Cont and Tankov (2004).

Variance Gamma Model

Set $\sigma = \sqrt{2C/(MG)}$, $\theta = C(1/M - 1/G)$ and $\nu = 1/C$. For an arbitrary time interval Δt , simulate a gamma (Z_1) and a normal (Z_2) random variable

$$\begin{aligned} Z_1 &\sim \Gamma(\Delta t/\nu, \nu) \\ Z_2 &\sim N(0, 1) \end{aligned}$$

An increment of the variance gamma process is obtained from

$$\Delta X = \theta Z_1 + \sigma \sqrt{Z_1} Z_2. \quad (\text{D.1})$$

Kou Model

Define $\Upsilon \in \{M, G\}$. For an arbitrary time interval Δt , the intensity parameter λ is given by

$$\lambda = (I^+ + I^-)\Delta t, \quad (\text{D.2})$$

where $I^\pm = C/\Upsilon$ follows from evaluating the integral $\int v(x)dx$ over the positive and negative half-line of the Kou Lévy density in (8).

Given an interval Δt , generate $N \sim Poi(\lambda)$ jump interval times. Next, set the dampening factor $\Upsilon = M$ and the jump sign $\eta = 1$. Simulate a uniform variable U_1 and reset $\Upsilon = G$, $\eta = -1$ if $U_1 < I^-/(I^+ + I^-)$.

The probability that a jump with sign η has size $|x| \geq \varepsilon$ is

$$V(\varepsilon) = C \int_{\eta\varepsilon \wedge 0}^{0 \vee \eta\varepsilon} e^{-\Upsilon|x|/I^\eta} = 1 - e^{-\eta\Upsilon\varepsilon}, \quad (\text{D.3})$$

By inversion of (D.3), the jump magnitude is

$$\varepsilon = -\log(U_2)/(\eta\Upsilon), \quad (\text{D.4})$$

where U_2 is another uniform random variable. An increment on the time-grid is given as the sum of the jumps on the interval Δt ,

$$\Delta X_t = \sum_{i=1}^N \varepsilon_i. \quad (\text{D.5})$$

References

- Andersen, Leif, 2008. "Simple and efficient simulation of the Heston stochastic volatility model." *Journal of Computational Finance* 11, 1-42.
- Bakshi, Gurdip, Cao, Charles and Chen, Zhiwu, 1997. "Empirical performance of alternative option pricing models." *The Journal of Finance* 52, 2003-2049.
- Bates, David S., 1996. "Jumps and stochastic volatility: exchange rate processes implicit in Deutsche mark options." *The Review of Financial Studies* 9, 69-107.
- Bates, David S., 2000. "Post-87' crash fears in the S&P 500 future option market." *Journal of Econometrics* 94, 181-238.
- Beaglehole, David R., Dybvig Philipp H. and Zhou, Guofu, 1997. "Going to extremes: Correcting simulation bias in exotic option valuation." *Financial Analyst Journal*, 62-68.
- Bertoin, Jean, 1996. "Lévy processes." Cambridge University Press, Cambridge.
- Black, Fischer and Scholes, Myron, 1973. "The pricing of options and corporate liabilities." *Journal of Political Economy* 81, 637-654.
- Broadie, Mark and Kaya, Oezguer, 2006. "Exact simulation of stochastic volatility and other affine jump diffusion processes." *Operations Research* 54, 217-231.
- Carr, Peter, Geman, Hélyette, Madan, Dilip and Yor, Marc, 2002. "The fine structure of asset returns: an empirical investigation." *Journal of Business* 75, 305-332.
- Carr, Peter, Geman, Hélyette, Madan, Dilip and Yor, Marc, 2003. "Stochastic volatility for Lévy processes." *Mathematical Finance* 13, 345-382.
- Carr, Peter and Madan, Dilip B., 1999. "Option valuation using the fast Fourier transform." *Journal of Computational Finance* 2, 61-73.
- Carr, Peter and Wu, Liuren, 2003. "The finite moment log stable process and option pricing." *The Journal of Finance* 58, 753-777.
- Carr, Peter and Wu, Liuren, 2004. "Time-changed Lévy processes and option pricing." *Journal of Financial Economics* 71, 113-141.
- Carr, Peter and Wu, Liuren, 2007. "Stochastic skew in currency options." *Journal of Financial Economics* 86, 213-247.

- Cont, Rama and Tankov, Peter, 2004. "Financial modelling with jump processes." Chapman & Hall/CRC, Boca Raton.
- Cox, John C., Ingersoll, Jonathan E. and Ross, Stephen A., 1985. "A theory of the term structure of interest rates." *Econometrica* 53, 385-407.
- Daal, Elton A. and Madan, Dilip B., 2005. "An empirical investigation of the variance-gamma model for foreign currency options." *Journal of Business* 78, 2121-2152.
- Fang, Fang and Oosterlee, Kees, 2008. "A novel pricing method for European options based on Fourier-cosine series expansions." *SIAM Journal on Scientific Computing* 31, 826-848.
- Garman, Mark B. and Kohlhagen, Steven W., 1983. "Foreign currency option values." *Journal of International Money and Finance* 2, 231-237.
- Glasserman, Paul, 2004. "Monte carlo methods in financial engineering." Springer Finance, New York.
- Heston, Steve L., 1993. "A closed-form solution for options with stochastic volatility with applications to bond and currency options." *The Review of Financial Studies* 6, 327-343.
- Higham, Desmond J. and Mao, Xuerong, 2005. "Convergence of Monte Carlo simulations involving the mean-reverting square root process." *Journal of Computational Finance* 8, 35-61.
- Huang, Jin-Zhi and Wu, Liuren, 2004. "Specification analysis of option pricing models based on time-changed Lévy processes." *The Journal of Finance* 59, 1405-1439.
- Ingersoll, Jonathan E., 1987. "Theory of financial decision making." Rowman & Littlefield Publishing Inc.
- Kou, Steven G., 2002. "A jump-diffusion model for option pricing." *Management Science* 48, 1086-1101.
- Lord, Roger, Koekkoek, Remmert and van Dijk, Dick, 2010. "A comparison of biased simulation schemes for stochastic volatility models." *Quantitative Finance* 10, 177-194.
- Madan, Dilip B. and Seneta, Eugene, 1990. "The variance gamma (V.G.) model for share market returns." *Journal of Business* 63, 511-524.

- Madan, Dilip B., Carr, Peter P. and Chang, Eric C., 1998. "The variance gamma process and option pricing." *European Finance Review* 2, 79-105.
- Merton, Robert C., 1976. "Option pricing when underlying stock returns are discontinuous." *Journal of Financial Economics* 3, 125-144.
- Newey, Whitney K. and West, Kenneth D., 1987. "A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix." *Econometrica* 55, 703-708.
- Reiswich, Dimitri and Wystup, Uwe, 2010. "A guide to FX options quoting conventions." *The Journal of Derivatives* 18, 58-68.
- Ribeiro, Claudia and Webber, Nick, 2006. "Correcting for simulation bias in Monte Carlo methods to value exotic options in models driven by Lévy processes." *Applied Mathematical Finance* 13, 333-352.
- Weber, Andreas and Wystup, Uwe, 2010. "Pricing formulae for foreign exchange options." In: *Encyclopedia of Quantitative Finance*, John Wiley & Sons Ltd. Chichester, UK, 1408-1418.
- Wu, Liuren, 2006. "Dampened power law: reconciling the tail behavior of financial security returns." *Journal of Business* 79, 1445-1473.
- Wu, Liuren, 2008. "Modeling financial security returns using Lévy processes." In: Birge, John R. and Linetsky, Vadim (Eds.), *Handbooks in Operations Research and Management Science*, Volume 15, Financial Engineering, Elsevier, North-Holland, 117-162.
- Wystup, Uwe, 2010, "Foreign exchange symmetries." In: *Encyclopedia of Quantitative Finance*, John Wiley & Sons Ltd. Chichester, UK, 752-759.

Curriculum Vitae

Personal Information

Name: Ralf Martin Büsser
Date of birth: 25 May 1980
Place of origin: Amden (St. Gallen)
Citizenship: Swiss

Professional Experience

2012 - **Swiss National Bank, Zurich, Switzerland**
Money Market and Foreign Exchange

2008 - 2011 **Algofin AG, St. Gallen, Switzerland**
Financial Markets Consulting and Research

2005 - 2008 **UBS Investment Bank, Opfikon, Switzerland**
Foreign Exchange (FX) Exotic Options Trading,
FX Structuring and Solutions

Education

2008 -2012 **University of St. Gallen, Switzerland**
PhD in Finance

2008 - 2010 Chartered Financial Analyst (CFA)

2000 - 2005 **University of St. Gallen, Switzerland**
Master of Arts in Quantitative Economics and Finance
(M.A. HSG)

2003 **Stockholm School of Economics, Sweden**

1993 - 1999 **Gymnasium Friedberg, Gossau SG, Switzerland**
Matura