

School of Finance



University of St.Gallen

VARIANCE RISK PREMIUMS IN FOREIGN EXCHANGE MARKETS

MANUEL AMMANN

RALF BUESSER

WORKING PAPERS ON FINANCE No. 2013/4

SWISS INSTITUTE OF BANKING AND FINANCE (S/BF – HSG)

APRIL 2013



Variance Risk Premiums in Foreign Exchange Markets

Manuel Ammann,^a Ralf Buesser^b

This version: April 20th, 2013

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.

Keywords: Foreign exchange, variance risk premium, variance swap, intraday data, risk-neutral expectation, jump risk

JEL Classification: C12, C13, F31, G12, G13, G15

^aSwiss Institute of Banking and Finance, University of St. Gallen, Rosenbergstrasse 52, CH-9000 St. Gallen, Switzerland. E-mail: manuel.ammann@unisg.ch

^bThis research was conducted while the author was affiliated with the University of St. Gallen, Dufourstrasse 50, CH-9000 St. Gallen, Switzerland. The author may be reached by phone: +41 78 866 22 14, or e-mail: ralf.buesser@alumni.unisg.ch. (corresponding author)

We thank Paul Söderlind, an anonymous referee and the associate editor for helpful comments and suggestions.

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 along the lines 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 approach of Britten-Jones and Neuberger (2000) 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 relation-

ship 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

¹We refer to Appendix A for a derivation.

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 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 timestamp on the option quotes, whereas the latter is identical to New York cut, i.e. the time when the options expire.

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.

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

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

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

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.

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).

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

⁶We use the term biproduct synonym for currency pair.

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 Reiswich and Wystup (2010).

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

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

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-modeling accuracy. Figure 1 illustrates this principle.

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 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.

⁷Note that the vanna of an option is again dependent on the delta convention. Some details have been relegated to Appendix C.




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.

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

		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
	<i>tstat</i>	-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
	<i>tstat</i>	-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
	<i>tstat</i>	-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
	<i>tstat</i>	-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
	<i>tstat</i>	-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
	<i>tstat</i>	-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.

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).

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 loose money on average to protect against a rise in variance. We test this hypothesis over the period

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

from January 2003 to August 2008, again for all biproducts.

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

	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.

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 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

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

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: 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.

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 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




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.

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

$$RV_t = \alpha + \beta RNV_t + \varepsilon_t. \quad (11)$$

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

	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
<i>tstat</i>	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
<i>tstat</i>	-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.

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 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

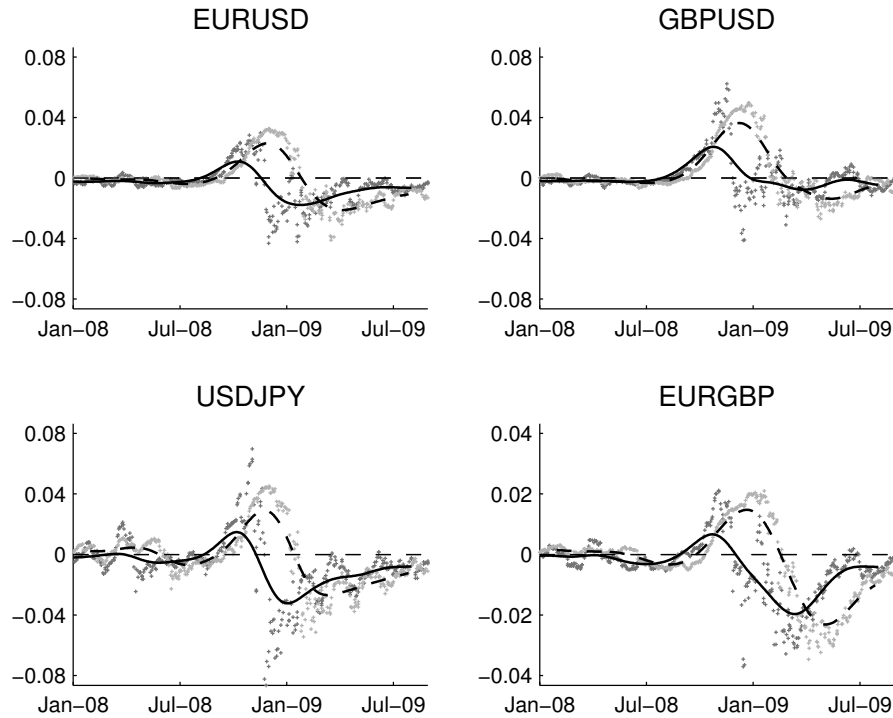


Figure 4: Time-Variation in Variance Risk Premiums

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 month.

reported dips into negative territory in Summer 2008 are hardly visible. At the onset 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	$\bar{L}RP$	-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	$\bar{L}RP$	-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	$\bar{L}RP$	-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, computed as the difference of the logarithm of the VIX series. The VIX reveals whether variance risk premiums in currency markets are attributable to the overall level of market volatility. Finally, we add the first difference of the 3-month TED spread series. 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. All data is obtained from Bloomberg.

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. The sampling frequency t in (12) is daily. To cope with the overlapping time intervals, 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 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

	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.

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

¹⁰We thank an anonymous referee for pointing this out.

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 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.

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

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 $E_t^{\mathbb{Q}}()$ 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

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

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

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 April 2013.

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

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

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

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

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).

C 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} ,

$$\mathbf{y} = \begin{pmatrix} \frac{\partial c(K, T, \sigma^K)}{\partial \sigma} \\ \frac{\partial^2 c(K, T, \sigma^K)}{\partial S \partial \sigma} \\ \frac{\partial^2 c(K, T, \sigma^K)}{\partial \sigma^2} \end{pmatrix}, \quad (\text{C.1})$$

and

$$\mathbf{A} = \begin{pmatrix} \frac{\partial c(K_{xDP}, T, \sigma^{xDP})}{\partial \sigma} & \frac{\partial c(K_{DN}, T, \sigma^{DN})}{\partial \sigma} & \frac{\partial c(K_{xDC}, T, \sigma^{xDC})}{\partial \sigma} \\ \frac{\partial^2 c(K_{xDP}, T, \sigma^{xDP})}{\partial S \partial \sigma} & \frac{\partial^2 c(K_{DN}, T, \sigma^{DN})}{\partial S \partial \sigma} & \frac{\partial^2 c(K_{xDC}, T, \sigma^{xDC})}{\partial S \partial \sigma} \\ \frac{\partial^2 c(K_{xDP}, T, \sigma^{xDP})}{\partial \sigma^2} & \frac{\partial^2 c(K_{DN}, T, \sigma^{DN})}{\partial \sigma^2} & \frac{\partial^2 c(K_{xDC}, T, \sigma^{xDC})}{\partial \sigma^2} \end{pmatrix}. \quad (\text{C.2})$$

The vega and volga of an option are given by

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

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

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

$$\frac{\partial^2 c(\cdot)^{(reg.)}}{\partial \sigma \partial S} = -e^{-r_f \tau} N'(d_1) \frac{d_2}{\sigma} \quad \frac{\partial^2 c(\cdot)^{(p.a.)}}{\partial \sigma \partial S} = \frac{\partial^2 c(\cdot)^{(reg.)}}{\partial \sigma \partial S} \frac{K}{F}, \quad (\text{C.4})$$

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

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

$$\mathbf{c}^{\text{me}} = \begin{pmatrix} c(K_{xDP}, T, \sigma^{xDP}) - c(K_{xDP}, T, \sigma^{BS}) \\ c(K_{DN}, T, \sigma^{DN}) - c(K_{DN}, T, \sigma^{BS}) \\ c(K_{xDC}, T, \sigma^{xDC}) - c(K_{xDC}, T, \sigma^{BS}) \end{pmatrix}. \quad (\text{C.5})$$

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

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: Intraday 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.