How Does the Volatility Risk Premium Affect the Informational Content of Currency Options?

Peter Breuer*

Research Department
International Monetary Fund
700 19th Street NW
Washington, DC 20431
United States of America
Tel.: +1-202-623-6364
Fax: +1-202-589-6364
e-mail: pbreuer@imf.org

Abstract

Interpretating probability density functions (PDFs) extracted from currency options data is ambiguous because PDFs combine risk neutral market views regarding the likelihood of particular exchange rate outcomes with investors' preferences towards risk. In order to disentangle the two effects, market expectations derived for option prices need to adjusted for the time-varying volatility risk premium that compensates risk averse option writers. Assuming rational expectations this risk premium can be extracted ex-post. The implied volatility bid-ask spread and volatility of implied volatility are considered here as proxies for the risk premium to enable the ex-ante adjustment of risk-neutral exchange rate expectations for risk preferences. The method is applied to demonstrate the impact of this adjustment on exchange rate expectations around Hong Kong SAR's equity market intervention in 1998. The risk premium explains part of the bias found in existing empirical studies of the predictability of future realized volatility by implied volatility.

JEL Classification Numbers: F31, G13, G15

Keywords: Forecasting, implied volatility, risk premium, volatility smile

* The author thanks Ranjit Teja and Charles Kramer for sharing very useful computer programs. Expert programming and research assistance were rendered by Brian Slatas and Yoon Kim. William Lee and Prebone Yamane were very helpful in obtaining data. The paper benefited tremendously from the comments by Jahangir Aziz, Torbjørn Becker, Burkhard Drees, Guy Meredith, Garry Schinasi, Ranjit Teja, and seminar participants at the 56th European Meeting of the Econometric Society, Lausanne, and the 2001 Far Eastern Meeting of the Econometric Society, Kobe.

HOW DOES THE VOLATILITY RISK PREMIUM AFFECT THE INFORMATIONAL CONTENT OF CURRENCY OPTIONS?

Abstract

Interpretating probability density functions (PDFs) extracted from currency options data is ambiguous because PDFs combine risk neutral market views regarding the likelihood of particular exchange rate outcomes with investors' preferences towards risk. In order to disentangle the two effects, market expectations derived for option prices need to adjusted for the time-varying volatility risk premium that compensates risk averse option writers. Assuming rational expectations this risk premium can be extracted ex-post. The implied volatility bid-ask spread and volatility of implied volatility are considered here as proxies for the risk premium to enable the ex-ante adjustment of risk-neutral exchange rate expectations for risk preferences. The method is applied to demonstrate the impact of this adjustment on exchange rate expectations around Hong Kong SAR's equity market intervention in 1998. The risk premium explains part of the bias found in existing empirical studies of the predictability of future realized volatility by implied volatility.

JEL Classification Numbers: F31, G13, G15

Keywords: Forecasting, implied volatility, risk premium, volatility smile

I. Introduction

Derivative instruments permit the extraction of market expectations for future asset price realizations. Forecasts based on current market information impounded in securities prices are founded in the notion that market expectations are embedded in an up-to-date way in current spot and derivatives prices. The price of derivative securities as state contingent claims depends in part on the perceived probability that one of these states will occur. In other words, if participants are suddenly convinced that the price of a particular security is likely to drop and they would like to insure themselves against losses stemming from such potential losses, the premium they need to pay for this insurance will rise. For example, the premium of a put option would rise.[1] Thus, the cost of insurance against price movements beyond different levels permits inferring the probability the market as a whole assigns to each of these events occurring.

Existing methods to extract probability density functions (PDF) from options data provide risk neutral frequency distributions of potential future outcomes. However, risk averse investors will demand a risk premium, driving a "wedge" between extracted risk neutral and actual underlying probabilities assigned by investors (see Appendix for an explanation of risk neutrality). The existence of a risk premium implies that the changing price of a derivative instrument may either reflect a change in the perceived underlying risk or a change in risk preferences towards holding such risk (or a combination of both). To the extent that the extracted risk-neutral PDFs are used as the "market's forecast" of a financial variable over a given time horizon, the user of such forecasts has to be aware that they are not a picture of actual expectations.

In this paper, a risk premium is extracted from over-the-counter currency option pricing information. Similar to the forward market, the risk premium drives a wedge between risk neutral expectations and future realized outcomes. The risk premium derived from options differs from the forward risk premium in that it is a risk premium on a higher moment, the volatility of asset prices (here currencies). The volatility risk premium is extracted ex-post and approximated with ex-ante available proxies in order to allow adjusting expectations ex-ante. Adjusting for the presence of a risk premium reduces some of the bias displayed in previous studies testing the predictability of actually realized volatility, using implied volatility as a forecast variable. Risk-neutral PDFs extracted from currency option prices are then adjusted for the volatility risk premium. The adjusted PDFs reflect the market's expectations of future exchange rate outcomes rather than the market's preferences towards holding the risks associated with these outcomes.

The paper is organized as follows. Apart from attempting to adjust risk-neutral expectations for a risk premium, this paper also contributes to the volatility forecasting literature because

[1] Note that the actual probabilities of a particular event occurring need not change for a change in the cost of insurance to occur.

it is the first to test predictability of volatility using over-the-counter currency option data (section II). The econometric model for estimating the volatility risk premia used in the adjustment is derived in section III. That section also approximates the ex-post observable risk premium with ex-ante instruments and re-estimates the original predictability equations with this adjustment. Section IV discusses adjusting probability density functions extracted from options data for the volatility risk premium. Section V presents an application of this method to market beliefs around the 1998 Hong Kong SAR intervention in equity markets. Section VI concludes.

II. Can Implied Volatility Accurately Forecast Future Realized Volatility?

Volatility implied by option prices is widely interpreted as the market's best forecast of future return volatility over the remaining life of the relevant option. If option markets are efficient and competitive, implied volatility should be an efficient forecast of future volatility as it impounds all information available at the time. If it were not, the argument goes, it would be possible to devise a profitable trading strategy that moves the option price to the level that reflects the best possible forecast of future volatility.

II.1. The Predictability Literature

A sizeable literature, however, suggests that implied volatility (derived under the assumption of risk-neutral investors) predicts future realized volatility poorly. Early papers (e.g. Latané and Rendleman, 1976, Chiras and Manaster, 1978, and Beckers, 1981) regress future realized volatility on weighted implied volatilities on a cross-section of stock options. These papers find that option prices contain volatility forecasts that are more accurate than historical measures, but are not reliable forecasts of actually realized volatility.

More recently, research has turned to the analysis of volatility in a time series framework using options written on only one underlying asset. Canina and Figlewski (1993) regress the volatility of the S&P 100 index against the implied volatility of index options between 1983 and 1986. They report that implied standard deviations (ISDs) have little predictive power for future volatility, and are significantly biased forecasts. Day and Lewis (1992) study S&P 100 index options with expiries between 1985 and 1989. Lamoureux and Lastrapes (1993) focus on individual stock options between 1982 and 1984. Both studies find that implied volatility is biased and inefficient: historical volatility of the underlying instrument contains predictive information about future volatility beyond that contained in implied volatilities. Other papers find less extreme results, but confirm the low predictive power of implied volatility. Jorion (1995), using data from exchange-traded currency futures options on the dollar exchange rates vis-à-vis the yen, mark and Swiss franc, finds that while ISDs outperform statistical time-series models, they appear to be biased volatility forecasts. Even the study most supportive of the forecasting power of implied volatilities, Christensen and Prabhala (1998), cannot reject the claim that implied volatilities are at best biased forecasts of future realized volatility.

There are at least three possible explanations for the bias found consistently in studies of the predictive power of implied volatility: (1) studies may have employed flawed econometric methods; (2) the market is inefficient; (3) the existence of a volatility risk premium biases the results.

Econometric misspecifications could bias results of existing studies through overlapping error terms or through orthogonal measurement error in the forecasting variable. Most time-series studies employ daily data from options with maturities over several months. The resulting overlap in the error terms biases the OLS standard errors downward. While the point estimates of the coefficients are still consistently estimated, the downward bias would lead to a rejection of the null hypothesis in instances where in fact one should fail to reject it. Most studies dealt with this problem by employing variants of Hansen's (1982) Generalied Method of Moments. The second problem could arise if there is an orthogonal measurement error in the implied volatility variable that biases its slope coefficient downward. Christensen and Prabhala (1998) point to non-synchronous measurement of option prices and index levels, early exercise and dividends, and other factors as possible causes for the presence of an orthogonal error. None of these factors is likely to affect the measurement of implied volatility in currency options.

The market inefficiency argument holds that options are not priced competitively. All studies find that the slope coefficient of implied volatility in the regression explaining actually realized volatility (often with some other regressors) is significantly less than unity and often indistinguishable from zero. Thus, along with a non-negative constant term the volatility assumed to prevail over the lifetime of the option (the implied volatility) consistently exceeds the volatility that actually realizes over the same period (the realized volatility). Since the price of an option is an increasing function of the implied volatility, this would imply that option buyers are consistently paying too much for their options. Figlewski (1997) points to frictions in the options markets in carrying out the necessary arbitrage strategies that may prevent the prices to move to their proper levels. This argument is more likely to hold in the case of index options, which involve recreating the index with several hundred underlying stocks. In the case of currency options with a single and frequently traded underlying instrument, it is less likely that such frictions exist. Another explanation could be collusive and anti-competitive behavior by option traders. Given the presence of many traders in most (currency) option markets the existence of such behavior is possible but unlikely.

The observed excess of implied over realized volatility could also be due to a volatility risk premium. Unlike the option buyer, the option writer faces an unlimited loss potential. The volatility assumed to prevail over the lifetime of the option—fixed at inception—could in fact be incorrect, adversely affecting the option writer's position. If option writers are risk averse they need to be compensated for taking on volatility risk. The associated risk premium is implicit in the observed option price and raises the level of the implied volatility compared to that charged by risk neutral option writers. To the extent that the model used to extract implied volatility from option prices does not take into account the volatility risk premium, the implied volatility will appear as a biased estimator of realized volatility. Several studies (e.g. Heston (1993), Green and Figlewski (1999), Buraschi and Jackwerth (1999)) have emphasized that the market price of variance risk may not be zero. Poteshman (2000) uses Heston's (1993) stochastic volatility model to extract implied volatility data from observed option prices on the S&P 500 index. Despite taking into account a non-zero market price of volatility risk in the extraction of implied volatility, the econometric model used to forecast volatility omits the volatility risk premium, causing the forecasting power of implied volatility to be rather low, in fact less than some historical predictors.

The present paper argues that in existing studies an orthogonal measurement error is present due to the misspecification of the econometric model by omitting the volatility risk premium in the predictability regression. Like the risk premium in the forward market that drives a wedge between the forward rate and the anticipated rate of domestic currency depreciation, the volatility risk premium has to be deducted from the observed implied volatility to arrive at the anticipated future realized volatility. By using data from over-the-counter market quoted in implied volatility we avoid problems related to the selection of the correct option pricing model to extract implied volatilities. Using currency options data avoids problems found in equity markets, which relate to various features of the underlying instrument that could cause the presence of an orthogonal error in the independent variable.

II.2. Method to Forecast Future Realized Volatility

A rational forecast of a variable is the expected value of the variable conditional on the available information set. By definition, the realized value of a stochastic variable is its rational expected value plus a zero mean random disturbance. If $\hat{\sigma}$ is to be a rational forecast of future realized volatility, σ_{RV}, it must be the conditional expected value of future realized volatility, given the market's information:

$$\sigma_{RV} = \hat{\sigma} + \varepsilon$$

with

$$E[\varepsilon] = 0; \qquad E[\hat{\sigma}\varepsilon] = 0$$

Different information sets will give rise to different conditional expectations. A better information set will yield a more accurate forecast by reducing the variance of the forecast error ε. Nevertheless, this specification must hold for every rationally formed forecast.

This approach leads to the following regression model. The predictive power of a volatility forecast can be estimated by regressing the realized volatility on forecast volatility:

$$\sigma_{t,T} = a + b\hat{\sigma}_t + \varepsilon_{t,T} \tag{1}$$

where $\sigma_{t,T}$ is the future realized volatility over the remaining life of the contract, measured as the annualized standard deviation of returns from day t to day T, and $\hat{\sigma}_t$ is the volatility forecast available on day t.[2] The typical test of market efficiency interprets the ISD as the volatility forecast and tests whether the intercept is equal to zero and the slope coefficient

[2] Since we are examining over-the-counter options there is no problem with shrinking time to maturity, as options commence their life on the day they are traded and mature at a fixed horizon that moves forward every day.

equal to unity. This specification compares the implied volatility of the option with the volatility that is actually realized ex-post. Since the implied volatility can be interpreted as the volatility that is assumed to prevail over the lifetime of the option contract and is an important determinant of its price, this estimation can be interpreted as a test of the accuracy of option pricing.

This framework can also be used to compare the predictive power of the ISD with that of other forecasts, such as time-series models using historical information. Here, we will simply compare the predictive power of implied volatility to that of realized historical volatility observed on the day of inception of the option, measured over the same time horizon as the volatility to be predicted:

$$\sigma_{t,T} = a + b_1\sigma_t^{ISD} + b_2\sigma_t^{HV} + \varepsilon_{t,T}, \tag{2}$$

where we expect the historical volatility, σ_t^{HV}, to have no incremental predictive power beyond that of implied volatility, σ_t^{ISD}. Hence, we expect the coefficient b_2 to be close to zero. This specification relies only on most recently available information, ignoring observations further in the past, as would be used in a time series specification.

Due to the relative short time series, it is necessary to use daily data. Overlapping error terms cause a downward bias in the usual ordinary least squares (OLS) standard errors. While the point estimates of the coefficients are still consistently estimated, the downward bias would lead to a rejection of the null hypothesis in instances where in fact we should fail to reject it.

We follow the GMM method suggested by Hansen (1982), providing a method to deal with heteroscedasticity and serial correlation. The variance-covariance matrix of the estimated coefficients is given by

$$\Sigma = (\mathrm{X'X})^{-1}\Omega\,(\mathrm{X'X})^{-1}, \tag{3}$$

where $\Omega = E[\mathrm{X}'\varepsilon\varepsilon'\mathrm{X}/T]$ is consistently estimated, using the OLS residuals $\hat{\varepsilon}$, by

$$\hat{\Omega} = \sum_t \hat{\varepsilon}_t^2 \mathrm{X'}_t \mathrm{X}_t + \sum_t \sum_t Q(s,t)\hat{\varepsilon}_s\hat{\varepsilon}_t(\mathrm{X'}_t\,\mathrm{X}_s + \mathrm{X'}_s\,\mathrm{X}_t), \tag{4}$$

with *Q(s, t)* defined as the indicator function equal to unity if there is overlap between returns at *s* and *t*, and zero otherwise. In cases where the residuals are homoscedastic and do not overlap, $E[\varepsilon_t^2] = s^2$, and *Q(s, t)* is always zero, so that the covariance matrix collapses to the OLS covariance matrix $s^2(\mathrm{X'X})^{-1}$. Simulations in Canina and Figlewski (1993) and Jorion (1995) indicate that in the absence of error in the variables, the GMM corrections make it unlikely that any bias would be caused by faulty standard error calculations.

II.3. Data

The data used in the present paper add to the predictability literature in two ways: (1) it employs over-the-counter (OTC) data; and (2) it presents currency option data over different maturity horizons. The data were obtained from the OTC market for currency options, as captured by the J.P. Morgan FX Options Trading System and by Prebone Yamane, a large currency option broker. The currencies covered are the U.S. dollar rates against the Japanese yen (January 25, 1996 to January 19, 2001), the euro (January 1, 1998 to January 25, 2001), the Thai baht (July 1, 1997 to May 2, 2000) and the Hong Kong Dollar (July 1, 1997 to May 2, 2000). Implied volatilities are for at-the-money forward options. Spot exchange rates were collected simultaneously. Realized volatility is the standard deviation of daily returns, scaled by $\sqrt{261}$ to annualize[3], over the remainder of the lifetime of the option.

Dealers in the OTC market quote option prices in terms of implied volatility, rather than in absolute currency units, as is the case on exchanges. When a deal is struck, they translate the agreed implied volatility to currency units by plugging it—along with other defining characteristics (e.g. interest rates, time to maturity, etc.)—into the Garman-Kohlhagen currency option pricing formula (a Black-Scholes variant). Dealers are fully aware that the option pricing model may by subject to misspecification error. The model is merely used as a vehicle to go back and forth between two ways of quotation. Implied volatility is a convenient way of comparing prices across options with differing features, such as varying strike prices or maturities. Every trader makes decisions at what level of implied volatility to buy and sell options based on an internal 'model', which may not necessarily be a formal model.

Using OTC implied volatilities as predictor for actually realized volatility has the advantage that no option pricing model needs to be specified. The OTC data are observed in the market place and reflect aggregate views about the future regardless of the model used to form these views. If an individual's view differs from that of the market he can take a position at the prevailing market price. Thus, this data is not subject to the assumptions inherent in the option price models used in previous studies to extract implied volatilities from observed option prices. This obviates the need for a stochastic volatility model, which Poteshman (2000) points out is necessary for exchange-traded data.

II.4. Estimating Predictive Power of Volatility Forecasts

Using the above method and data, equations (1) and (2) were estimated. The findings confirm that implied volatility by itself is a biased predictor of future realized volatility (Tables 1-4). The slope coefficient for implied volatility in regressions for the euro, baht and Hong Kong dollar are all significantly less than unity. The coefficient is close to unity only in the case of some yen options. At the same time the constant term is positive in almost all cases. This suggests that implied volatility consistently exceeds realized volatility and confirms the bias found in the literature. As none of the intercepts are equal to zero as expected under the null hypothesis, there may be a missing variable problem.

---

[3] Currencies are traded on more days per year than other underlying instruments, such as equity, necessitating a slightly larger scaling factor than is used for those assets.

Comparing the results with Jorion's (1995) we notice that the adjusted R^2 in the current study tends to be higher. The highest R^2 reported in Jorion's study was 0.16 (for the German mark), compared to 0.77 (for the euro) in this study. In general, the explanatory power of OTC implied volatility seems to be much greater than that of exchange-traded options. This could be due to the noise introduced by implicit assumptions in the pricing model used to extract implied volatilities from exchange traded data, such as the lack of a non-zero market price of variance risk.

While implied volatility explains a large amount of the variation of realized volatility of the Hong Kong Dollar, its slope coefficients are very small. This, together with the positive intercept, suggests that implied volatility exceeds realized volatility by far more in the case of the Hong Kong Dollar than in any of the other currencies. This is largely a result of the success of the peg, which has meant that the currency remained extremely stable with very little volatility, whereas implied volatility varies with market sentiment and reflects the implied risk of the peg breaking. In the flexible exchange rate environment of the other currencies implied volatility is a better forecast of future realized volatility.

Observed historical volatility on the day of inception of the option over horizons equivalent to those of the options appears to be a better predictor in individual regressions than implied volatility. This observation does not hold for the euro at short maturities and the yen at all maturities. Moreover, this result is not robust to the inclusion of both (implied and historical volatility) variables on the right hand side. When regressing future realized volatility on implied and historical volatility, the latter does not have much additional explanatory power in most regressions.

The explanatory power of ISD seems to decline in the maturity of the options. In addition, the constant is larger as the maturity of the forecasting horizon increases. This observation is consistent with the hypothesis that forecasts are bound to be less accurate over a longer horizon than over a shorter horizon.

III. The Volatility Risk Premium

While arbitrage based option pricing models, such as Garman-Kohlhagen, are consistent with any set of risk preferences, observed equilibirum option prices may contain risk premia that compensate risk averse investors. In particular, the option-pricing model requires an estimate of expected volatility as an input that may depend on risk preferences. The observed market price is the outcome of a trade between traders and embodies their risk preferences. In return for providing insurance, the observed option premium compensates the option writer for at least two types of risk: (1) the underlying price risk predicated on the assumption of constant volatility over the lifetime of the option; and (2) the risk that volatility may change over the lifetime of the option. While the underlying price risk can be hedged by dynamically replicating the option, the volatility risk cannot. A risk averse option writer will demand a risk premium for taking on this risk and would thus charge a higher option price (i.e. higher implied volatility) than a risk neutral option writer. Option writers need to be compensated for taking on volatility risk because they—unlike option buyers—are exposed to an unlimited loss potential. Since option prices in the over-the-counter market are quoted in terms of implied volatility rather than currency units the volatility risk premium can be inferred ex-post under the assumption of rational expectations as the difference between implied volatility and actually realized volatility.

Table 1. Thai Baht: Dependent Variable is Future Realized Volatility

Maturity	a	Slopes On		Adj. R^2	Observations
		ISD++	HV		
1 month	6.74*	0.51+		0.45	305
	(1.30)	(0.03)			
	4.19*		0.66*	0.44	1133
	(0.42)		(0.02)		
	7.15*	0.53+	-0.03	0.45	305
	(1.35)	(0.3)	(0.03)		
2 month	12.54*	0.38+		0.26	305
	(1.41)	(0.03)			
	5.88*		0.56*	0.32	1107
	(0.46)		(0.03)		
	14.61*	0.40+	-0.09	0.27	305
	(1.63)	(0.04)	(0.04)		
3 month	17.77*	0.24+		0.09	305
	(1.53)	(0.04)			
	7.42*		0.49*	0.24	1062
	(0.52)		(0.03)		
	23.18*	0.26+	-0.19*	0.13	305
	(2.07)	(0.04)	(0.05)		
6 month	26.89*	-0.05+		0.01	305
	(1.58)	(0.05)			
	11.09*		0.37*	0.15	932
	(0.60)		(0.03)		
	-10.99*	0.02	1.05*	0.18	305
	(4.98)	(0.04)	(0.13)		
1 year	26.06*	-0.17+		0.06	305
	(1.34)	(0.04)			
	28.08*		-0.42*	0.20	671
	(0.80)		(0.03)		
	60.15*	0.14+	-1.45*	0.60	305
	(1.90)	(0.03)	(0.07)		

* Significantly different from zero at the 5 percent level.
+ Significantly different from unity at 5 percent level.
++ Implied volatility at ask price.

Table 2. Hong Kong Dollar: Dependent Variable is Future Realized Volatility

Maturity	a	Slopes On		Adj. R^2	Observations
		ISD++	HV		
1 month	0.07*	0.07+		0.36	604
	(0.01)	(0.03)			
	0.16*		0.23*	0.09	958
	(0.01)		(0.02)		
	0.08*	0.08+	-0.11*	0.38	604
	(0.02)	(0.01)	(0.03)		
2 month	0.10*	0.06+		0.40	604
	(0.01)	(0.01)			
	0.14*		0.35*	0.20	958
	(0.01)		(0.03)		
	0.09*	0.05+	0.05	0.41	604
	(0.01)	(0.01)	(0.03)		
3 month	0.08*	0.05+		0.56	604
	(0.01)	(0.02)			
	0.13*		0.39*	0.24	958
	(0.01)		(0.02)		
	0.07*	0.05+	0.02	0.56	604
	(0.01)	(0.01)	(0..03)		
6 month	0.13*	0.03+		0.31	604
	(0.01)	(0.01)			
	0.15*		0.37*	0.21	925
	(0.01)		(0.02)		
	0.10*	0.03+	0.13*	0.33	604
	(0.02)	(0.01)	(0.04)		
1 year	0.13*	0.02+		0.19	578
	(0.01)	(0.01)			
	0.46*		-0.38*	0.16	795
	(0.13)		(0.03)		
	0.51*	0.02+	-0.93*	0.50	578
	(0.02)	(0.01)	(0.05)		

* Significantly different from zero at the 5 percent level.
+ Significantly different from unity at 5 percent level.
++ Implied volatility at ask price.

Table 3. Japanese Yen: Dependent Variable is Future Realized Volatility

Maturity	a	Slopes On ISD++	Slopes On HV	Adj. R^2	Observations
1 month	0.29	0.93+		0.40	1225
	(0.42)	(0.03)			
	5.88*		0.51*	0.26	1250
	(0.32)		(0.02)		
	0.35	0.90	0.02	0.40	1205
	(0.43)	(0.05)	(0.04)		
2 month	0.46	0.93+		0.38	1201
	(0.44)	(0.03)			
	5.26*		0.57*	0.32	1202
	(0.32)		(0.02)		
	1.26*	0.68+	0.19*	0.39	1159
	(0.45)	(0.05)	(0.04)		
3 month	0.95*	0.88+		0.36	1218
	(0.45)	(0.03)			
	5.00*		0.60*	0.36	1164
	(0.32)		(0.02)		
	2.68*	0.42+	0.36*	0.38	1154
	(0.48)	(0.06)	(0.05)		
6 month	3.96*	0.66+		0.20	1157
	(0.52)	(0.04)			
	10.10*		0.25*	0.07	1032
	(0.38)		(0.03)		
	4.52*	0.99	-0.34*	0.17	1027
	(0.61)	(0.09)	(0.03)		
1 year	8.53*	0.37+		0.09	1029
	(0.50)	(0.04)			
	17.60*		-0.21*	0.07	772
	(0.42)		(0.03)		
	13.56+	0.53+	-0.47*	0.15	769
	(0.60)	(0.06)	(0.04)		

* Significantly different from zero at the 5 percent level.
+ Significantly different from unity at 5 percent level.
++ Implied volatility at ask price.

Table 4. Euro: Dependent Variable is Future Realized Volatility

Maturity	a	Slopes On ISD++	Slopes On HV	Adj. R^2	Observations
1 month	3.74*	0.57+		0.39	537
	(0.42)	(0.03)			
	5.08*		0.51*	0.21	698
	(0.38)		(0.03)		
	3.96*	0.64+	-0.10*	0.34	529
	(0.44)	(0.05)	(0.06)		
2 month	2.93*	0.66+		0.52	534
	(0.33)	(0.03)			
	4.51*		0.60*	0.27	631
	(0.39)		(0.04)		
	3.92	0.80+	-0.25*	0.51	506
	(0.38)	(0.05)	(0.06)		
3 month	2.26*	0.73+		0.67	537
	(0.26)	(0.02)			
	3.19*		0.76*	0.49	569
	(0.33)		(0.30)		
	2.60*	0.56+	0.18*	0.64	486
	0.32	(0.04)	(0.05)		
6 month	3.01*	0.71+		0.74	536
	(0.21)	(0.02)			
	3.92*		0.75*	0.56	420
	(0.39)		(0.03)		
	4.40*	0.49+	0.14	0.66	420
	(0.30)	(0.05)	(0.60)		
1 year	4.75*	0.60		0.77	537
	(0.16)	(0.01)			
	4.65*		0.75*	0.56	537
	(0.41)		(0.04)		
	4.66*	0.21+	0.49*	0.58	290
	(0.41)	(0.06)	(0.08)		

* Significantly different from zero at the 5 percent level.
+ Significantly different from unity at 5 percent level.
++ Implied volatility at ask price.

While this premium is termed volatility risk premium here, it could reflect at least two other factors: (1) the competitiveness of the option market; (2) a compensation for dynamic hedging services. If the option market is not competitive, market makers would be able to charge a premium above the price derived from replicating the option, which could be mistaken for the volatility risk premium. Given the relative ease with which any financial institution, which has already incurred large fixed costs to set up trading operations, can enter the market for any particular currency pair, we will assume here that the option markets are competitive. Option buyers may also have to pay an additional fee above the arbitrage-based price to the option writer as compensation for performing dynamic hedging services on behalf of the option buyer who would have to conduct these operations himself in the absence of an option market.

III.1. The Econometric Model

We define the volatility risk premium akin to the forward premium, which drives a wedge between the forward rate and the anticipated rate of domestic currency depreciation. Thus, the volatility risk premium is the difference between the observed implied volatility and the anticipated future realized volatility:

$$\sigma_t^{VRP} = \sigma_t^{IV} - E[\sigma_{t,T}^{RV}]\Big|_t, \tag{5}$$

where σ_t^{VRP} is the volatility risk premium the option writer needs to be paid in order to take on volatility risk at time *t*, σ_t^{IV} is the implied volatility at time *t*, and $E[\sigma_{t,T}^{RV}]\Big|_t$ is the anticipated (at time *t*) future realized volatility between time *t* and the maturity horizon of the option, *T*. Since we cannot observe the anticipated future realized volatility, we need to assume rational expectations, i.e. that market participants anticipate on average at time *t* the volatility that is actually realized ex-post at time *T*:

$$E[\sigma_{t,T}^{RV}]\Big|_t = \sigma_{t,T}^{RV}\Big|_T. \tag{6}$$

Omitting the time index denoting the observation time, we can extract the volatility risk premium ex-post:

$$\sigma_t^{VRP} = \sigma_t^{IV} - \sigma_{t,T}^{RV} \tag{7}$$

This specification implies that the predictability regressions we estimated in section II were biased. Accounting for the volatility risk premium, equation (1) needs to be reformulated:

$$\sigma_{t,T}^{RV} = \alpha + \beta_1 \sigma_t^{IV} + \beta_2 \sigma_t^{VRP} + \varepsilon_{t,T,}. \tag{8}$$

Based on equation (7) we will be interested in testing whether $\beta_1 = 1$ and $\beta_2 = -1$.

We can now examine how the results under model (1) would be affected by the omitted variable, the volatility risk premium. Since the volatility risk premium, σ_t^{VRP}, is likely to be positively correlated with implied volatility, σ_t^{IV}, there is likely a bias. Substituting for the volatility risk premium from (7) in (8) and re-arranging the equation the following result obtains:

$$\sigma_{t,T}^{RV} = \frac{\alpha}{1+\beta_2} + \left(\frac{\beta_1+\beta_2}{1+\beta_2}\right)\sigma_t^{IV} + \frac{\varepsilon}{1+\beta_2} \tag{9}$$

Since $\beta_1 > 0$, $\beta_2 < 0$ and $b < 1$, the coefficients, *a* and *b*, in equation (1) were likely biased downward:

$$\alpha = a(1+\beta_2) \text{ and } \beta_1 = b + (b-1)\beta_2,$$

so that $\alpha > a$ and $\beta_1 > b$. It follows that the regression coefficients in Tables 1 to 4 are smaller than they would be if the volatility risk premium had been accounted for.

III.2. Finding an Ex-ante Proxy for the Risk Premium

While the actual volatility risk premium under the rational expectations assumption is observable ex-post, we need a measure that is available ex-ante in order to re-test the predictability of realized volatility with forward looking market based indicators. There are at least two possible proxies that would reflect the risk premium: (1) the implied volatility bid-ask spread, and (2) the volatility of implied volatility

Like implied volatility itself, its bid-ask spread is of a forward looking nature, reflecting expectations of greater risk in the future. Wei (1994) establishes that an increase in anticipated foreign exchange volatility (of the pound, mark, yen and Swiss franc)—as proxied by implied volatility—raises the spot bid-ask spread. Unanticipated realized volatility does not have this effect, confirming that risk anticipation influences the bid-ask spread. Using tick-by-tick inter-bank quotes for the dollar/mark exchange rate, Huang and Masulis (1996) show that the spot bid-ask spread increases as there is greater disagreement among dealers and/or rapidly changing dealer expectations of the spot exchange rate. Becker, Chadha and Sy (1998) emphasize that the main determinant of emerging market currency bid-ask spreads is the exchange rate risk facing the trader. Galati (2000) confirms that due to the dealer's inventory cost of holding foreign exchange the spot bid-ask spread rises when exchange rate volatility increases. Pasquariello (2000) points out that the bid-ask spread varies intra-day with the number of present dealers, which suggests to use daily data.[4]

[4] There could of course be structural changes in the competitiveness, which would also be present in daily bid-ask spread data.

Based on the literature it seems reasonable to postulate that—similar to the spot bid-ask spread—the implied volatility bid-ask spread reflects, among other things, option inventory costs and is likely to widen when exchange rate volatility is anticipated to increase.[5] Thus, the implied volatility bid-ask spread can be interpreted as a proxy for the volatility risk premium. While changes in the level of implied volatility reflect changes in risk itself, changes in the bid-ask spread can be interpreted to reflect changes in risk preferences: In other words, the option premium increases when there is more market risk and the cost of insuring this risk is higher; the bid-ask spread around this option premium increases when option writers are more risk-averse and demand larger compensation for providing this insurance.

Another possible candidate variable as proxy for the volatility risk premium is the volatility of implied volatility, in market parlance known as "vol of vol".[6] It is computed analogously to historical volatility. Traders pay close attention to the vol of vol as it generates much of the risk of holding delta-hedged trading portfolios containing options. Malz (2000) uses it as a measure of how fast implied volatility is rising. A potential drawback of this instrument is that the vol of vol is backward looking. By contrast, one expects the trader to impound all available information about the future in the bid-ask spread.

The two suggested instruments are used to test whether their ex-ante observable value can predict the contemporaneous volatility risk premium, which, however, is only observable ex-post:

$$\sigma_t^{VRP} = \varpi + \theta\, \sigma_t^{VOV} + \vartheta\, \sigma_t^{IVBA} + \xi_t\,, \tag{10}$$

where σ_t^{VOV} is the volatility of implied volatility at time *t*, and σ_t^{IVBA} is the bid-ask spread of implied volatility at time *t*.

Tables 5 to 8 report the results of the instrument regressions in which the ex-post observable volatility risk premium is explained with the ex-ante available potential instruments. Based on the adjusted R^2, the two instruments together have very high predictive power for the two emerging market currencies, the Thai Baht and the Hong Kong Dollar, with R^2 ranging from 49 percent to 83 percent.[7] For the mature market currencies, the Japanese yen and the euro, it tends to be more difficult to predict the current (unobservable) volatility risk premium with currently available information. The adjusted R^2 ranges between 2 and 56 percent. Despite a low adjusted R^2, coefficients on both predictors are significantly different from zero with the exception of the implied volatility bid-ask spread of the euro at shorter time horizons. This

[5] While we expect the bid-ask spread to be a good predictor of the volatility risk premium, we have to be aware that it may also reflect liquidity conditions in the market.

[6] The measure is similar to "volatility cones", a technique used by traders to assess whether implied volatility is expensive or cheap relative to recent experience.

[7] A notable exception are the estimates for the Hong Kong Dollar at the one year horizon.

suggests that even in cases where predictability of the volatility risk premium appears low as measured by R^2, using the instruments will improve over omitting the variables as is the case for model (1).

Surprisingly, the predictability of the volatility risk premium increases with an increasing time horizon for the baht, the yen and the euro (with the exception of the one month horizon). This is not the case for the Hong Kong Dollar, where predictability of the volatility risk premium declines with the time horizon.

Dropping vol of vol in equation (10) and using only the forward looking bid-ask spread as predictor of contemporaneous volatility risk premium worsens the explanatory power of the instrument regression in all cases. The amount by which R^2 drops is larger for longer time horizons. We can conclude that in their formulation of risk preferences market participants pay a considerable amount of attention to comparing the current level of implied volatility to its recent history. This result is intuitive.

III.3. Re-Estimating the Predictability Equations

Using the instrument for the ex-ante unobservable volatility risk premium, it is possible to re-estimate the original predictability equations, addressing the omitted variable problem encountered in section II. Equation (1) is re-estimated, taking into account the volatility risk premium, by obtaining the fitted value of equation (10), $\hat{\sigma}_t^{VRP}$, and substituting it in equation (8):

$$\sigma_{t,T}^{RV} = \alpha + \beta_1 \sigma_t^{ISD} + \beta_2 \hat{\sigma}_t^{VRP} + \varepsilon_{t,T,} \qquad (11)$$

Since σ_t^{ISD} and σ_t^{VRP} are correlated variables, omitting one of them biases the coefficient of the remaining variable. Including σ_t^{VRP} in the regression ensures that β_1 is a consistent estimator and eliminates the bias present when omitting the volatility risk premium. The risk premium may be measured with error (due to the use of a proxy) and may thus lead to lower explanatory power and larger standard errors than would be the case if it could be measured with precision ex-ante. However, this is still preferable to omitting the variable altogether.

Table 5. Thai Baht: Dependent Variable is Future Realized Volatility Risk Premium

Maturity	a	Slopes On: Vol of Vol	Slopes On: B-A Spread	Adj. R^2	Observations
1 month	0.94*	-0.04*	1.13*	0.59	719
	(0.32)	(0.01)	(0.06)		
	0.44*		1.08*	0.56	741
	(0.31)		(0.04)		
2 month	-0.59*	-0.02*	1.10*	0.56	698
	(0.34)	(0.01)	(0.07)		
	-0.53*		1.17*	0.53	741
	(0.34)		(0.04)		
3 month	-0.60*	0.00	1.22*	0.49	676
	(0.40)	(0.01)	(0.08)		
	-0.76*		1.19*	0.46	741
	(0.39)		(0.50)		
6 month	0.99*	0.01	1.28*	0.82	544
	(0.26)	(0.01)	(0.04)		
	-1.68*		1.26*	0.36	675
	(0.55)		(0.07)		
1 year	4.66*	0.05	1.18*	0.83	286
	(0.70)	(0.01)	(0.06)		
	-2.38*		1.55*	0.42	546
	(0.72)		(0.08)		

* Significantly different from zero at the 5 percent level.

Table 6. Hong Kong Dollar: Dependent Variable is Future Realized Volatility Risk Premi

Maturity	a	Slopes On: Vol of Vol	Slopes On: B-A Spread	Adj. R^2	Observations
1 month	0.36*	0.01*	1.14*	0.83	719
	(0.05)	(0.00)	(0.02)		
	0.42*		1.17*	0.82	741
	(0.05)		(0.02)		
2 month	-0.15*	0.02	1.51	0.8	698
	(0.07)	(0.004)	(0.05)		
	-0.15*		1.77*	0.77	741
	(0.07)		(0.03)		
3 month	-0.2	0.04*	1.72*	0.75	676
	(0.10)	(0.00)	(0.06)		
	0.22*		1.99*	0.69	741
	(0.09)		(0.05)		
6 month	-0.67*	0.08*	1.85*	0.69	578
	(0.22)	(0.01)	(0.08)		
	1.87*		2.12*	0.57	708
	(0.14)		(0.07)		
1 year	23.75*	-0.25	0.08*	0.11	318
	(0.48)	(0.04)	(0.04)		
	8.52*		0.18*	0.02	578
	(0.19)		(0.04)		

* Significantly different from zero at the 5 percent level.

Table 7. Japanese Yen: Dependent Variable is Future Realized Volatility Risk Premium

		Slopes On			
Maturity	a	R Vol of Vol	B-A Spread	Adj. R^2	Observations
1 month	-1.17*	0.04*	2.95*	0.03	1144
	(0.32)	(0.01)	(0.86)		
	-0.76		3.97	0.02	1225
	(0.28)		(0.73)		
2 month	-0.87*	0.04*	0.96*	0.02	1085
	(0.25)	-0.01	(0.33)		
	0.08		1.28	0.02	1202
	(0.14)		(0.33)		
3 month	-1.53	0.01	6.16*	0.07	1029
	(0.31)	(0.01)	(0.75)		
	-0.92*		4.70*	0.04	1219
	(0.21)		(0.63)		
6 month	-3.85*	0.11*	6.85*	0.17	1027
	(0.31)	(0.01)	(0.62)		
	-1.25		6.27*	0.08	1157
	(0.20)		(0.62)		
1 year	-8.90	0.28*	5.67*	0.56	770
	(0.28)	(0.01)	(0.56)		
	-2.20*		7.95*	0.13	1029
	(0.20)		(0.62)		

* Significantly different from zero at the 5 percent level.

Table 8. Euro: Dependent Variable is Future Realized Volatility Risk Premium

		Slopes On			
Maturity	a	R Vol of Vol	B-A Spread	Adj. R^2	Observations
1 month	-0.28	0.16*	-0.68	0.25	512
	(0.25)	(0.01)	(0.49)		
	1.43*		-0.25	0.00	537
	(0.24)		(0.53)		
2 month	-0.03	0.09*	0.06	0.05	491
	(0.23)	(0.01)	(0.43)		
	1.05*		0.003	0.00	534
	(0.17)		(0.41)		
3 month	-0.39*	0.10*	0.24	0.11	470
	(0.19)	(0.01)	(0.34)		
	0.67*		0.46	0.00	537
	(0.13)		(0.31)		
6 month	-1.57*	0.16*	0.48*	0.21	404
	(0.21)	(0.02)	(0.28)		
	-0.15		1.17*	0.04	536
	(0.09)		(0.25)		
1 year	-5.30*	0.31*	0.88*	0.44	275
	(0.39)	(0.02)	(0.22)		
	-0.63*		1.48*	0.07	537
	-0.09		(0.23)		

* Significantly different from zero at the 5 percent level.

The results of this estimation are reported in Tables 9 to12. As expected, β_1 in equation (11), using one or both instruments, increases significantly and is in some cases not distinguishable from unity at the 95 percent confidence level. The model performs particularly well for the Japanese yen, where at least for near and medium term horizons not only the ISD coefficients, but also the coefficients for the volatility risk premium variable, β_2, behave as expected. Improvements are also noticeable for the Thai baht and the euro, where re-estimated coefficients may not equal unity, but are significantly larger and closer to unity than the original estimates.[8] In these cases, the constant term is—while still significantly different from zero—smaller. Since the explanatory power of the predictability depends on the amount of additional information introduced through the instrument variable, the improvement in adjusted R^2 of the re-estimated predictability regressions is directly proportional to the adjusted R^2 in the instrument regressions (10). The re-estimation procedure did not alter the coefficients for the Hong Kong Dollar. The success of the peg meant that neither implied volatility nor the implied volatility risk premium are good predictors of exchange rate volatility.

To summarize the results so far, the existence of a volatility risk premium has biased the results of studies attempting to test the predictability of future realized volatility using implied volatility. While this risk premium can be extracted ex-post under the rational expectations assumption, it needs to be approximated with ex-ante available instruments. The forward looking bid-ask spread together with the vol of vol, which compares the current level of implied volatility with its recent history, provide a reasonable instrument to approximate the contemporaneous volatility risk premium. Accounting for the premium reduces this bias significantly, and in some cases removes it entirely. This result has important consequences for another strand of the literature, which focuses on extracting risk neutral density functions from option prices.

IV. Extracting Exchange Rate Expectations from Currency Option Data

Options provide the opportunity to infer a risk neutral probability distribution around the forward rate. The forward rate as the market's risk neutral expectation of the anticipated exchange rate can be thought of as the mathematical expectation of a whole series of possible future spot exchange rate outcomes, each associated with a different probability. Similar to the forward rate, the market price of a call option can be thought of as the risk neutral mathematical expectation of its payoff at maturity, discounted to the present.[9] Breeden and Litzenberger (1978) show that if the prices of put and call options with many different strike prices can be simultaneously observed, one can trace out the entire risk-neutral distribution, since the discounted risk-neutral density function of the asset price equals the second derivative of the call option price with respect to the exercise price.

[8] The medium term estimates for the euro are an exception to this observation.

[9] See Appendix 1 for an introduction to risk-neutrality.

Table 9. Thai Baht: Dependent Variable is Future Realized Volatility

Maturity	a	Slopes On ISD++	HV	IV B-A Hat	VV BA Hat	Adj. R^2	Observations
1 month	7.15*	0.53+	-0.03+			0.45	305
	(1.35)	(0.3)	(0.03)				
	4.27*	0.86		-0.78		0.53	305
	(1.26)	(0.06)		(0.12)			
	4.31*	0.85+			-0.83	0.56	283
	(1.24)	(0.06)			(0.13)		
2 month	14.61*	0.40+	-0.09+			0.27	305
	(1.63)	(0.04)	(0.04)				
	9.03	0.84+		-1.24		0.37	305
	(1.38)	(.07)		(.17)			
	6.93*	1.00			-1.29	0.42	262
	(1.55)	(0.08)			(o.16)		
3 month	23.18*	0.26+	-0.19+			0.13	305
	(2.07)	(0.04)	(0.05)				
	13.83*	0.72+		-1.24		0.22	305
	(1.53)	(0.08)		(0.18)			
	12.07*	0.89			-1.27	0.24	240
	(1.98)	(0.10)			(0.24)		
6 month	-10.99*	0.02	1.05+			0.18	305
	(4.98)	(0.04)	(0.13)				
	26.0*	0.09		-0.36+		0.01	305
	(1.67)	(0.10)		(0.22)			
	2.33*	0.60			-0.33+	0.65	174
	(0.95)	(0.05)			(0.09)		
1 year	60.15*	0.14+	-1.45+			0.60	305
	(1.90)	(0.03)	(0.07)				
	27.00*	-0.39+		0.57+		0.08	305
	(1.36)	(0.08)		(0.18)			
	6.87*	0.06			-0.02+	0.50	145
	(0.31)	(0.02)			(0.02)		

* Significantly different from zero at the 5 percent level.
\+ Significantly different from unity at 5 percent level.
++ Implied volatility at ask price.

Table 10. Hong Kong Dollar: Dependent Variable is Future Realized Volatility
Model Accounts for Volatility Risk Premium

		Slopes On					
Maturity	a	ISD++	HV	IV B-A Hat	VV BA Hat	R-Squared	Observations
1 month	0.08*	0.08+	-0.11+			0.38	604
	(0.02)	(0.01)	(0.03)				
	0.08*	0.08+		-0.12+		0.37	604
	(0.02)	(0.01)		(0.01)			
	0.08*	0.08+			-0.02+	0.37	582
	(0.02)	(0.01)			(0.01)		
2 month	0.09*	0.05+	0.05+			0.41	604
	(0.01)	(0.01)	(0.03)				
	0.07*	0.04+		0.05+		0.42	604
	(0.01)	(0.01)		(0.01)			
	0.06*	0.08+			0.04+	0.44	561
	(0.01)	(0.006)			(0.007)		
3 month	0.07*	0.05+	0.02+			0.56	604
	(0.01)	(0.01)	(0.03)				
	0.05*	0.04+		0.06+		0.56	604
	(0.01)	(0.00)		(0.01)			
	0.03*	0.03+			0.04+	0.61	539
	(0.01)	(0.003)			(0.004)		
6 month	0.10*	0.03+	0.13			0.33	604
	(0.02)	(0.01)	0.04				
	0.11*	0.01+		0.08+		0.42	604
	(0.11)	(0.00)		(0.01)			
	0.06*	0.02+			0.01+	0.42	474
	(0.01)	(0.002)			(0.004)		
1 year	0.51*	0.02+	-0.93+			0.50	578
	(0.02)	(0.01)	(0.05)				
	0.14*	0.02+		0.00+		0.21	578
	(0.01)	(0.00)		(0.01)			
	0.10*	0.01+			-0.008+	0.64	318
	(0.01)	(0.0006)			(0.002)		

* Significantly different from zero at the 5 percent level.
+ Significantly different from unity at 5 percent level.
++ Implied volatility at ask price.

Table 11. Japanese Yen: Dependent Variable is Future Realized Volatility

		Slopes On					
Maturity	a	ISD++	HV	IV B-A	VV BA Hat	Adj. R^2	Observations
1 month	0.35	0.90	0.02+			0.40	1205
	(0.43)	(0.05)	(0.04)				
	0.18	0.98		-0.97		0.41	1225
	(0.42)	(0.03)		(0.28)			
	-1.53*	1.13+			-1.42	0.44	1166
	(0.48)	(0.04)			(0.22)		
2 month	1.26*	0.68+	0.19+			0.39	1159
	(0.45)	(0.05)	(0.04)				
	0.56	0.95		0.92		0.39	1201
	(0.44)	(0.03)		(0.26)			
	-0.69	1.05			-1.18	0.42	1084
	(0.49)	(0.04)			(0.24)		
3 month	2.68*	0.42+	0.36+			0.38	1154
	(0.48)	(0.06)	(0.05)				
	0.37	0.97		0.96		0.38	1218
	(0.45)	(0.04)		(0.14)			
	0.02	0.99			-1.00	0.40	1027
	(.049)	(0.04)			(0.12)		
6 month	4.52*	0.99	-0.34+			0.17	1027
	(0.61)	(0.09)	(0.03)				
	3.01*	0.76+		0.81+		0.24	1157
	(0.52)	(0.04)		(0.10)			
	4.01*	0.70+			-0.80+	0.24	1027
	(0.55)	(0.04)			(0.07)		
1 year	13.56+	0.53+	-0.47+			0.15	769
	(0.60)	(0.06)	(0.04)				
	6.98*	0.48+		-0.64+		0.15	1029
	(0.52)	(0.04)		(0.08)			
	6.98*	0.51+			-0.79+	0.41	770
	(0.53)	(0.04)			(0.03)		

* Significantly different from zero at the 5 percent level.
+ Significantly different from unity at 5 percent level.
++ Implied volatility at ask price.

Table 12. Euro: Dependent Variable is Future Realized Volatility

Maturity	a	Slopes On ISD++	HV	BA Hat	VV BA Hat	Adj. R^2	Observations
1 month	3.96*	0.64+	-0.10*			0.34	529
	(0.44)	(0.05)	(0.06)				
	7.72*	0.57+		-2.98		0.34	537
	-2.61	(0.03)		(1.92)			
	2.82*	0.72+			-0.61+	0.36	512
	(0.46)	(0.04)			(0.09)		
2 month	3.92	0.80+	-0.25*			0.51	506
	(0.38)	(0.05)	(0.06)				
	-12.93*	0.66+		15.09			
	(128.68)	(0.03)		(122.05)			
	3.54*	0.60+			0.19+	0.47	491
	(0.39)	(0.04)			(0.18)		
3 month	2.60*	0.56+	0.18*			0.64	486
	0.32	(0.04)	(0.05)				
	3.16*	0.73+		-1.05		0.67	537
	(0.57)	(0.20)		(0.59)			
	4.67*	0.48+			0.95+	0.63	470
	(0.34)	(0.04)			(0.17)		
6 month	4.40*	0.49+	0.14			0.66	420
	(0.30)	(0.05)	(0.60)				
	3.09*	0.71+		-0.47+		0.74	536
	(0.21)	(0.02)		(0.18)			
	5.92*	0.46+			0.47+	0.62	404
	(0.40)	(0.03)			(0.12)		
1 year	4.66*	0.21+	0.49*			0.58	290
	(0.41)	(0.06)	(0.08)				
	4.55*	0.61+		-0.43+		0.78	537
	(0.17)	(0.01)		(0.10)			
	4.08*	0.67+			-0.65	0.61	275
	(0.42)	(0.03)			(0.07)		

* Significantly different from zero at the 5 percent level.
\+ Significantly different from unity at 5 percent level.
++ Implied volatility at ask price.

IV.1. Risk-Neutral Probability Distributions

The market price of a European call option, *c(t, X, T)*, is the difference between the expected value of the future exchange rate and the exercise price, with the probability weights drawn from the risk neutral distribution, $\pi(x)$:

$$c(t,X,T) = e^{-r\tau} E^{*}\left[\max(S_T - X)\right] = e^{-r\tau} \int_X^{\infty} (S_T - X)\pi(S_T) dS_T\,, \tag{12}$$

where X is the exercise price, t and T are the current and option maturity dates, $\tau \equiv T - t$, r is the risk-free interest rate, S_t is the asset price at time t, E^{*} is the expectation operator taken under the risk neutral probability distribution, and $\int_a^b \pi(S_T) dS_T \equiv P^{*}(a \leq S_T \leq b)$, where P^{*} denotes a risk neutral probability. Notice that all variables in equation (12) are observable, except for the risk neutral distribution, which is to be identified. Akin to the way the risk-neutral probabilities change as market conditions fluctuate (Appendix 1, case 2), $\pi(X)$ is the set of probabilities that changes as other observable variables change, so as to equate both sides of the equation.

In order to uncover the risk neutral probability distribution, we twice differentiate the price of the option with respect to the exercise price:

$$\frac{dc(t,X,T)}{dX} = -e^{-r\tau}[1 - \Pi(x)] \tag{13}$$

where $\Pi(x) \equiv P^{*}(S_T \leq x)$ is the risk neutral cumulative distribution function; and

$$\frac{d^2 c(t,X,T)}{dX^2} = e^{-r\tau}\pi(X)\,. \tag{14}$$

Theoretically, one could trace out the entire probability distribution using options with a series of very closely spaced exercise prices. In practice, only a few strike prices are observable, typically at least one for at-the-money, out-of-the money and in-the-money options. Given the scarcity of data, one possible solution is to assume that the risk neutral probability distribution of the future asset price belongs to a particular parametric family (see for example Melick and Thomas, 1996).[10] Other methods for estimating implied PDFs include: stochastic process methods, implied binomial trees, PDF approximating function methods, finite-difference methods, and implied volatility smoothing methods (for a survey of various methods see Chang and Melick (1999)).

[10] For a survey of methods to recover probability density functions from options data see Chang and Melick (1999).

Here we will use the smoothed volatility method, which is an alternative involving less restrictive assumptions. It uses over-the-counter options market data to interpolate between observable strike prices. A detailed procedure which effectively fits a polynomial function through the observed points to extract the probability distribution implied by a set of option prices is described in Malz (1997) and applied in this paper.[11]

The method takes advantage of the observed implied volatility smile. Contrary to the Garman-Kohlhagen assumptions, out-of-the money options often have higher implied volatilities than at-the-money options, indicating that the market perceives exchange rate returns to be leptokurtic; that is the risk-neutral likelihood of large exchange rate moves is greater than is consistent with the lognormal distribution assumed under Garman-Kohlhagen. Moreover, out-of-the money call options often have implied volatilities that differ from those of equally out-of-the money puts, indicating that the market perceives the distribution of future exchange rates to display skewness, or that market participants are willing to pay more for protection against sharp currency moves in one direction than in the other. The relationship between the implied volatility and the moneyness (as expressed by the delta) is called the "volatility smile" —due to the smile-like pattern that is often caused by higher implied volatilities for deep in-the-money and out-of-the money options.

The volatility smile is manifested in two readily available quotes of option combinations. Besides the at-the-money implied volatility, dealers frequently trade strangles and risk reversals, which involve two out-of-the-money options with the same delta (i.e. equal out-of-the moneyness) and the same maturity. The strangle is a position consisting of a long out-of-the money put and call. It is quoted as the spread volatility over the at-the-money forward volatility, and thus indicates the degree of curvature of the volatility smile. The risk reversal consists of a long call and a short put option. It is quoted as the spread of the foreign currency call over the foreign currency put option; and thus influences the skew of the volatility smile. A positive risk reversal quote implies that the foreign currency is more likely to appreciate under the risk neutral probability measure.

Malz's (1997) method approximates the volatility smile by expressing implied volatility as a function of the option's delta:

$$\hat{\sigma}_{\delta}(\delta) = atm_t - 2rr_t(\delta - 0.50) + 16str_t(\delta - 0.50)^2 , \qquad (15)$$

where atm_t is the level of the at-the-money implied volatility at time t, rr_t is the quoted risk reversal at time t, str_t is the quoted strangle at time t, and δ is the delta for which the strangle and risk reversals are being quoted (typically the 25 delta or 10 delta levels).[12] Once the

[11] Recent studies by Bliss and Panigirtzoglou (1999) and Jondeau and Rockinger (2000) show that the smoothed implied volatility method used here is more robust to small perturbations of the quoted prices than other methods.

[12] The delta of a put option, Δ_t^p, is related to the delta of a call option: $\Delta_t^p = \Delta_t^c - 1$.

market's schedule of implied volatilities is approximated by $\hat{\sigma}_\delta(\delta)$, the implied volatility corresponding to each exercise price, X, rather than each delta, needs to be found by substituting the mathematical expression for delta—itself a function of σ — into (15) and solving it for σ as a function of X. While there is no closed-form solution, the equation can be solved numerically.

Substituting the implied volatility function, $\sigma_X(X)$, into the Garman-Kohlhagen call option formula, yields a generalized Garman-Kohlhagen formula in which the implied volatility depends on the exercise price. The cumulative distribution and probability density function can then be derived following the method described in equations (13) and (14).
The Garman-Kohlhagen model is used simply as a transformation or mapping from one measurement space (implied volatilities) to another (option prices). The smoothed implied volatility smile method does not assume that the underlying exchange rate process is lognormal.

In order to implement this method, we fit the implied volatility function (15) through the three observed points of the volatility smile, which interpolates values between any two points. For values before the first known data point and beyond the last known data point a parabolic spline curve is fitted through the two closest data points. The resulting volatility smile is a continuous function, expressing implied volatility as a function of delta (Figure 1A). The volatility smile is transformed in order to express implied volatility as a function of the corresponding strike price, expressed in foreign currency per U.S. dollar (Figure 1B). This function is used to calculate a continuum of option prices, which can be twice differentiated. The resulting probabilities are plotted against the strike price and allow inferences about the relative probabilities assigned by the market on various exchange rate outcomes (Figure 1C).[13] If the exchange rate would indeed follow a geometric Brownian motion process as assumed by the Black-Scholes option pricing model, it would be log normally distributed with a standard deviation equal to the implied volatility of an at-the-money option. This distribution is displayed in Figure 1C to permit comparisons with risk-neutral market expectations.

IV.2. Risk Preferences

In an efficient market, current market information contains not only the market's view of the likelihood of future outcomes, but also captures risk preferences of market participants. Forecasts based on current market prices of options reflect aggregate views about the future level of the exchange rate and impound aggregate risk preferences in the form of risk neutral probabilities. For example, the payoff to an option writer doesn't only depend on an individual participant's view of the real probabilities, but also on how his view differs from collective views and risk preferences of all participants. If an individual's view differs from that of the market he can take on or lay off risk at the prevailing price.

[13] Absolute probabilities are not indicated on the vertical axis of the density graphs because they depend on the bin size chosen for the interpolation of the spline curve.

Figure 1. Euro: Volatility Smile, 3 Month Option, December 8, 2000

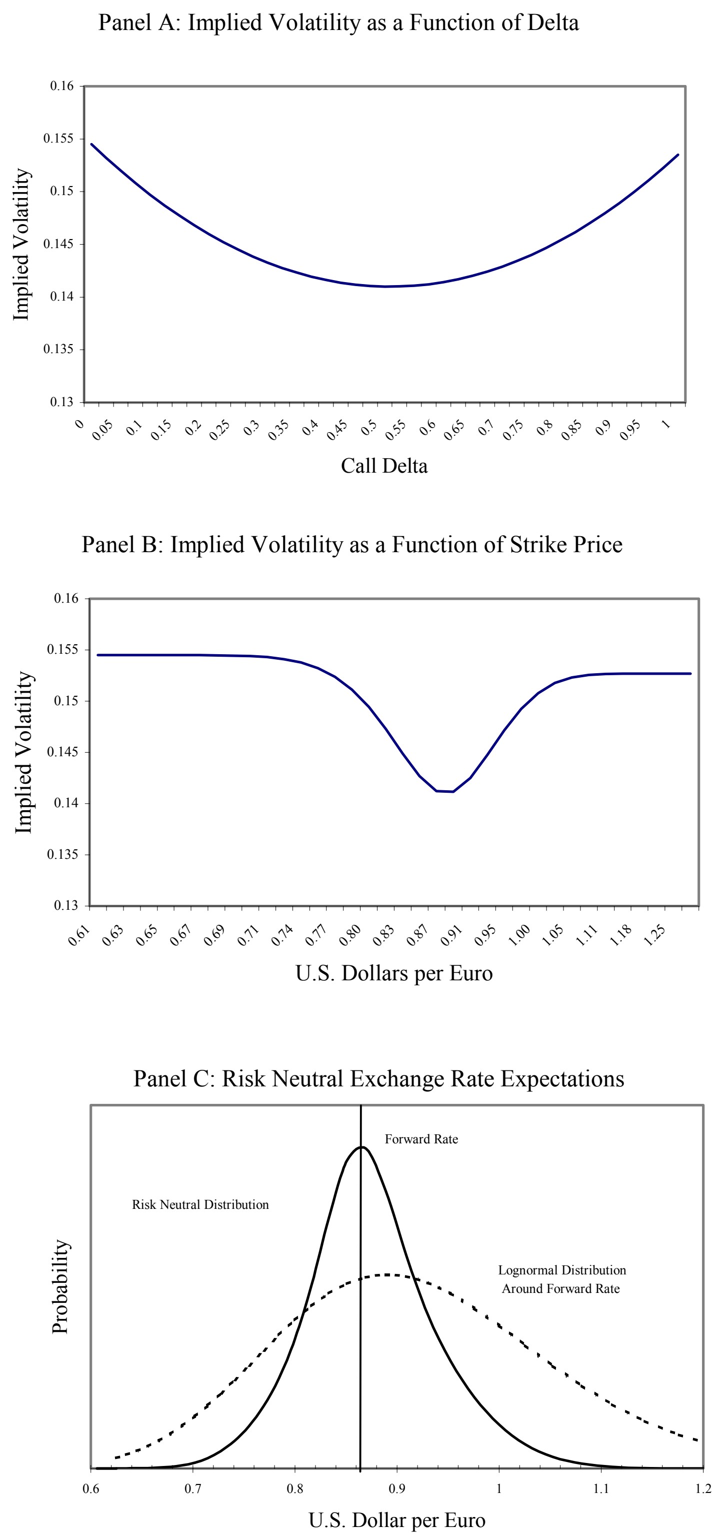

A risk neutral probability is determined by the market and does not reflect true expected probabilities (see Appendix 1). For example, the forward rate is often taken as the "market's" consensus forecast of the anticipated future exchange rate at the horizon of the maturity of the forward contract. This is only accurate in a risk neutral world. If agents eschew risk, investing in a foreign currency bond and eliminating exchange rate risk by hedging the currency risk exposure with a forward contract will have a different payoff than investing in a foreign currency bond and repatriating the proceeds at the prevailing spot market rate at the time of the bond's maturity. If agents are risk averse, they need to be compensated for taking on uncovered positions in the form of a risk premium. The forward premia literature established that forward rates do differ significantly from future realized spot rates because of the presence of a risk premium.[14] If there is a sudden increase in risk aversion, the risk premium will increase, making the gap between the forward rate and the anticipated rate of depreciation even larger. While the increase in the risk premium changes the forward rate, it may not necessarily affect the anticipated rate of depreciation.

Having shown that implied volatility by itself is a biased predictor of future realized volatility due to the presence of a time-varying risk premium, we need to adjust the exchange rate probability distributions extracted to allow for risk preferences. Risk-neutral PDFs in the literature are drawn for the mid-point between the implied volatility bid and offer prices. In order to reflect the risk premium the option writer charges for selling volatility insurance, each point of the PDF needs to be adjusted upward, as shown below. Since ex-ante the volatility risk premium is not observable, it is necessary to use a proxy. As an ex-ante available proxy $\hat{\sigma}_t^{VRP}$, the fitted value of equation (10) for a given day t is used to adjust the PDF.

The presence of a volatility risk premium affects the probability distribution. From (13) we know that:

$$\Pi(x) = \frac{dc(t,X,T)}{dX} e^{r\tau} + 1 \tag{16}$$

Increasing the implied volatility by the risk premium in order to adjust for risk preferences will affect the cumulative distribution:

$$\begin{aligned} \frac{d\Pi(X)}{d\sigma} &= e^{r\tau} \frac{dc(t,X,T)\big/dX}{d\sigma} \\ &= \sqrt{\tau} \quad \pi(x) > 0 \end{aligned} \tag{17}$$

An increase in the implied volatility increases the cumulative distribution function (CDF). Whether this shifts the mean of the density function depends on whether the cumulative distribution under the risk premium, $\Pi(X)^{VRP}$, dominates the cumulative distribution at the

---

[14] See Engel (1995) for a survey of the literature.

mid-point, $\Pi(X)^{MP}$, in the sense of first order stochastic dominance, i.e. $\Pi(X)^{VRP} > \Pi(X)^{MP}$ everywhere. If it does, the mean will shift. If $\Pi(X)^{VRP} \geq \Pi(X)^{MP}$, i.e. the adjusted CDF dominates the CDF at the mid-point in the sense of second order stochastic dominance, the mean may stay constant while the shape of the associated PDF changes.

IV.3. Two Illustrations

The probability distributions in Figure 1 were drawn using the mid points of the implied volatility quotes. They will adjusted them by the estimated volatility risk premium. Figure 2 shows implied probability distributions for the Thai baht/ U.S. Dollar exchange rate on February 26, 2001 at the one month horizon, drawn for risk-neutral and risk-adjusted distributions. The difference between the mean of the distribution and the spot rate of the day (42.935 baht/$) indicates the degree of baht depreciation expected over the next month. It can be seen that the risk-adjusted distribution implies a greater baht depreciation than the risk neutral distribution. The distribution around the mean reflects the degree of uncertainty associated with that point forecast and the relative probabilities of other outcomes. In this case, the adjusted and the risk-neutral distributions display nearly the same amount of uncertainty in expectations. The negative skew of both distributions indicates that on balance the market thinks further baht depreciation is more likely than appreciation. The risk-adjusted distribution is more negatively skewed, reflecting that a long risk-reversal position involves the option writer buying a THB call option and selling a THB put option, thus acquiring a risk profile that is equivalent to being long baht. Given the overall expectation of further baht depreciation on that day, the risk premium necessary to induce the dealer to accept this position is larger than it would be in a risk-neutral environment. The kurrtosis refers to the probability of extreme price movements occurring. Here, the risk-adjusted distribution factors in higher probabilities of extreme events than the risk-neutral distribution.

In the second illustration, time series of implied moments of the risk-neutral and risk-adjusted probability distributions at the 12 month horizon were extracted for the yen-dollar exchange rate from January 25, 1996 to January 19, 2001. Relative movements in the 12 month mean expected exchange rate, as well as the skewness of the risk-neutral distribution reveal that the yen was expected to depreciate for much of 1997 and 1998 (Figure 3). This trend reversed sharply in September 1998 in the wake of the unilateral Russian debt moratorium and the near-collapse of the American hedge fund, Long Term Capital Management. The ensuing crisis prompted many highly leveraged institutions to unwind the "carry trade"—the borrowing of funds at low interest rates in Japan and investing overseas at higher rates of return. As investors were struggling to meet margin calls, they needed to purchase yen in order to pay back their loans, resulting in a sharp appreciation of the yen. In contrast to the forward market, currency option market participants had been anticipating that the balance of risks was on a yen appreciation: the skew of the probability distribution reversed from negative values (implying further yen depreciation) prior to this event, even as the mean expected exchange rate was still depreciating (Figure 3).

Figure 2: Implied Probability Distribution, Thai Baht, February 26, 2001, 1 month horizon

	Mean	*S.D.*	*skew*	*kurtosis*
Risk-Adju	44.4026	2.6776	-0.7196	3.6953
Risk-Neut	44.3895	2.6737	-0.2851	3.4493

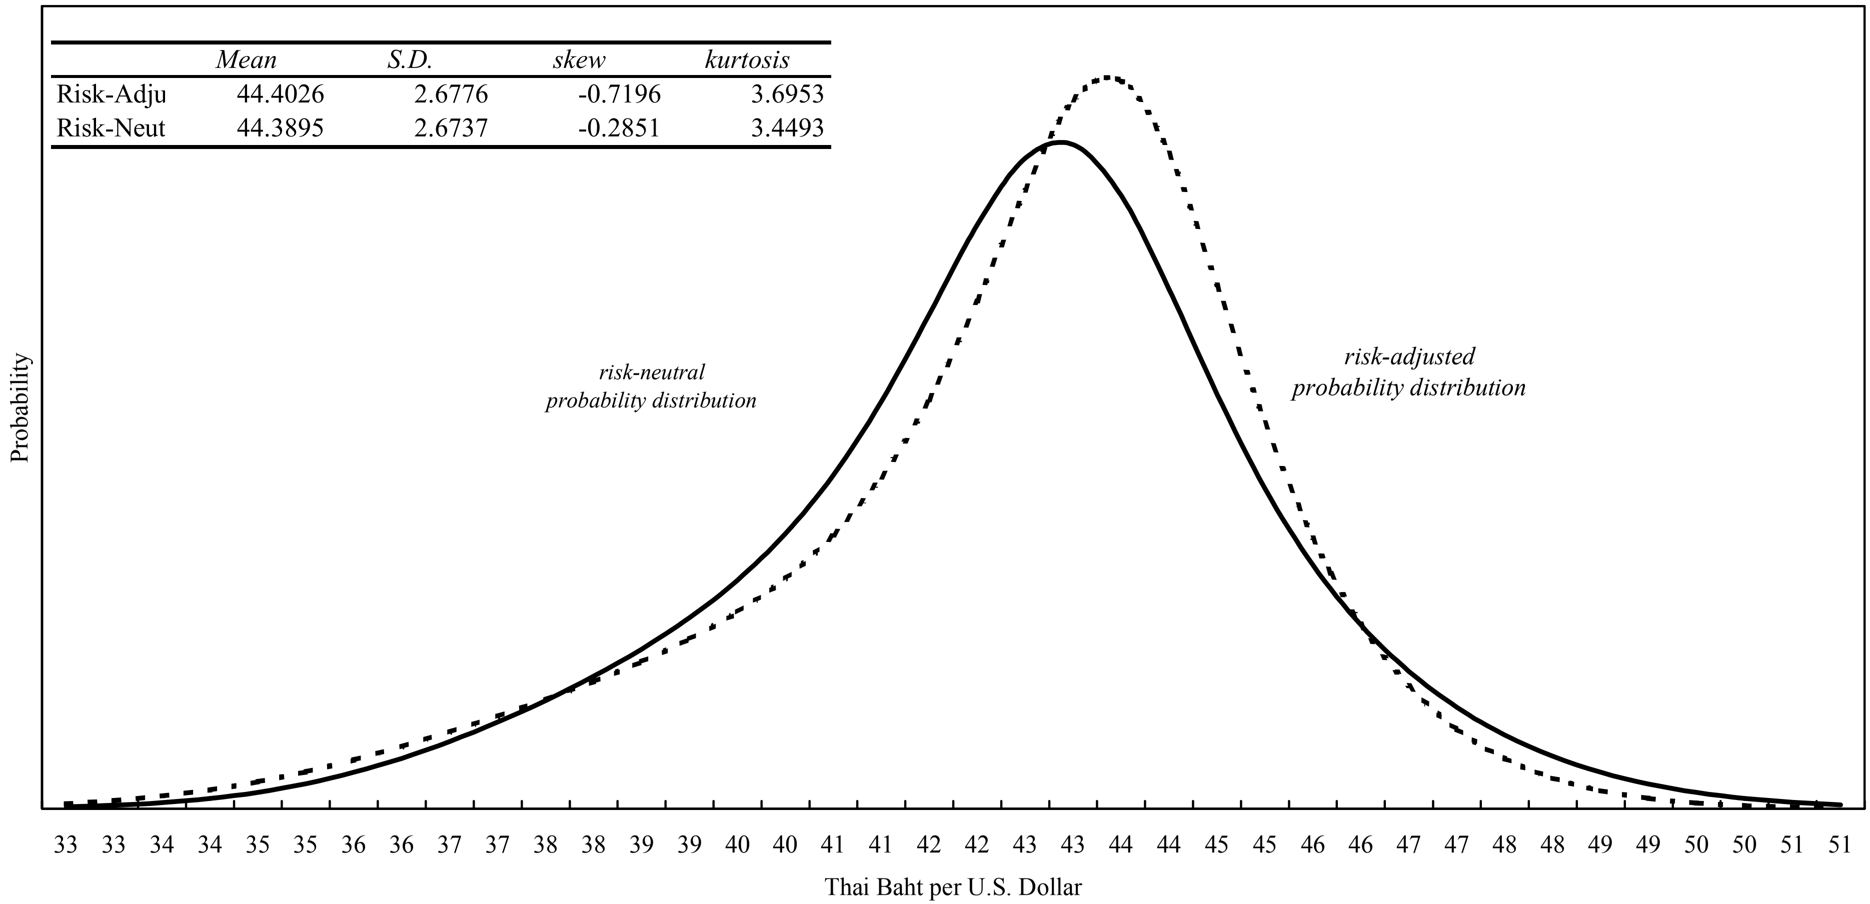

Had the risk-neutral probability distribution been adjusted for the volatility risk premium, the expected yen appreciation would have also been reflected in other moments as well. The difference in moments between the risk-neutral and the risk-adjusted distributions reflects the relative performance (Figure 4). The difference in mean expected exchange rates reveals that prior to the sharp yen appreciation in October1998 the risk-adjusted probability distribution predicted a lower rate of expected yen depreciation than the risk-neutral probability distribution (Figure 4). In addition, the (negative) skew of the risk-adjusted distribution is less inclined towards a further yen depreciation than that of the risk-neutral distribution (Figure 4).

In order to fully adjust the probability distribution, we need to know the bid-ask spread around each quote across the entire volatility smile. However, it appears that market makers and brokers do not keep bid-ask spread data for away from the money options. Only the bid ask spread around the at-the-money implied volatility is available in a historical time series. Hence, we need to make the assumption that the bid ask spread is constant across the volatility smile. In practice, it is likely that the bid-ask spread—and thus our indicator of the volatility risk premium—increases the further away from the money the quote is. This would make the results in this paper more pronounced.

V. Market Beliefs Around the Hong Kong SAR Equity Market Intervention

Applying the methodology developed in this paper to a real-world example allows differentiating the conclusions one may draw by merely using risk-neutral probability distribution functions of expectations from those conclusions one would draw taking into account not only risk, but also risk preferences.

In the period immediately before the August 1998 equity market intervention by the Hong Kong SAR authorities, the Hong Kong dollar came under several waves of speculative attacks. Amidst the Asian financial crisis, the depreciation of trading partners' and competitors' currencies had left the pegged Hong Kong dollar overvalued. As the asset price bubble deflated and market sentiment weakened owing to deteriorating economic conditions in the region, the Hong Kong dollar had already come under speculative pressure on several occasions in 1997. The depreciation of the yen against the U.S. dollar in mid-1998 led to renewed pressures against the Hong Kong dollar. At the same time, stock and futures prices plummeted, with the Hang Seng index sinking by almost 25 percent from mid-July, to 40 percent of its mid-1997 peak level. The authorities, arguing that the markets were being manipulated, and concerned that domestic confidence could be seriously weakened, reacted by intervening in the stock and futures markets on August 14 and again on August 28, acquiring 7 percent of market capitalization and nearly 30 percent of the traded stock volume.

Figure 3. Risk-Neutral Market Expectations, JPY-USD, 1996-2001, 12 month horizon

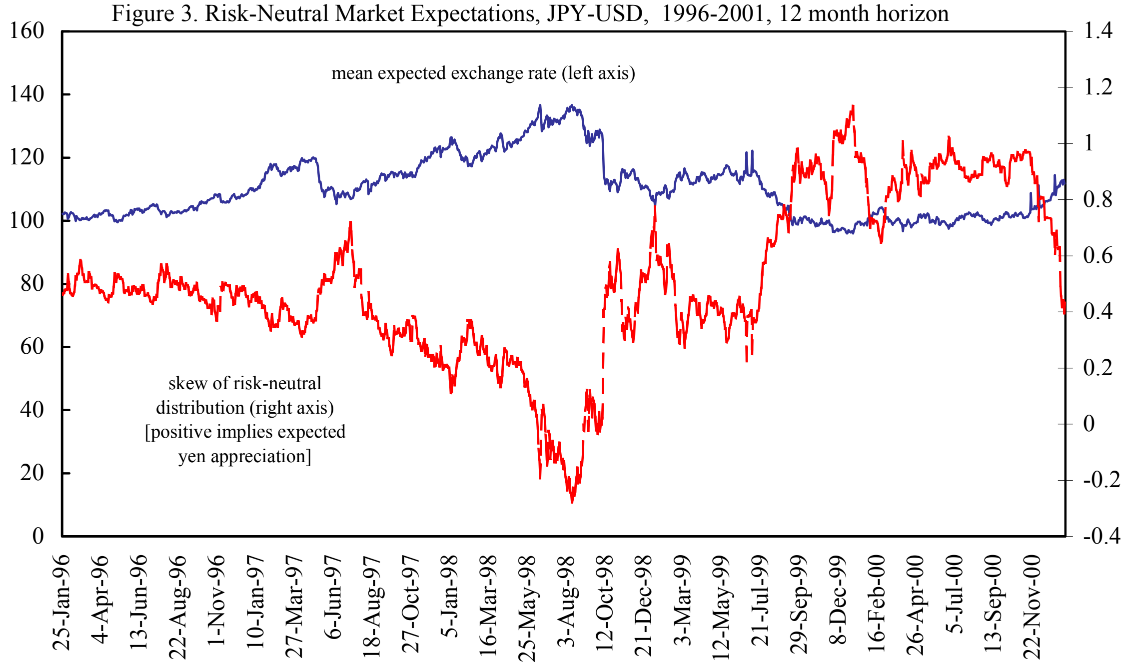

Figure 4. Difference between moments of risk-adjusted and risk-neutral distributions, JPY-USD, 1996-2001, 12 month horizon

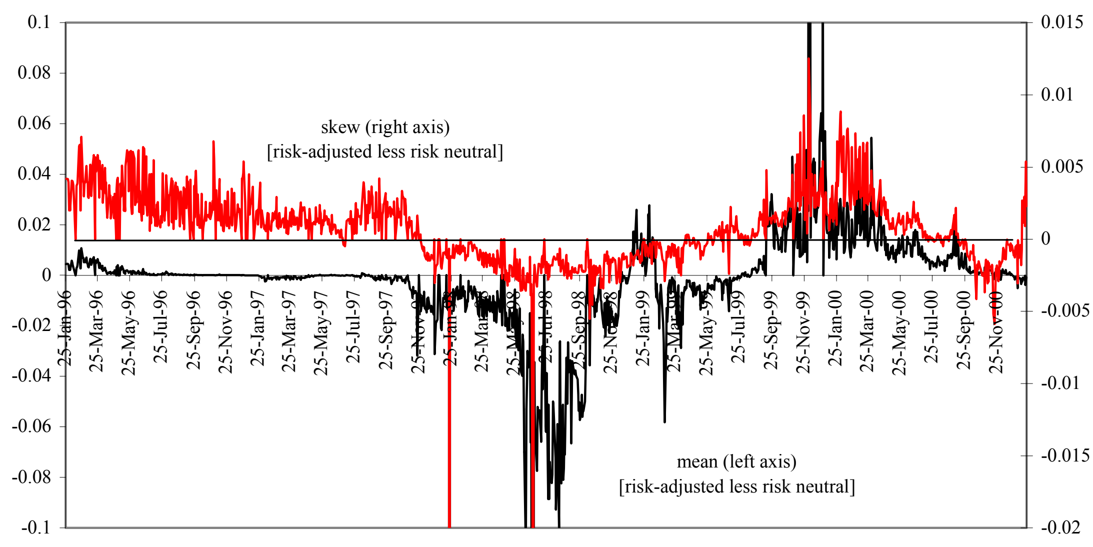

The intervention in the equity market was intended as the ultimate defense of the pegged exchange rate regime. According to the Hong Kong Monetary Authority, some large players were taking large short positions in the spot and futures markets for Hong Kong SAR equities and its currency. They then engaged in abrupt sales of Hong Kong dollars, driving up interest rates as the aggregate balance—the balance banks maintain with the Monetary Authority for clearing Hong Kong dollar transactions—contracted. This spike in interest rates drove down equity market prices, allegedly allowing speculators to close out their short positions at a profit. Had they also succeeded in breaking the peg, their short positions on the Hong Kong dollar would have generated an additional profit. While this account has been disputed[15], it can be argued that intervening in the equity market was an alternative form of defending the peg, alleviating speculative pressures on the Hong Kong dollar. Instead of investing the domestic currency proceeds from selling reserves in the money market, they were invested in domestic equities.

While the exchange rate peg remained intact, shifts in the extracted risk-neutral probability distribution reflected market sentiment changing with the intervention. By August 7 there was almost a consensus in the market that the Hong Kong dollar peg would break within 3 months (Figure 5). The mean of the risk-neutral exchange rate probability distribution was HK$8.08=US$1 (Table 1), against the linked rate of HK$7.8=US$1. The first intervention of August 14—when reportedly US$6 billion worth of equity was purchased by the authorities—calmed fears of a devaluation temporarily. By August 21 (the first observation after the August 14 intervention), the mean of the risk-neutral exchange rate distribution had moved to HK$7.99=US$1. However by August 28, although widely dispersed, market beliefs again expected a collapse of the peg. The August 28 intervention of US$9 billion, however, significantly allayed devaluation concerns. By September 4, market sentiment gravitated almost uniformly back to the linked rate.

How closely do these extracted probability functions reflect market beliefs? The procedure used in this exercise extract the risk-neutral probability function and not the true statistical distribution of expectations. Risk-neutral distributions do not expunge other characteristics of the market, such as the risk aversion of market participants and the liquidity of the market from the calculated probabilities. Consequently, it is difficult to unambiguously conclude that a shift in the implied probability distribution reflects a change in beliefs about the future value of the exchange rate or is due to changes in the underlying market structure, such as market liquidity or risk aversion.

[15] See International Monetary Fund, 1999.

Figure 5: Implied Risk-Neutral Probability Distributions, 3 Month Horizon, Hong Kong SAR, 1998

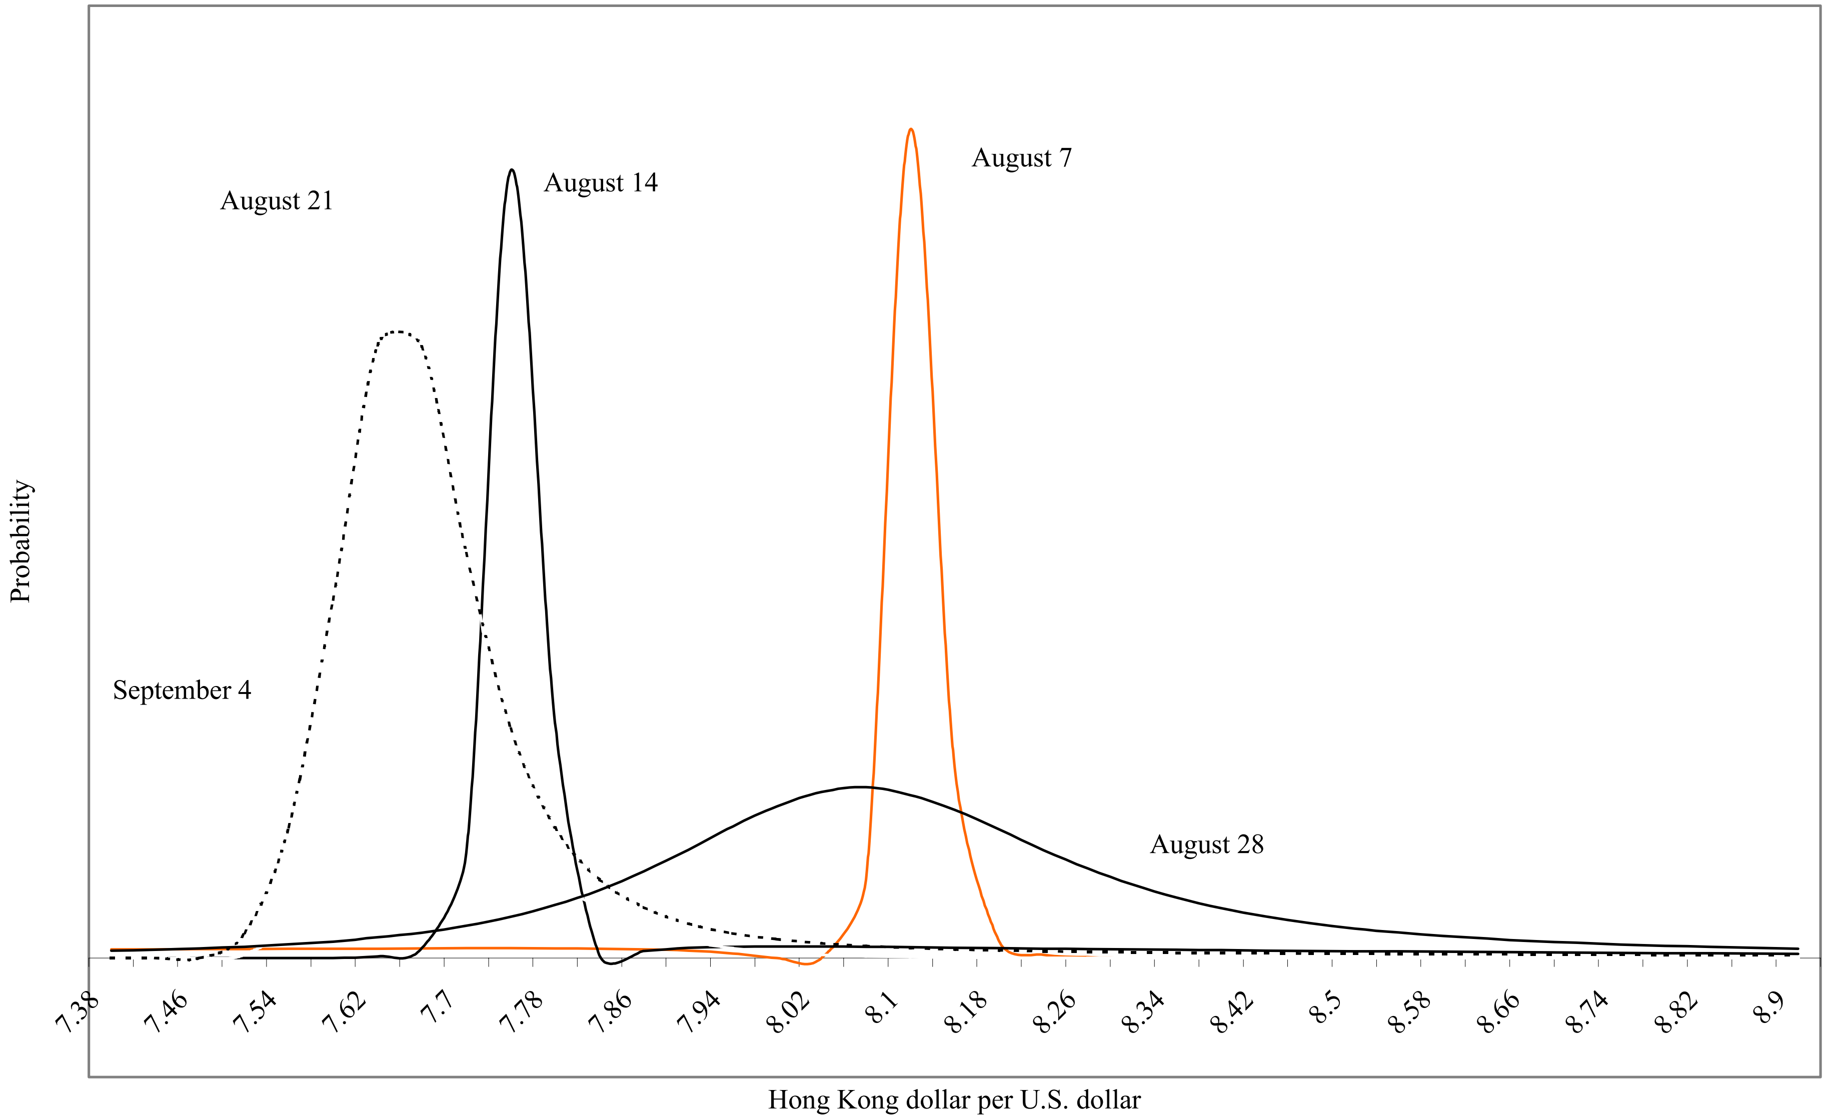

The extracted probability distributions, after adjusting for the market's risk aversion generally reveal a higher implied probability of a devaluation. This can be gleaned from the fact that the adjusted probability distributions are mostly shifted to the right of the unadjusted ones (Figure 6 and 7).[16] The dispersion of the adjusted probability distributions are larger, indicating a higher degree of uncertainty about future outcomes (Table 1).

Moments of Risk-Neutral Probability Density Distributions
(Hong Kong dollar)

	7-Aug-98	14-Aug-98	21-Aug-98	28-Aug-98	4-Sep-98
		Risk Neutral			
Mean	8.06	7.83	7.67	8.09	7.71
Standard deviation	0.15	0.23	0.07	0.27	0.17
Skewness	-2.99	2.80	1.90	0.37	3.95
Kurtosis	11.05	10.18	8.23	3.62	20.99
		Risk Adjusted			
Mean	8.02	7.82	7.73	8.13	7.77
Standard deviation	0.14	0.25	0.14	0.31	0.22
Skewness	-3.06	2.40	0.72	0.11	2.14
Kurtosis	11.56	8.31	5.40	2.93	9.42
		Difference Between Risk Adjusted and Risk Neutral			
Mean	-0.04	-0.01	0.06	0.03	0.06
Standard deviation	-0.01	0.02	0.07	0.04	0.05
Skewness	-0.07	-0.40	-1.18	-0.26	-1.81
Kurtosis	0.51	-1.86	-2.84	-0.70	-11.57

Source: Staff estimates.

[16] The exchange rate probability distributions were extracted by adjusting the quoted implied volatility for the market risk aversion using the estimated volatility risk premium as a proxy. However, since only the bid-ask spread around the at-the-money implied volatility was available, it was assumed for the exercise that the spread remained constant over other implied volatilities. In practice, it is likely that the bid-ask spread increases the further away from the money the quote is, which would strengthen the results of this exercise.

Figure 6. Hong Kong Dollar: Impact of the August 14 Intervention on Market Beliefs,
3 Month Options, August 1998.

Panel A

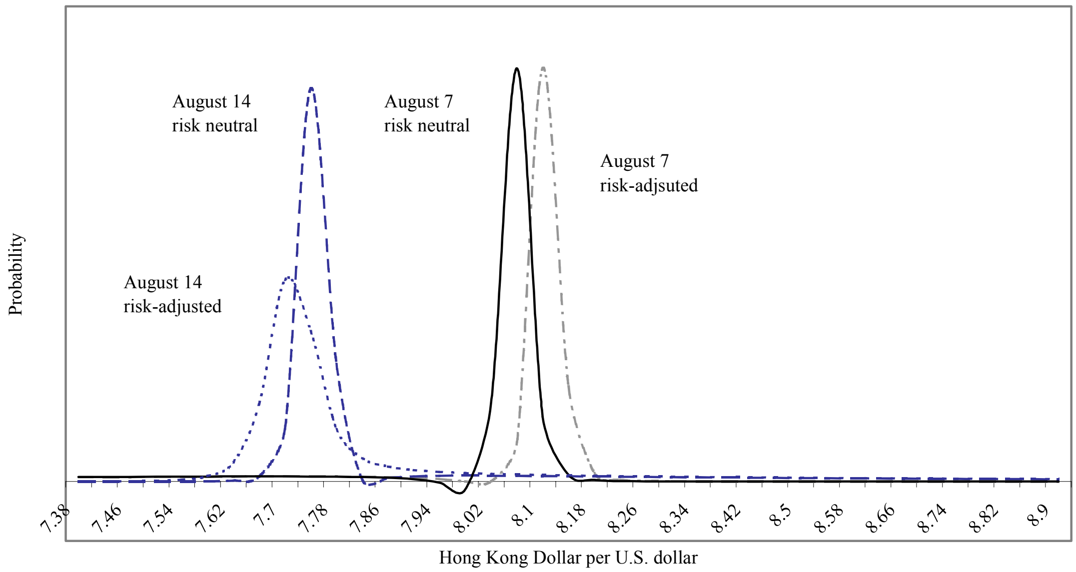

Panel B

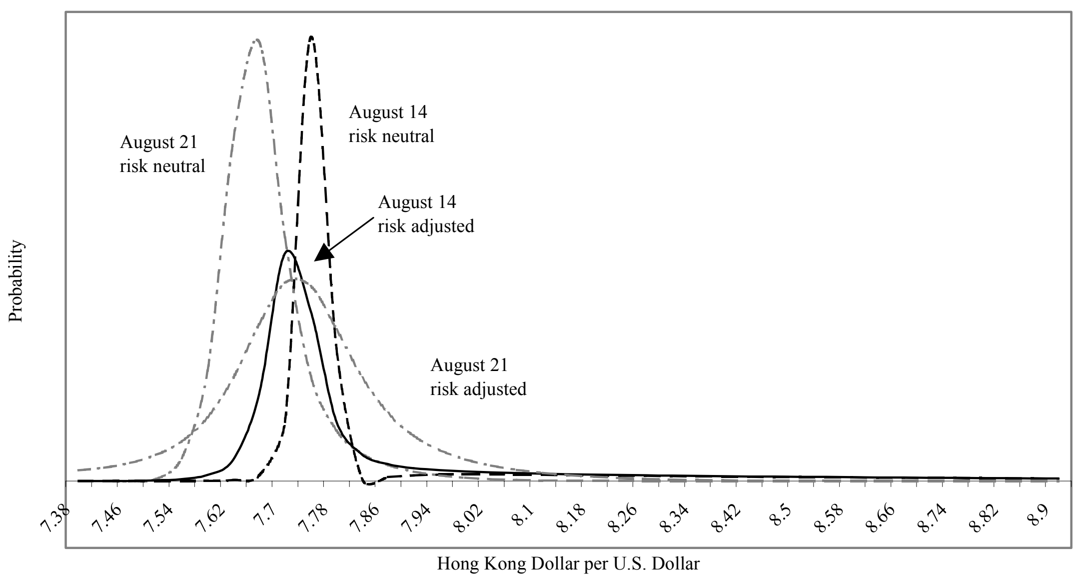

Figure 7. Hong Kong Dollar: Impact of August 28 Intervention on Market Beliefs, 3 Month Options, August 1998.

Panel A

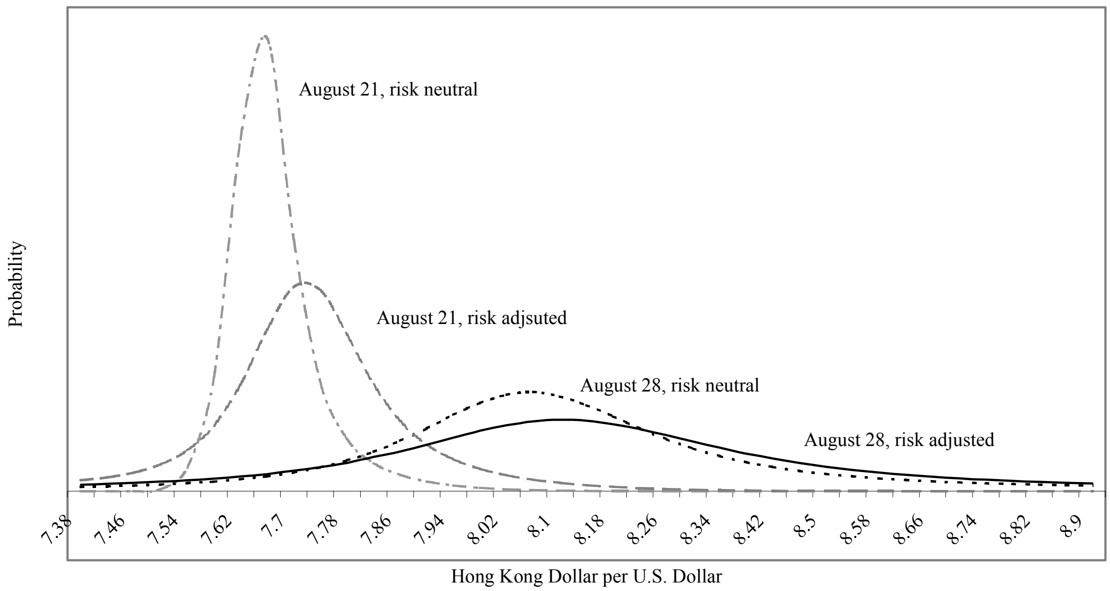

Panel B

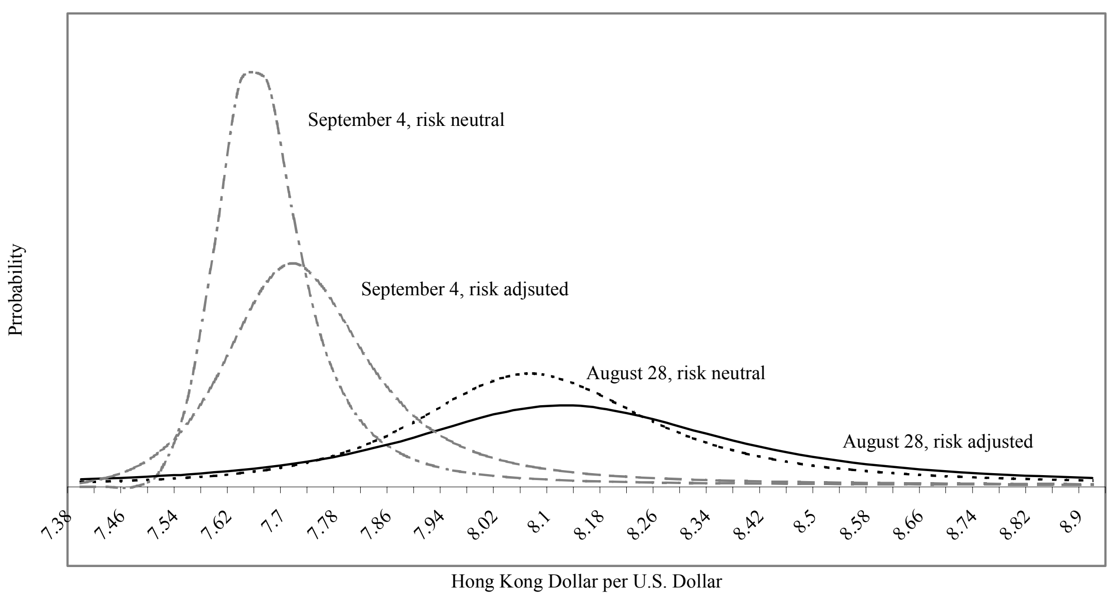

Importantly, the adjusted probability distributions indicate that the impact of August 1998 intervention in calming market sentiment was muted (Figures 6 and 7). The changes in the means of the implied distributions after intervention dates were smaller in the risk-adjusted case than in the unadjusted case (Table 1). At the same time, the standard deviation for risk-adjusted implied distributions did not decline by as much after intervention as it did in the unadjusted case. Thus, the market was not as reassured after the intervention as would be apparent by just considering risk unadjusted probability distribution.

The results can be interpreted to indicate that the intervention did succeed in convincing market makers that the probability of a devaluation per se had receded, reducing market risk. However, there was still a good degree of uncertainty regarding future volatility and liquidity risk in the market, which kept the volatility risk premium high. Thus, while risk itself may have been mitigated, risk preferences had not yet followed as market participants remained cautious.

VI. Conclusion

In answer to the question posed in the title, the volatility risk premium affects the informational content in two important ways: (1) it affects the forecasting power of implied volatility for future realized volatility; and (2) it may change the shape and location of probability distributions extracted from options data to obtain the "market's forecast" of possible future outcomes.

This paper confirmed that implied volatilities are a biased predictor of future realized volatility. Extracting a volatility risk premium ex-post and approximating it with ex-ante observable instruments, allows adjusting the econometric forecasting model for the presence of the risk premium. This adjustment reduces some of the bias found in the original predictability regressions.

While it is possible to use a cross-section of implied volatilities to extract risk neutral exchange rate expectations in the market, it is necessary to adjust for risk preferences to arrive at a measure that is useful for exchange rate forecasting. The method suggested here is to approximate the ex-ante unobservable volatility risk premium with the bid ask spread on implied volatility quotes.

This technique is used to analyze market sentiment on the Hong Kong Dollar around the time of the equity market intervention by the Hong Kong Monetary Authority. Adjusting for risk preferences reveals that market expectations were not as reassured after the intervention as would be apparent by just considering risk neutral expectations. It appears that while the intervention did succeed in mitigating market risk, it did not immediately affect risk preferences.

References

Becker, Torbjørn, Bankim Chadha, and Amadou Sy (1998), "Foreign Exchange Bid-Ask Spreads in the Asian Crises," mimeo, International Monetary Fund (Washington).

Beckers, S., 1981, "Standard Deviations Implied in Option Prices as Predictors of Future Stock Price Variability," Journal of banking and Finance, pp. 363-381.

Bliss, Robert R. and Nicolaos Panigirtzoglou, 1999, "Testing the Stability of Implied Probability Density Functions," Federal Reserve Bank of Chicago Working Paper Series WP 99-21 (Chicago, November).

Breeden, D.T. and R.H. Litzenberger, 1978, "Prices of State-contingent Claims Implicit in Option Prices," Journal of Business, Vol. 51, No. 4, pp. 621-651.

Buraschi, Andrea and Jens Jackwert, 1999, "Is Volatility Risk Priced in Options?" mimeo, London Business School.

Canina, Linda and Stephen Figlewski, 1993, "The Informational Content of Implied Volatility," Review of Financial Studies, Vol. 6, No. 3, pp. 659-681.

Chang, Kevin P. H. and William R. Melick, 1999, "Background Note," Workshop on Estimating and Interpreting Probability Density Functions, Bank for International Settlements, June 14, 1999, available online at: www.bis.org.

Chiras, D., and S. Manaster, 1978, "The Information Content of Option Prices and a Test of Market Efficiency," Journal of Financial Economics, Vol. 10, pp. 29-58.

Christensen, B.J. and N. R. Prabhala, 1998, "The Relation Between Implied and Realized Volatility," Journal of Financial Economics, Vol. 50, pp. 125-150.

Day, T. and C. Lewis, 1992, "Stockmarket Volatility and the Information Content of Stock Index Options," Journal of Econometrics, Vol. 52, pp. 267-287.

Engel, Charles, 1995, "The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence," Journal of Empirical Finance, Vol. 3, pp. 123-192.

Figlewski, Stephen, 1997, Forecasting Volatility, Financial Markets, Institutions & Instruments, Vol. 6, pp. 1-88.

Galati, Gabriele, 2000, "Trading Volumes, Volatility, and Spreads in Foreign Exchange Markets: Evidence from Emerging Markets," Bank for International Settlements Working Paper (Basel, October).

Garman, Mark B. and Steven W. Kohlhagen, 1983, "Foreign Currency Option Values," Journal of International Money and Finance, Vol. 2, pp. 231-237.

Green, Clifton T. and Stephen Figlewski, 1999, “Market Risk and Model Risk for a Financial Institution Writing Options,” Journal of Finance, Vol. 54, No. 4, pp. 1165-1499.

Hansen, L. P., 1982, “Large Sample Properties of Generalized Method of Moments Estimators,” Econometrica, Vol. 50, pp. 1029-1054.

Heston, S.L., 1993, “A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options,” Review of Financial Studies, Vol. 6, pp. 327-344.

Huang, Roger D. and Ronald W. Masulis, 1996, “FX Spreads and Dealer Competition Across the 24 Hour Trading Day,” mimeo, Owen Graduate School of Management, Vanderbilt University (Nashville, December).

Hull, J., and A. White, 1987, “The Pricing of Options on Assets with Stochastic Volatility,” Journal of Finance, Vol. 42, pp. 281-300.

International Monetary Fund, 1999, International Capital Markets: Developments, Prospects, and Key Policy Issues, World Economic and Financial Surveys (Washington, September).

Jondeau, Eric and Michael Rockinger, 2000, “Reading the Smile: The Message Conveyed by Methods which Infer Risk Neutral Densities,” Journal of International Money and Finance, Vol. 19, pp. 885-915.

Jorion, Philippe, 1995, “Predicting Volatility in the Foreign Exchange Market,” The Journal of Finance, Vol. 50, No. 2, pp. 507-528.

Lamoureux, D., and W. Lastrapes, 1993, “Forecasting Stock-Return Variance: Toward an Understanding of Stochastic Implied Volatilities, Review of Financial Studies, Vol. 6, pp. 293-326.

Latané, H. and R. Rendleman, 1976, “Standard Deviation of Stock Price Ratios Implied by Option Premia,” Journal of Finance, Vol. 31, pp. 369-382.

Malz, Allan M., 1997, “Estimating the Probability Distribution of the Future Exchange Rate from Option Prices,” Journal of Derivatives, Winter, pp. 18-36.

Malz, Allan M., 2000, “Do Implied Volatilities Provide Early Warning of Market Stress?” RiskMetrics Journal, Vol. 1, pp. 41-60.

Melick, William R. and Charles P. Thomas, 1997, “Recovering an Asset’s Implied PDF from Option Prices: An Application to Oil Prices During the Gulf Crises,” Journal of Financial and Quantitative Analysis, Vol. 31, No. 1, pp. 91-115.

Pasquariello, Paolo, 2000, "The Microstructure of Currency Markets: An Empirical Model of Intra-day Returns and Bid-Ask Spread Behavior," mimeo, Stern School of Business, New York University (New York, August).

Poteshman, Allen M., 2000, "Forecasting Future Volatility from Option Prices," mimeo, University of Illinois at Urbana-Champaign.

Wei, Shang-Jin, 1994, "Anticipation of Foreign Exchange Volatility and Bid-Ask Spreads," NBER Working Paper No. 4737 (Cambridge, MA, May).

Appendix 1: Risk Neutrality

To understand the concept of risk neutrality consider the following example. Suppose a non-profit bookmaker accepting bets on both teams in a football game knows that the true probabilities that each team wins are equal, i.e. ½. However, among the 10 people placing bets opinions are divided differently. Eight people would like to bet that team A will win and two people are convinced that team B will win. If the bookmaker were to accept bets based on the true probabilities—paying out \$2 for every successful and \$0 for every unsuccessful \$1 bet—he may incur a \$6 loss if team A wins: He needs to pay out \$16 on the winning bets, but received only \$10 for all bets placed (Table 1). If team B wins, the bookmaker would make a profit of \$6, as he only needs to pay out \$4 for the winning tickets from the \$10 collected (Case 1). While on average the bookmaker would break even if the game is repeated several times, he bears the risk for any individual game.

In order to break even for each game and thus not bear any risk, the market maker may adjust the payout ratios to reflect aggregate risk preferences in the "market". In this case, he would pay only \$1.25 on the popular bet that team A wins, but would reward the successful bet on team B with \$5. Adjusting the payouts such that the bookmaker has a risk-neutral position ensures that he always breaks even: If team A wins, his payout (\$1.25 x 8 bets) is exactly covered by his revenue (\$10). If team B wins, his payout (\$5 x 2 bets) is again covered by the wagers received (\$10) (Case 2). The relative supplies of bets on the two teams prescribes the "market consensus" and implies the risk neutral probabilities of victory—80 percent for team A and 20 percent for team B. However, these expectations differ from the true probabilities.

Table 1. Implied Probabilities in Hypothetical Bet.

<table>
<tr><th colspan="2" rowspan="2"></th><th rowspan="2">Team A wins</th><th rowspan="2">Team B wins</th><th colspan="2">Implied Probability</th></tr>
<tr><th>Team A</th><th>Team B</th></tr>
<tr><td rowspan="2">Case 1: Market maker uses true probabilities</td><td>payout</td><td>$2 x 8 bets = $16</td><td>$2 x 2 bets = $4</td><td rowspan="2">1÷2 = 50%</td><td rowspan="2">1÷2 = 50%</td></tr>
<tr><td>profit</td><td>- $6</td><td>+ $6</td></tr>
<tr><td rowspan="2">Case 2: Market maker uses risk neutral probabilities</td><td>payout</td><td>$1.25 x 8 bets = $10</td><td>$5 x 2 bets = $10</td><td rowspan="2">1÷1.25 = 80%</td><td rowspan="2">1÷5 = 20%</td></tr>
<tr><td>profit</td><td>$0</td><td>$0</td></tr>
</table>