Is there a premium for currencies correlated with volatility? Some evidence from risk reversals[1]

Henri Pagès

Introduction

Are option price data useful in predicting exchange rate changes? Option prices reflect the market perceptions of the underlying asset's distribution, so they may reveal information about the exchange rate's future moves. Among option-based indicators of market sentiment in the foreign currency market, two seem particularly relevant for central banks. One is *at-the-money* volatility, the market's implicit volatility forecast for those options whose strike price is closest to being at the forward rate. The other is given by the price of *risk reversals*. Risk reversals are derivative instruments constructed as a linear combination of two out-of-the-money put and call options, written on the same currency and expiring at the same date. Their payoffs can turn either positive or negative for large deviations of the exchange rate from the forward rate, depending on the direction of the move. For this reason, they are often interpreted as the market's best guess about the directional bias of future exchange rate moves. This paper sets out to test whether actual changes in the future spot rate are indeed related to developments in the foreign currency options market.

Predicting asset returns is one of the main concerns of the "efficient-markets hypothesis" literature. In foreign exchange markets, forecasts are generally measured in variation from the corresponding forward rates, where the difference is referred to as the "expected return to speculation" or, up to a change in sign, the "forward bias". The simplest version of efficiency states that the mean return to speculation in the foreign exchange market, conditioned on available information, is zero. Even though this simple test is not in general borne out by the data, there are still some advantages in carrying it out. First, forward rate prediction biases may be correlated with available information which is not clearly identified by theory, but nevertheless helps predict future spot rates. Second, the metrics of efficiency tests is a convenient way to ascertain whether certain variables can be brought to bear on the rejection of rational expectations, and thus provide evidence in favour of specific alternative hypotheses: for instance, that rejection is due to the existence of a risk premium or to the role of some expectational errors.

The evidence gathered in this paper tends to support the view that information revealed by option prices helps improve forecasts of future spot rates. However, it is different from the traditional view regarding risk reversals' directional bias, according to which high positive (resp. negative) risk reversals are attributable to market perceptions that the leading currency is likely to surge (resp. plunge) in value. Rather, it points to a consistent correlation between risk reversals and the forward bias: when the price of risk reversals goes up, the leading currency's forward rate tends to *increase* with respect to future realisations of the spot rate. If there is a downward forward bias, implying that on average the forward rate is below the future spot rates, the bias will be reduced. Conversely, if the bias is upwards, indicating an overestimation of realised future spot rates, the bias will be increased. Because higher risk reversals tend to *lower* the future spot rate relative to the forward rate, they are *not* in general associated with an appreciation of the leading currency: the net result depends on the concurrent shifts in the forward rate (possibly spurred by central bank

1 Preliminary and incomplete draft. I am indebted to the Bank of France's Direction Générale des Services Étrangers for providing the data set. I benefited from stimulating discussions with A. Duchateau, M.-O. Strauss-Kahn, participants from an internal workshop as well as the 1995 Autumn Meeting of Central Bank Economists at the BIS. All errors and opinions expressed are mine.

intervention). However, they help narrow down the forecast error, and so uncover more precisely the mean future spot rate from the observed forward rate.

One interpretation envisaged in the paper, as in many others, is that the bias of forecast errors may stem in part from a risk premium required by risk-averse investors. In determining their forecasts of a risky currency, investors may include a risk premium in the return differential, which would cause the forward rate to be a biased estimate of the future spot rate. Of course, a stronger case for this view could be made if the forward bias was tied to variables which theory links to the risk premium. The empirical results reported here obviously call for some structure to interpret them. Because the very existence of a risk premium reveals departure from the efficient-markets hypothesis, its identification requires in principle a specification of consumers' preferences and information sets, of the technology available for producers, and of the risks inherent in the economy. Although the paper falls short of providing any argument rooted in general equilibrium theory, it uses a fairly simplified version of the portfolio balance model developed by Kouri (1977) [12] and Dornbusch (1983) [6], with an important new feature: it recognises the previously neglected dimension of *volatility risk*, much in the spirit of recent stochastic volatility models.

Stochastic volatility arises when conditional second moments are not only variable, but also follow some dynamics in their own right. In this case there are two distinct sources of risk, one which relates to innovations in the exchange rate (exchange rate risk) and the other to innovations in its volatility (volatility risk). They will, in general, be only partially correlated with each other, and the correlation may vary over time. Hence the dynamics of the exchange rate can be characterised by a time-varying volatility and a time-varying correlation between the exchange rate and its volatility. Since the risk premium can be theoretically expressed in terms of those last two variables, the paper's results can be viewed as a test of the existence of a risk premium, where the information revealed by foreign currency options is exploited to measure the market expectations of the instantaneous volatility and its comovements with the spot rate, respectively.

A key question raised by the paper's efficiency tests is whether information imparted from the foreign currency options market may be considered as properly parametrising the risk premium. The two kinds of variables used in the paper's econometric analysis should be viewed only as an approximation: relying on such market-based indicators is a short-cut to avoid the technical difficulties of estimating the dynamics of the exchange rate volatility pair. Despite its drawback, this approach has some merits. It uses observed implicit volatilities which, in contrast to an ARCH modelling, do not depend on specific assumptions about squared errors, and it rests on the market's own method of conditioning. Moreover, it takes advantage of modern advances in the theory of stochastic volatility, according to which the distinction between volatility per se and spot/volatility comovements can explain the biases found in the prices of options. It turns out that instantaneous volatility is well captured by at-the-money volatility, and that comovements between the spot rate and its volatility produce *skewness* of the distribution, which in turn causes the price of risk reversals to adjust. This vindicates the paper's use of at-the-money volatility as a measure of the exchange rate time-varying volatility, and of risk reversals as a measure of the time-varying spot/volatility correlation.

There exists by now a vast literature testing the foreign exchange risk premium. Frankel (1982) [7] has estimated a coefficient of relative risk aversion under the assumption of constant conditional second moments, and was unable to reject the null of no risk premium. More recently, Lyons (1988) [14] has used option-implied volatilities for three currencies to identify a time-varying risk premium, and found evidence for it, although his data did not give strong support to the balance portfolio approach of the risk premium.

A substantial body of empirical research has also been aimed at testing the efficiency of option prices as predictors of future exchange rate *second* moments. The null hypothesis is that implied volatilities from the Black-Scholes model map well into the (square root of the) observed future variance. The overall conclusion is that at-the-money volatility is excessively variable (Wei and Frankel (1991) [20]) and that out-of-the-money volatilities are overvalued for both call and put options (Bodurtha and Courtadon (1987) [4], Borensztein and Dooley (1987) [5]), implying an apparent over-estimation of the likelihood of exchange rate changes.

Given the failure of second moment efficiency tests for options markets, researchers have sought to match the excess volatility biases with alternative option pricing models. Bates (1990) [1] fitted option prices to an asymmetric jump-diffusion process with constant volatility and argued that non-zero risk reversal prices are attributable to a crash premium, reflecting the probability that there will be a jump depreciation in the dollar. Malz (1994) [15] employed a similar method to calculate realignment probabilities for the French franc and pound sterling, but provided no confidence interval for the estimates derived. Bates (1993) [2] has developed a general stochastic volatility/jump diffusion model, and found that there were no significant expectations of exchange rate jumps in the dollar/Mark market.

The balance of the paper is organised as follows. In Section 1, some linear forecasting equations of the forward bias are presented, where at-the-money volatility and risk reversals prices compete with some other traditional variables such as past forward biases or past changes in forward rates. In Section 2, the canonical stochastic volatility model is briefly introduced and some informal arguments are made to convey the idea that risk reversals capture the *comovements* between spot rates and their volatility. Then, it is argued, the typical balance portfolio approach of Kouri and Dornbusch may be invoked to show that the risk premium is not only a function of the underlying variance, as it should be, but also of the correlation between spot returns and volatility, which risk reversals are known to reflect. In the conclusion, a reinterpretation of Section 1's results is offered regarding the impact of risk reversals on forward rates by contending that some currencies, like the dollar vis-à-vis the Deutsche Mark or the yen, may have been more sought after by international investors because they allowed some hedging against volatility risk.

1. Estimating the forward bias

1.1 Data description

Option-implied data were retrieved from over-the-counter markets. These markets have developed since the early 1980s and have become larger than organised exchange markets since the mid-1980s. Risk reversals were traded as option-based derivative instruments before the end of the 1980s, but reported data are not considered reliable for European cross-currencies before 1992 or 1993: at the time, ERM crises contributed to drawing attention to these instruments, increasingly used by end-users as a low-cost way to hedge against large changes in exchange rates and by speculators to take leveraged positions. The data used in this paper cover all trading days from 2nd November 1994 to 29th September 1995 on the Mark/franc, dollar/Mark and dollar/yen markets. Quotes are expressed in terms of implied volatilities ("vols") in percentage per year, so that no transformation on the basis of the Black-Scholes formula [3] was necessary.

The exact definition of at-the-money volatility and risk reversals is relegated to an appendix, where the Garman-Kohlhagen model [8] commonly used in calculating European currency options is also provided. At the Bank of France, at-the-money volatilities are read from a Reuters screen edited by Société Générale. As for risk reversals, Société Générale and another bank fax quotations each trading day, also in volatilities, as representative of the prices of the previous evening; they correspond to, respectively, the one-month (Mark/franc, dollar/Mark and dollar/yen) and three-month (Mark/franc) time to expiry. Naturally, all implied volatilities were selected so as to match the maturity of risk reversals.

Spot exchange rates for these currencies are those reported every day at 2.15 p.m. by central banks participating in the "concertation group". They correspond to the average of bid-ask prices. Forward exchange rates were derived from the spot rates on the basis of the relevant interest rates, taken from Reuters on Euro-currency markets.

Some transformation was necessary to compute lag or lead values of the variables. This was necessary because they frequently fall during weekends or holidays. For options, the following convention is adopted by the market: when the settlement date occurs during a day off, it is assigned to the first following trading day, except if this would change the settlement month, when it is assigned to the last trading day. The same convention has been adopted there, but there are some problems. First, the settlement day is different from the exercise date (it comes generally two days later). Second, the convention adopted is unjustified for the forward rate. Both will create biases, but it is unclear how important they are. The question has not been addressed in the paper.

1.2 Empirical results

The hypothesis tested in this section is that the option-based indicators defined above have predictive content with regard to the forward bias. Let $fb_{t+k} = f_{t,k} - s_{t+k}$ be the forward bias, where $f_{t,k}$ is the log of the k-step ahead forward rate set at time t, s_{t+k} the log of the realised spot rate at time $t+k$, and k the forecast horizon, equal to one or three months. Forward rate unbiasedness implies that fb_{t+k} has zero mean and is uncorrelated with I_t, the information set available at time t. The problem is thus to estimate the parameters a, b, c and β in the k-step ahead linear forecasting equation:

$$E[fb_{t+k} | I_t] = a + b\text{var}_t + c\,\text{rrv}_t + x_t\beta\,, \tag{1}$$

where var_t is the square of the option-implied volatility from t to $t+k$ observed in the market at time t, rrv_t is the price at t, in volatility terms, of 25-δ risk reversals expiring at time $t+k$, and x_t is a row vector of variables contained in I_t, like past forecast errors or past rates of change of forward rates.[2] To get some preliminary insights into the basic correlations between the selected market indicators and the forecast error, some "plain vanilla" tests are first carried out when other effects are assumed away, i.e. when $\beta = 0$. Alternative regressions are then presented. Using terminology from the efficient-markets literature, both weak and semi-strong forms are considered, depending on whether or not data from other exchange markets are included in the regression.

In all the regressions presented, the sample data are tightly overlapping. As a result, consecutive forecast errors will be serially correlated. Ordinary least squares would yield consistent estimates, because (1) implies that forecast errors are not contemporaneously correlated with the right-hand-side variables. However, the estimated variances would be biased. Generalised least squares would yield inconsistent estimates, because the transformed variables would violate the orthogonality conditions implied by (1). As shown by Hansen and Hodrick (1980) [9], this arises because the forecast error is correlated with *future* right-hand-side variables. Intuitively, recent errors made by the forecaster, which cannot yet be detected at the current time, contaminate both the *current* and the *future* predictions, resulting in inconsistent estimation. Finally, the usual practice of extracting non-overlapping data from the given sample to circumvent the problem of serial correlation would have dramatic consequences here in terms of loss of information. In this context, the Generalised Method of Moments estimation has been implemented with a window equal to the maximum number of trading days in the forecast horizon (i.e. 22 or 66 observations).[3]

Consider first the regression equation:

$$fb^i_{t+k} = a^i + b^i\,\text{var}^i_t + c^i\,\text{rrv}^i_t + \varepsilon^i_{t,k}, \tag{2}$$

2 As usual, the maintained assumption throughout the analysis is that the conditional expectation of the left-hand-side variable is a linear combination of elements in I_t, and that all relevant variables are those included on the right-hand side.

3 To avoid computational difficulties in estimating the covariance matrix, the "damp" parameter was set to 1 in all regressions, which is the smallest value which guarantees a positive-definite matrix, even though a smaller number was sometimes sufficient.

where $\varepsilon^i_{t,k}$ is the k-step ahead forecast error, and i one of the four exchange rates considered. The results of regression 2 are presented in Table 1. In general, all point estimates are insignificant even at high significance levels, except for the Mark/franc rate. In this case, the two variance estimates are significant at the 95% confidence level. The coefficient is positive, implying that a higher variance raises the forward rate with respect to the future spot rate and so increases its magnitude, given that the mean forward bias is positive (the upward bias is approximately 1%). In this case, speculators are penalised by buying the Deutsche Mark forward, indicating that the leading currency is rather the "safe" currency. The case of risk reversals is more confusing. For the Mark/franc regressions, the coefficient is significant at the 10% level when k=3, but not when k=1. This casts doubts on the reliability of the Mark/franc one-month risk reversal series. The estimates for the other two markets are not significant. In all, these simple regressions do not provide much insight, especially with respect to risk reversals. Evidence against the null that all coefficients are zero cannot be found for the dollar/Mark and the dollar/yen markets, even at very high significance levels (37% and 34% respectively).

Table 1

Plain vanilla tests

($fb^i_{t+k} = a^i + b^i \, var^i_t + c^i \, rrv^i_t + \varepsilon^i_{t,k}$ (2) for currency i, subscripts as below)

Currency	$\hat{a}$	$\hat{b}$	$\hat{c}$	R^2	SEE	obs
DM/French franc	- 4.0	0.39	- 4.0			
(1 month)	(5.2)	(0.18)	(9.8)	0.22	13.0	207
	0.44	0.03	0.68			
DM/French franc	- 10.3	0.33	5.6			
(3 months)	(2.0)	(0.05)	(3.3)	0.78	4.3	165
	<0.01	<0.01	0.09			
US$/DM	28.4	- 0.12	7.3			
(1 month)	(24.7)	(0.11)	(10.5)	0.10	41.0	207
	0.25	0.32	0.49			
US$/yen	18.0	- 0.21	- 10.7			
(1 month)	(27.1)	(0.13)	(10.3)	0.08	62.6	207
	0.50	0.11	0.30			

Note: SE in parentheses, then marginal significance level.

Consider now the weak form tests, in which only the information from the own exchange rate is allowed to have non-zero coefficients. In testing exchange market efficiency, researchers typically include past forecast errors or past realisations of the forward bias. To avoid unbalanced regressions in the presence of non-stationarity, one may invoke cointegration between s_{t+k} and $f_{t,k}$, a necessary condition for market efficiency. It is thus logical to include the error-correction term $fb_{t-k,k} = f_{t-k,k} - s_t$,[4] the forward bias that results from the forecast k-months ago. Assume for simplicity that the forecasting equation can be written as:

$$s_{t+k} - s_t = \alpha(f_{t,k} - f_{t-k,k}) + \beta(f_{t-k,k} - s_t) + \text{other terms.}$$

4 In principle, one could test that the cointegrating vector is one, i.e. that the forward rate and future spot rates never drift apart. Unfortunately, due to the non-stationary nature of the spot and forward rates, it is difficult to implement a formal test because the standard error of the cointegrated vector is not consistent. The author did not correct the reported standard errors as in West (1986) or use any other method, because all the procedures involved are quite sensitive to assumptions regarding the data-generating process.

This can be put in terms of past forward biases and past forward changes as:

$$fb_{t+k} = (1-\alpha)(f_{t,k} - f_{t-k,k}) + (1-\beta) fb_t + \text{same terms.}$$

Table 2 reports estimation of the regressions:

$$fb^i_{t+k} = a^i + b^i \, \mathrm{var}^i_t + c^i \, \mathrm{rrv}^i_t + d^i\left(f^i_{t,k} - f^i_{t-k,k}\right) + e^i fb^i_t + \varepsilon^i_t, \tag{3}$$

for all i = 1a, ... pairs considered. The null of unbiasedness can now be rejected for all exchange rates with the appropriate χ^2 (5)-test except for the dollar/Mark rate, where the significance level is still as high as 24%. Table 3 presents some complementary tests on exclusion restrictions. Compared with the simple regressions, there is a dramatic increase in confidence of *both* indicators in the case of the dollar/Mark and the dollar/yen rates, although the dollar/Mark coefficients are at the verge of significance at the 10% level. The signs of the variance parameters still point to the dollar and the franc as the relatively "risky" currencies, although there appears to be some conflicting evidence between the one-month and three-month ahead regressions in the case of the franc. Risk reversals, when they are significant, have a positive impact on forward biases, implying that higher values of the variable raise the current forward rate relative to the future spot rates. Again, the Mark/franc case stands out because the forward bias seems impervious to movements in the risk reversals prices. The only pattern that arises is the strong significance of the forward discount, $f_{t,k} - s_t$, in the rejection of the unbiasedness hypothesis for the one-month regression. Since the sum of the last two variables, $f_{t,k} - f_{t-k,k}$ and fb_t, is precisely $f_{t,k} - s_t$, one may interpret the result as evidence that a 1% rise in the one-month return differential between France and Germany would *raise* the current forward rate relative to the future spot rate by approximately 1.6%. Thus, in the French/German case, the interest differential tends to obscure the predictive content of risk reversals. This is in sharp contrast with the other two exchange rates, for which lagged forward biases and forward rates of changes are dominated by both option-implied variables.

Table 2

Weak form tests

($fb^i_{t+k} = a^i + b^i \, \mathrm{var}^i_t + c^i \, \mathrm{rrv}^i_t + d^i\left(f^i_{t,k} - f^i_{t-k,k}\right) + e^i fb^i_t + \varepsilon^i_t$ (3) for market i, subscripts as below)

Currency	$\hat{a}$	$\hat{b}$	$\hat{c}$	$\hat{d}$	$\hat{e}$	R^2	SEE	obs
DM/French franc	- 4.4	0.32	- 5.5	1.6	1.6			
(1 month)	(5.4)	(0.21)	(0.007)	(0.6)	(0.56)	0.27	13.3	186
	0.42	0.12	0.52	<0.01	<0.01			
DM/French franc	- 6.0	0.42	0.3	0.04	0.5			
(3 months)	(1.1)	(0.03)	(0.5)	(0.3)	(0.3)	0.86	3.6	101
	<0.01	<0.01	0.54	0.89	0.07			
US$/DM	42.9	- 0.15	18.9	0.5	0.7			
(1 month)	(23.9)	(0.09)	(11.0)	(0.5)	(0.5)	0.20	39.7	186
	0.07	0.11	0.09	0.33	0.19			
US$/yen	63.1	- 0.31	27.2	- 0.3	0.7			
(1 month)	(20.9)	(0.10)	(11.1)	(0.5)	(0.6)	0.49	49.1	186
	<0.01	<0.01	0.01	0.63	0.20			

Note: SE in parentheses, then marginal significance level.

The data presented in Table 2 reflect certain episodes which were marked by changing market perceptions about the prevailing economic environment. For example, the Japanese Government announced in early August a programme to overhaul the economy, which was followed by the Bank of Japan's own interventions to support the weak dollar. Structural stability tests over the full sample with dummies times regressors indicate that this may have been the case. The results are reported in Table 4.

Table 3

Indicators exclusion restriction

(Test that $b^i = c^i = 0$ for market i, subscripts as below)

Currency	Vanilla	Weak form	Semi-strong
DM/French franc	7.6	2.6	9.7
(1 month)	0.02	0.27	<0.01
DM/French franc	110	166	326
(3 months)	<0.01	<0.01	<0.01
US$/DM	3.1	5.7	7.5
(1 month)	0.21	0.06	0.02
US$/yen	3.2	15.5	13.2
(1 month)	0.21	<0.01	<0.01

Note: χ^2 (3), then marginal significance level.

For the yen the cut was set at 1st July 1995 and both the χ^2-test and the Bonferroni test on the separately induced hypothesis that all coefficients are zero in the subsample provide strong evidence against stability. Interestingly, the risk reversal variable appears to be also responsible for rejection, now with a *negative* coefficient, indicating that the relative position of the forward rate has now become a decreasing function of risk reversals. The dollar/Mark and Mark/French franc forecasting equations are also unstable, but do not single out risk reversals as a cause for rejection. The case of the French franc is a bit more contentious, for the χ^2 and Bonferroni tests conflict at the 95% confidence level, with a marginal significance level of only 0.05/5 = 0.01 for the separately induced tests of the hypothesis that the maximum of all coefficients is zero in the subsample. Hence, one can barely conclude that the Mark/franc evinces instability with the given data at the 95% confidence level, unless one has prior knowledge about the possible causes for rejection. Weak form tests appear to be more powerful tests of the option prices predictive content, but the forecasting equations are unstable.

Table 4

Structural stability tests

(Test that all coefficients are zero in the subsample - weak form, equation (3))

Currency	$\hat{a}^i$	$\hat{b}^i$	$\hat{c}^i$	$\hat{d}^i$	$\hat{e}^i$	$\chi^2(5)$	Subsample
DM/French franc	27	- 0.14	- 4.3	- 0.9	- 0.4	115	95:5:1-95:29:9
(1 month)	(10)	(0.23)	(11)	(0.5)	(0.5)		
	0.01	0.55	0.70	0.08	0.49	<0.01	
US$/DM	- 168	0.68	- 29.0	- 0.6	- 1.8	43	95:5:1-95:29:9
(1 month)	(33)	(0.20)	(28)	(0.5)	(0.5)		
	<0.01	<0.01	0.30	0.23	<0.01	<0.01	
US$/yen	- 191	- 0.02	- 52.0	1.9	0.2	109	95:7:1-95:29:9
(1 month)	(29)	(0.20)	(12.3)	(0.8)	(0.8)		
	<0.01	0.88	<0.01	0.02	0.83	<0.01	

Note: SE in parentheses, then marginal significance level.

Finally, the tests are expanded to include information from all exchange rates. The semi-strong form is here written as:

$$fb^i_{t+k} = a^i + b^i \,\mathrm{var}^i_t + c^i \,\mathrm{rrv}^i_t + \sum_j d^{ij}\left(f^j_{t,k} - f^j_{t-k,k}\right) + \sum_j e^{ij}\, fb^j_t + \varepsilon^i_t, \qquad (4)$$

where the j index refers to the one-month own and other two forward rate markets. (The three-month Mark/franc variables have not been included in the regressions, except in the three-month Mark/franc own forecasting equation; conversely, the three-month Mark/franc forecasting equation does not include the one-month Mark/franc data.) The results seem more satisfactory and are reported in Table 5. The last but one column gives the standard error of estimate of the four regressions; the corresponding standard error of dependent variable is 15.3, 7.1, 43.0 and 67.7, in annual percentage. The coefficients of the option-implied variables are all significant, except that of the one-month Mark/franc risk reversal. The magnitude of the variance coefficients is consistent with the balance portfolio approach alluded to in Section 2.2, which equates them to the product of a risk aversion parameter and of the difference between two currency shares. Their signs indicate that the dollar and the franc are the two relatively risky currencies in the sense that increases in variance translate into a lower forward rate for the dollar (excess supply of dollars relative to the Mark and the yen), and into a higher forward rate for the Mark (excess demand for Marks relative to the franc). The risk reversals coefficients are also quite significant. They conform to the pattern of the former regressions, although the Mark/franc forward bias appears to be less sensitive to variations in risk reversals than the dollar/Mark or the dollar/yen. The evidence is weaker for the one-month Mark/franc risk reversals, where the null hypothesis of no predictive content is not rejected at the 10% significance level.

In all, the inclusion of more expectational variables, like past forward biases or past forward rates, seems to enhance the role of market-based indicators. The results indicate that risk reversals were positively correlated with the forward bias over the whole period. In the following section, a simple model is presented to help explain the influence of comovements between spot rates and volatility on the forward bias.

Table 5

Semi-strong tests

($fb^i_{t+k} = a^i + b^i \,\mathrm{var}^i_t + c^i \,\mathrm{rrv}^i_t + \sum_j d^{ij}\left(f^j_{t,k} - f^j_{t-k,k}\right) + \sum_j e^{ij}\, fb^j_t + \varepsilon^i_t$ (4) where ij refers to the regression coefficient of currency i on currency j; subscripts as below)

Currency	$\hat{a}^i$	$\hat{b}^i$	$\hat{c}^i$	$\hat{d}^{i1}$	$\hat{d}^{i2}$	$\hat{d}^{i3}$	$\hat{e}^{i1}$	$\hat{e}^{i2}$	$\hat{e}^{i3}$	R^2	SEE	obs
DM/French franc	-20.4	0.34	14.4	1.4	- 0.2	0.4	1.4	0.1	0.1			
(1 month)	(7.8)	(0.17)	(10.0)	(0.4)	(0.2)	(0.2)	(0.4)	(0.2)	(0.2)	0.48	11.3	186
	0.01	0.05	0.15	<0.01	0.19	0.04	<0.01	0.79	0.48			
DM/French franc	- 8.9	0.29	6.7	1.4	- 0.2	0.2	1.7	- 0.2	0.2			
(3 months)	(1.1)	(0.03)	(0.6)	(0.2)	(0.1)	(0.1)	(0.2)	(0.1)	(0.1)	0.86	2.7	101
	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	0.02			
US$/DM	64.1	- 0.19	35.5	- 1.2	1.5	- 1.0	- 2.3	0.9	- 0.4			
(1 month)	(21.8)	(0.10)	(16.1)	(1.3)	(0.7)	(0.5)	(1.6)	(0.8)	(0.5)	0.48	32.3	177
	<0.01	0.06	0.03	0.35	0.02	0.05	0.14	0.24	0.43			
US$/yen	55.2	- 0.22	21.6	0.2	1.4	- 0.8	- 2.3	0.6	0.3			
(1 month)	(23.7)	(0.11)	(6.4)	(1.8)	(0.8)	(0.5)	(2.0)	(0.9)	(0.5)	0.69	38.4	186
	0.02	0.04	<0.01	0.92	0.07	0.09	0.25	0.49	0.62			

Note: SE in parentheses, then marginal significance level.

2. Interpretation of the results

The econometric results just presented raise two questions: do they provide evidence that a risk premium exists? And, if so, how is the risk premium related to option-based indicators? In this section, a stochastic volatility model is sketched to show how risk reversals can mirror variations in comovements between spot returns and volatility, as opposed to reflecting realignment risks. Then the model is specialised in an effort to appeal to balance portfolio arguments and derive the equilibrium *spot risk premium*.

2.1 What do risk reversals measure?

Risk reversals, it is argued, capture the *skewness* of the exchange rate distribution. Exchange rate shocks seem to be asymmetric, in that their magnitude tends to vary according to whether the spot rate appreciates or depreciates. But there are different interpretations. The theory of stochastic volatility explains where this asymmetry may come from and how it can be observed in the markets.

Classical option pricing theory requires that exchange rates follow a geometric Brownian motion, which implies in particular constant second moments. Market-makers have known for a long time that this key assumption is flawed and use standard implied volatilities only as a convenient way to price options. Most empirical investigations concerning nominal returns, starting with Westerfield (1977) [22], have found that nominal exchange returns violate the normality assumption in at least three respects. Their distribution appears to have a time-varying variance, with contiguous periods of high and low volatility; it has fat tails, implying that for a given variance there is a higher probability of large deviations from the mean; and, finally, it is skewed, in that an appreciation and a depreciation of a given size are not equally likely. As a result, alternative models have been developed to generalise the Black-Scholes formula by allowing volatility to change randomly; e.g. Hull and White (1987) [11], Scott (1987) [19] and Wiggins (1987) [23]. Melino and Turnbull (1991) [16] found that these models explain well the price of currency options, although they tend to overestimate volatility. Since then, they have been extended by Melino and Turnbull (1991) [10] and Leblanc (1994) [13] to provide closed-form solutions for arbitrary correlation between asset returns and volatility. Nelson and Foster (1994) [17] show how they can be best be approximated by optimally chosen univariate ARCH models. They are briefly outlined below.

Let the spot rate be S_t, which gives units of domestic currency per unit of foreign currency, and Y_t the *unobservable* process controlling the instantaneous conditional variance of S_t, namely $\sigma(Y_t)^2$. The general formulation which includes stochastic volatility can be written as:

$$\frac{dS_t}{S_t} = \left(\mu_t - i_t^*\right)dt + \sigma(Y_t)\left(\sqrt{1-\rho^2(t,Y_t)}dW_t^1 + \rho(t,Y_t)dW_t^2\right),$$

$$dY_t = \eta_t + \gamma(t,Y_t)dW_t^2, \tag{5}$$

where μ_t is the expected instantaneous return on the leading currency, i_t^* the nominal interest on foreign bonds, $\gamma(t,Y_t)$ is the volatility of volatility, and $\rho(t,Y_t)$ the instantaneous correlation between the spot process and its volatility. Uncertainty is generated by the bivariate standard Brownian motion $(W_t^1, W_t^2)'$. This is an incomplete market framework since, when the domestic currency is chosen as the numeraire, there is only one asset for two sources of shocks. It is well known that in this setup the absence of arbitrage is equivalent to the existence of a risk-neutral probability, under which the spot process S, discounted at rate i (the nominal interest on bonds denominated in home currency), and paying dividend i^*, is a martingale. This implies that the foreign currency *risk premium* $RP_t = \mu_t - i_t$ verifies:

$$\frac{RP_t}{\sigma_t} = \lambda_t\sqrt{1-\rho_t^2} + \nu_t\rho_t,$$

where λ and ν are the risk premia associated with W_1 and W_2, respectively. It is important to recognise that, as a consequence of incompleteness, there are infinitely many arbitrage-free option prices, each corresponding to a particular choice of the volatility risk premium ν. Several solutions have been proposed to solve the indeterminacy problem of the volatility risk premium. It is in fact an intertemporal equilibrium problem, as shown in Pham and Touzi (1993) [18].

The nice feature about model (5) is that it can reproduce fairly well the empirical regularities of implied volatilities observed on options with varying strike prices and the same maturity. Figure 1 shows how changes in the correlation parameter ρ bend the volatility smile of a European call.[5] An increase in ρ raises the relative price of *out-of-the-money* calls and so tilts the smile to the left. As argued by Heston (1993) [10], the interpretation is fairly intuitive. When high returns are positively correlated with volatility (high ρ), the distribution of spot returns is spread out to the right and spread in to the left. The induced skewness raises the price, in volatility terms, of options which benefit more from a fat right tail.

Figure 1

Volatility smile of a European call

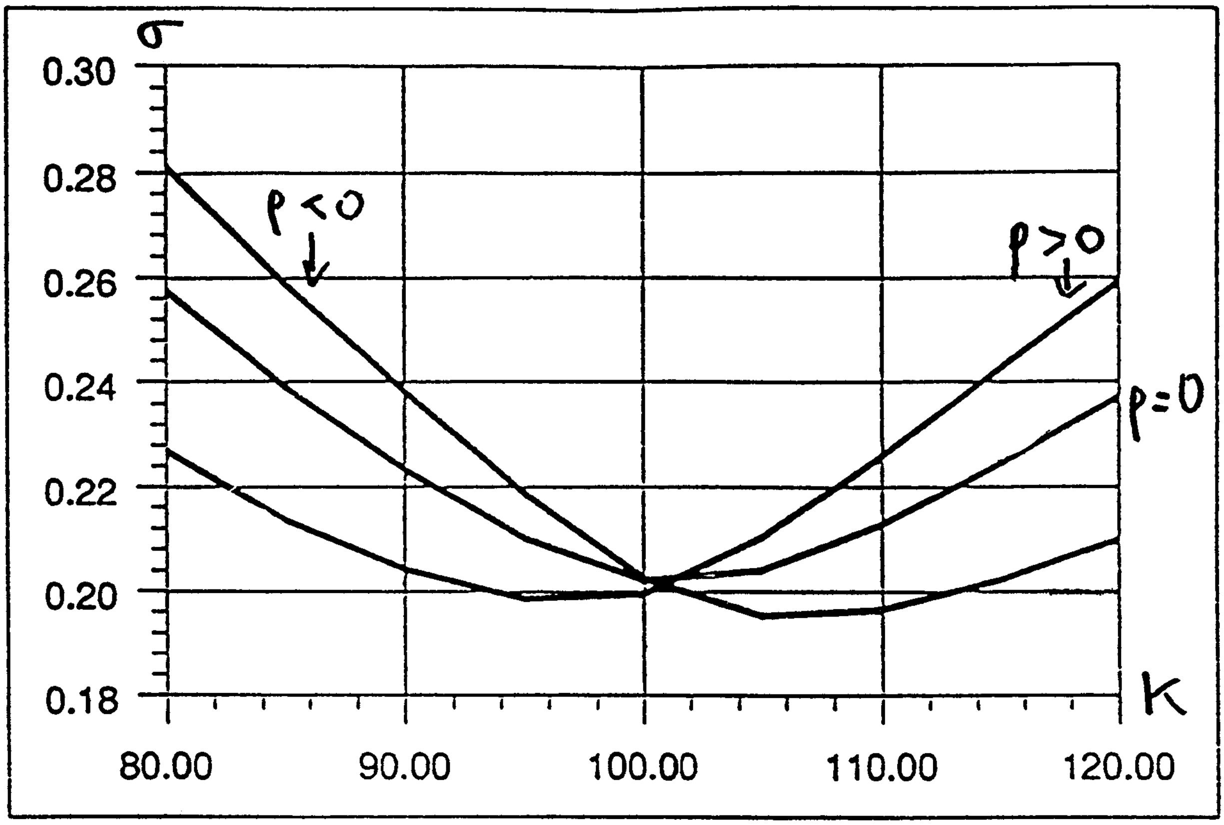

Since a right-skewed distribution raises the call and lowers the put components of a risk reversal, it has unequivocal effects on its price. Indeed, skewness of the exchange rate distribution is precisely what makes risk reversals valuable. This reasoning is confirmed even by a casual comparison between spot versus volatility charts, and the corresponding risk reversals prices; cf. Figures 2-9. The evidence is particularly striking for the dollar/Mark and dollar/yen markets. In both cases, spot returns were first negatively correlated with volatility, and particularly so in February-March, when the dollar plunged in value and its volatility rocketed. At the same time, risk reversals which were already negative experienced a sharp drop. Starting in June or July, spot rates became progressively more correlated with volatility; in the meantime, risk reversals increased and eventually turned positive, before peaking in August. In September, the correlation was again reversed, and this was followed by a drop in risk reversals prices. The same observations apply to the Mark/franc rate, although the positive correlation was never reversed, producing no change of sign in risk reversals.

5 Simulations presented by Nizar Touzi at a Seminar of the École Normale Supérieure, entitled "Méthodes Non Linéaires en Finance", June 1995.

Hence, it seems reasonable to interpret at-the-money volatility as the market's best guess about instantaneous volatility, and risk reversals as a market's measure of the comovements between spot rates and their volatility. The first identification results from the Black-Scholes and stochastic volatility models producing comparable prices for at-the-money options. The second is based on an informal calculation given in the appendix according to which, up to a homographic transformation, risk reversals indicate in "risk-neutral" terms the difference in likeliness between large upward and large downward deviations of the exchange rate from the forward rate. More precisely, they depend on $P(S_T \geq K_c) - P(S_T \leq K_p)$, where $K_p < K_c$ are the strike prices of the put and call components of the option, respectively, and P is the risk-neutral probability under which the future spot rate is centred on the forward rate. Since, under all stochastic volatility models, the instantaneous correlation between spot rates and their volatility control for the skewness of the exchange rate distribution, the price of risk reversals can be regarded as a time-varying, market-oriented and forward-looking measure of the spot/volatility correlation. In the following section, an elementary balance portfolio model, which captures the gist of stochastic volatility, is presented with a view to examining the impact of changing volatility and skewness on the exchange rate risk premium.

2.2 A balance portfolio model with volatility

Consider the following two-period specialisation of model (5):

$$d + i^* = \mu + \sigma(v)\left(\sqrt{1-\rho^2}u + \rho v\right),$$

$$\text{with } \sigma(v) = \sigma(1+\gamma v), \tag{6}$$

where d is the depreciation (appreciation) rate of the home (foreign) currency. As before, shocks to the exchange rate have an autonomous component, called u, and a volatility component, called v, with both zero mean and unit variance. To simplify the model, it is further assumed that v takes on the values +1 and -1 only with probability ½. The variables μ, ρ and γ receive similar interpretations as before. As expected, the random component of depreciation is uncentred, with mean $\gamma\sigma\rho$ and variance $\sigma^2(1+\gamma^2(1-\rho^2))$. The correlation parameter ρ controls the skewness of the distribution. In order to get an interesting theory of the risk premium, it is necessary to let the nominal interest rates on domestic and foreign bonds depend on volatility. Again, a linear schedule is assumed, with $i(v) = i+\varepsilon v$ and $i^*(v) = i^*+\eta v$, so that ε and η can be interpreted as the sensitivity of the home and foreign nominal interest rates to volatility.

The risk premium on the *foreign* currency is by definition $RP = E\,d-(i(v) - i^*(v)) = \mu - i + \gamma\sigma\rho$, so that (6) can be written as:

$$d = i(v) - i^*(v) + RP + \tilde{\sigma}(v) \tag{7}$$

$$\text{with } \tilde{\sigma}(v) = \sigma(v)\left(\sqrt{1-\rho^2}u + \rho v\right) - \gamma\sigma\rho.$$

In the textbook portfolio balance model, it is assumed that risk-averse investors minimise consumption risk, for a given share λ of consumption in the home currency. Inflation of the composite consumption good is given by $\tilde{\pi} = \lambda\pi + (1-\lambda)(\pi^* + d) = \bar{\pi} + (1-\lambda)d$, where $\bar{\pi} = \lambda\pi + (1-\lambda)\pi^*$ is weighted inflation. The real returns on the home and foreign currencies in terms of the composite consumption good are then found to be, respectively:

$$\tilde{r} = i(v) - \bar{\pi} - (1-\lambda)d, \tag{8}$$

$$\tilde{r}^* = i^*(v) - \bar{\pi} + \lambda d. \tag{9}$$

Figure 2

Risk reversals (in implied volatilities) - DM/French franc

(one-month)

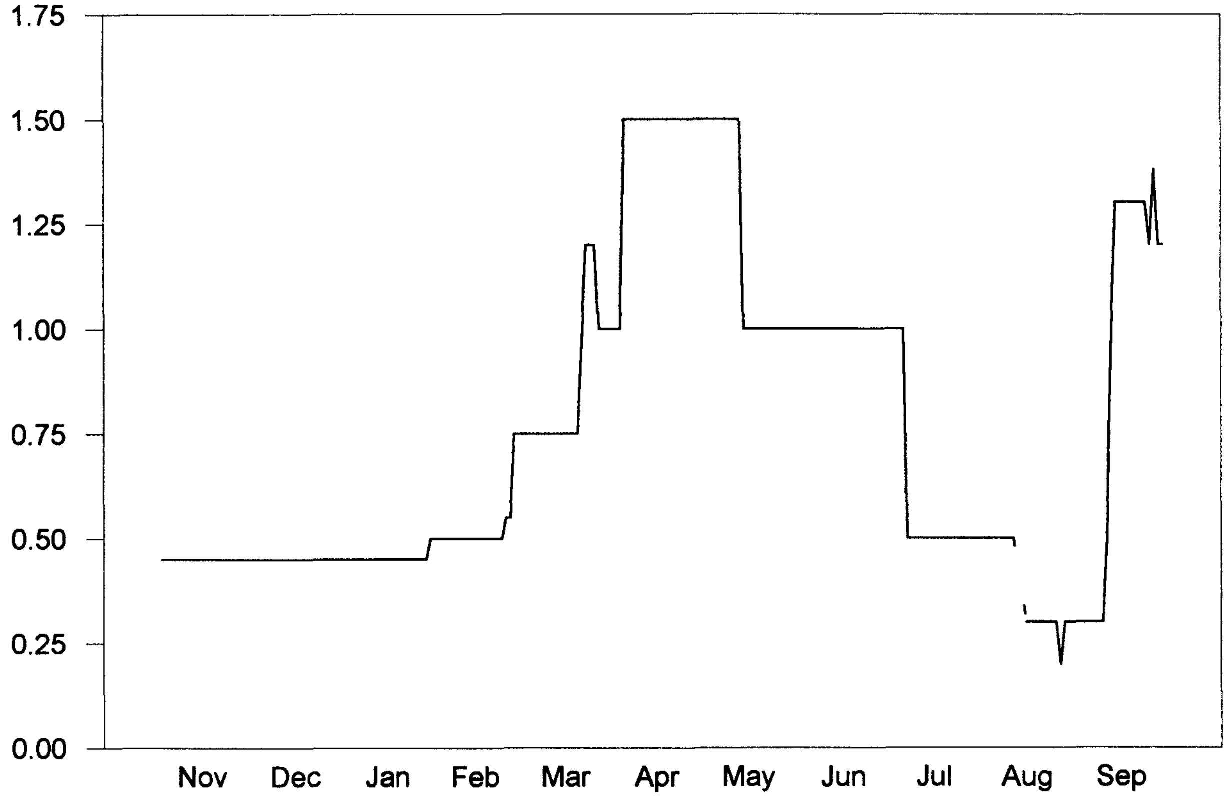

Figure 3

Spot rate and at-the-money volatility - DM/French franc

(one-month)

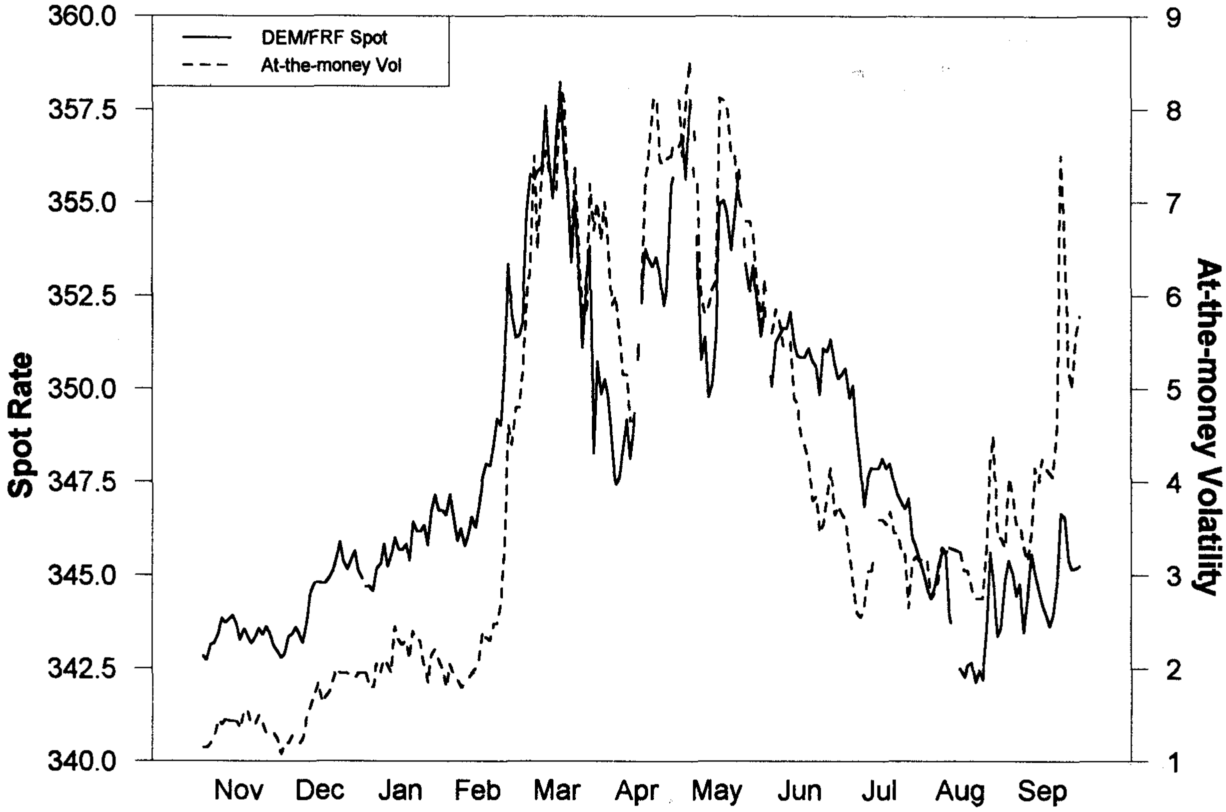

Figure 4

Risk reversals (in implied volatilities) - DM/French franc
(three-month)

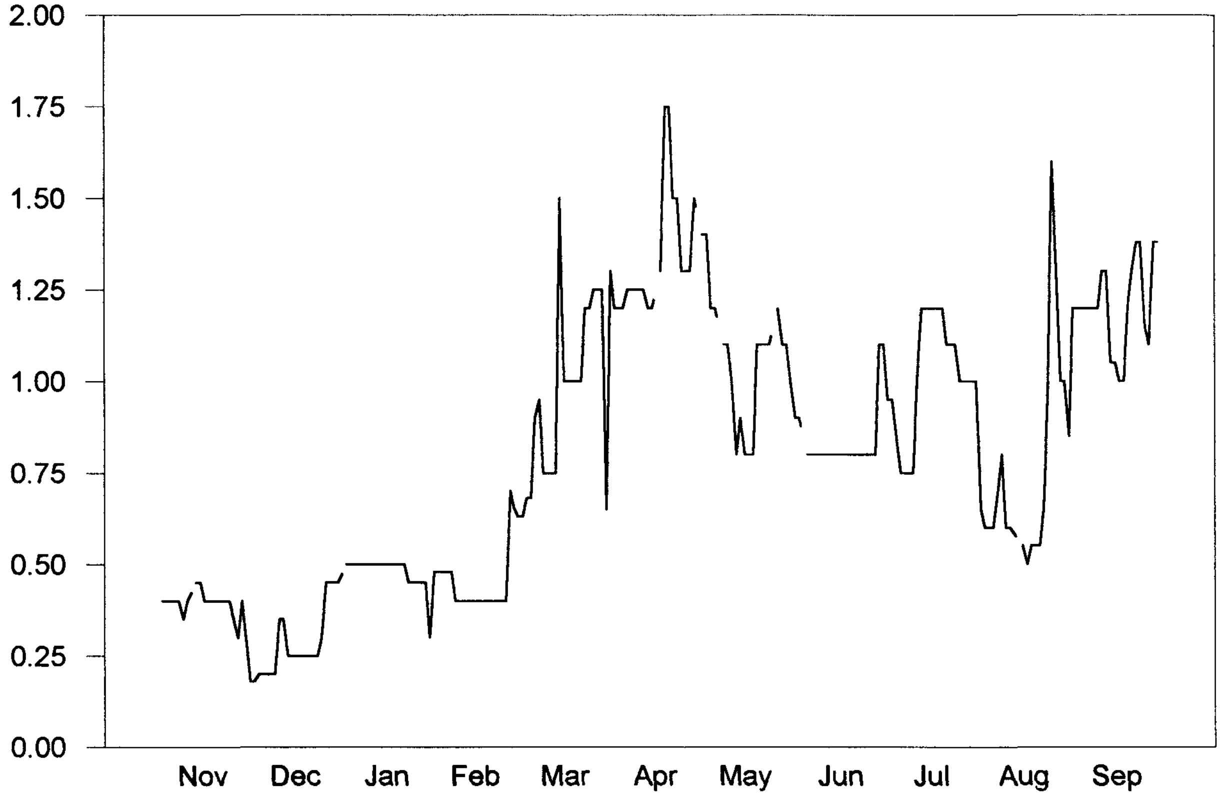

Figure 5

Spot rate and at-the-money volatility - DM/French franc
(three-month)

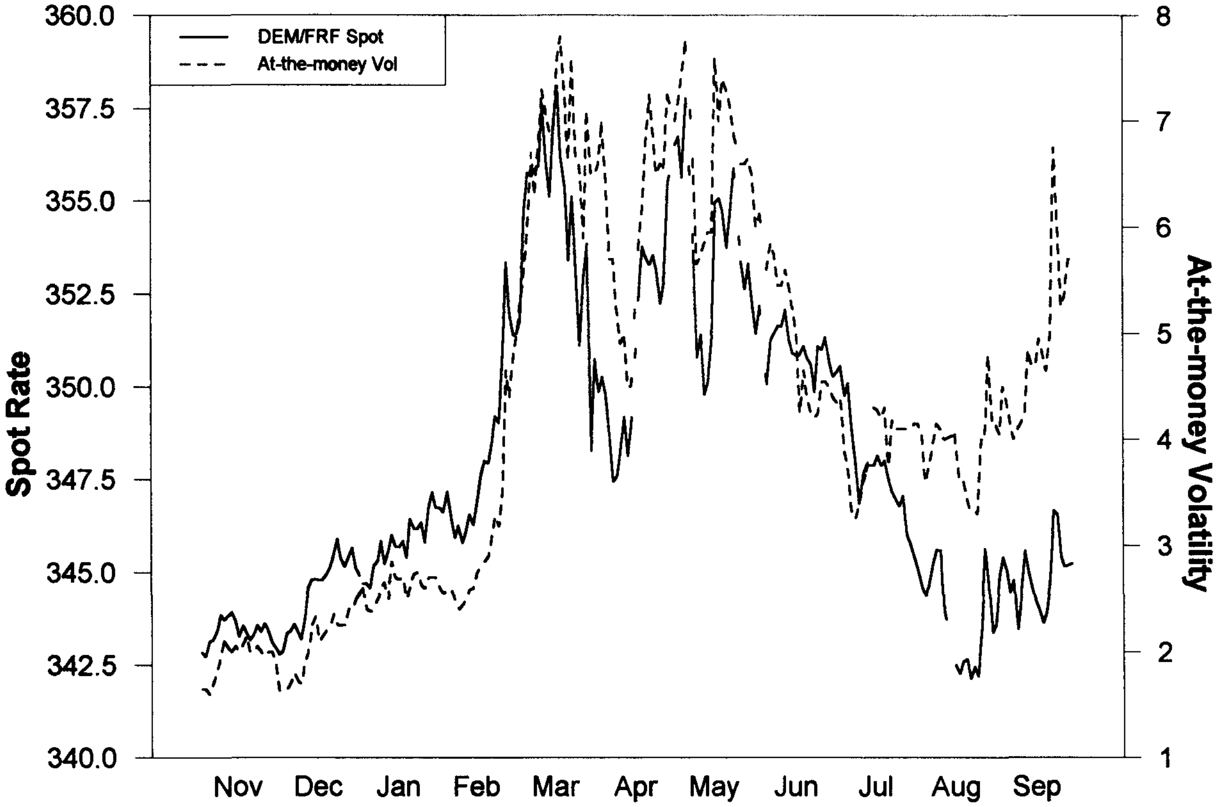

Figure 6

Risk reversals (in implied volatilities) - US dollar/DM
(one-month)

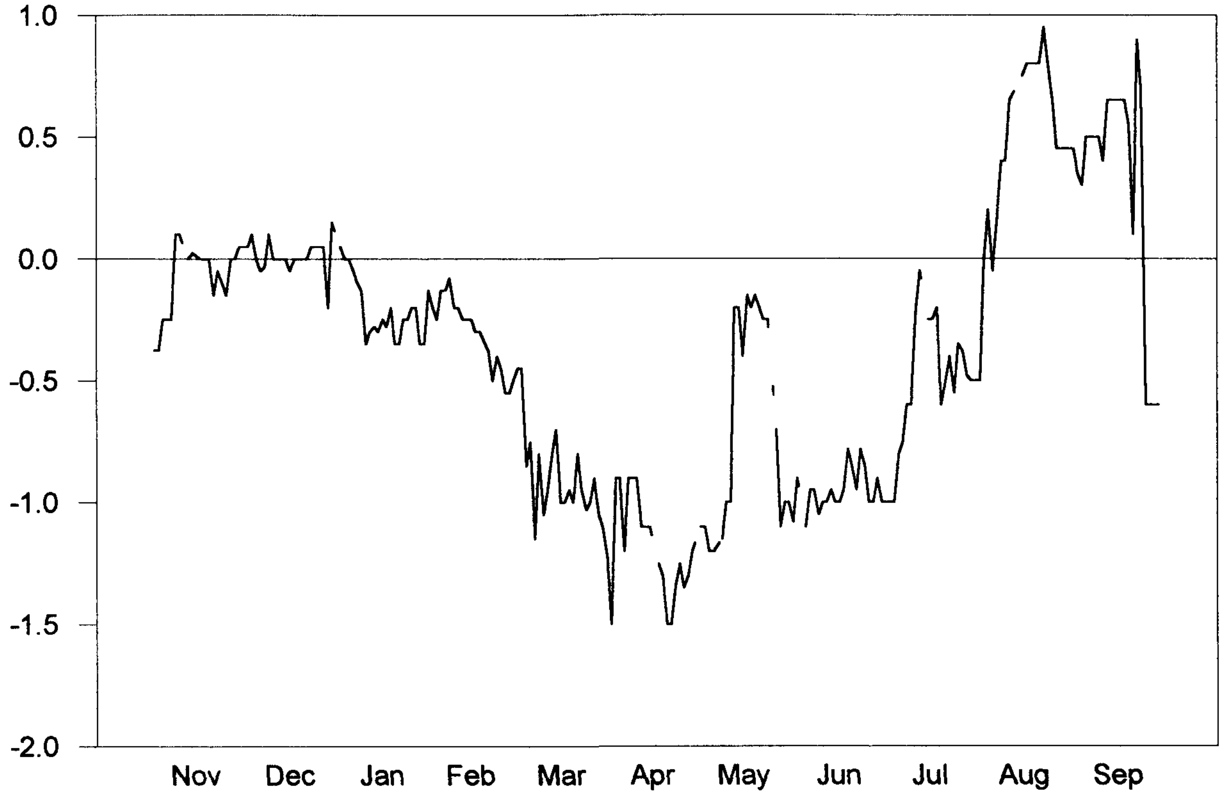

Figure 7

Spot rate and at-the-money volatility - US dollar/DM
(one-month)

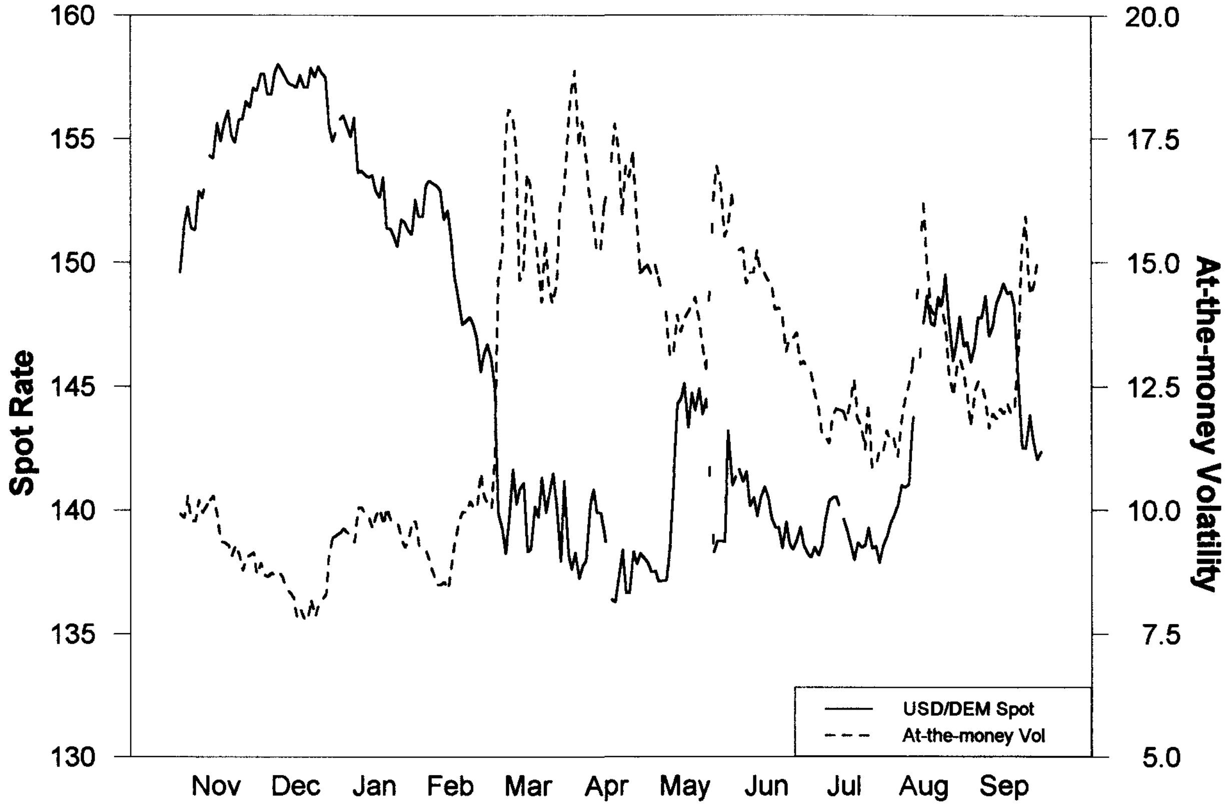

Figure 8

Risk reversals (in implied volatilities) - US dollar/DM

(three-month)

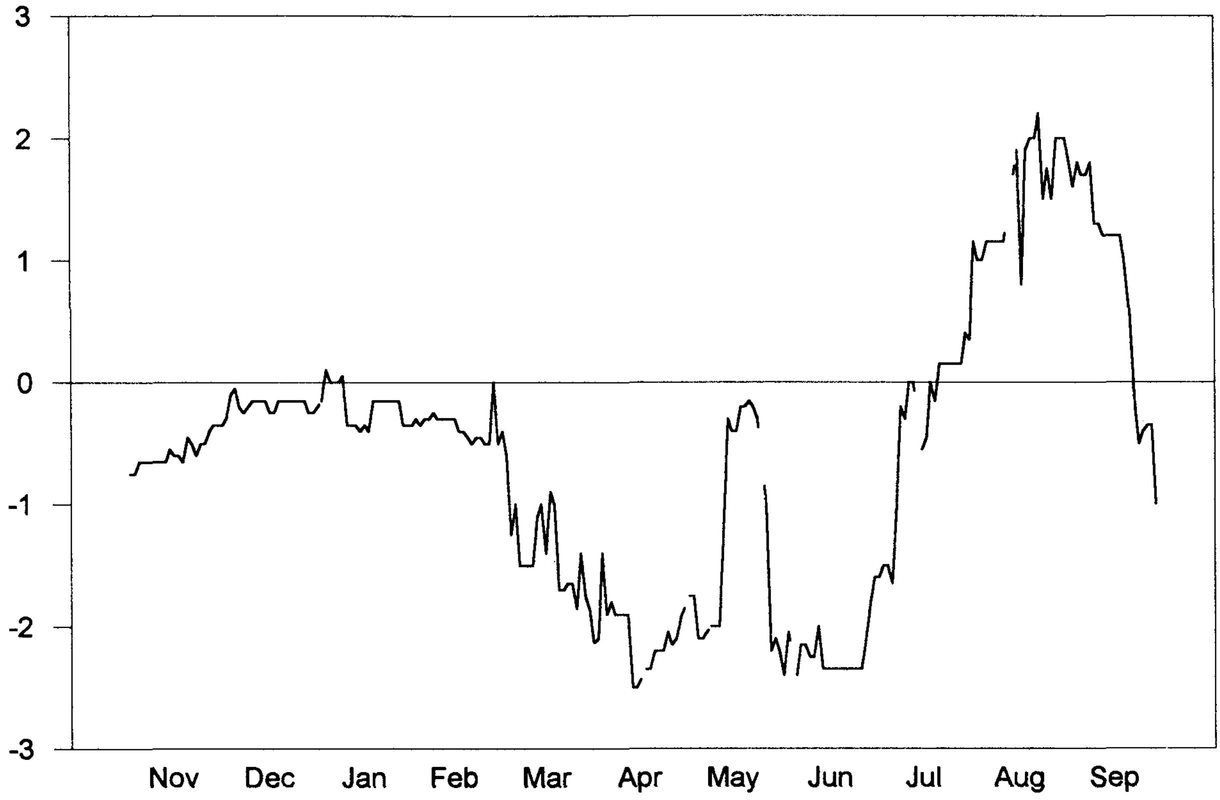

Figure 9

Spot rate and at-the-money volatility - US dollar/DM

(three-month)

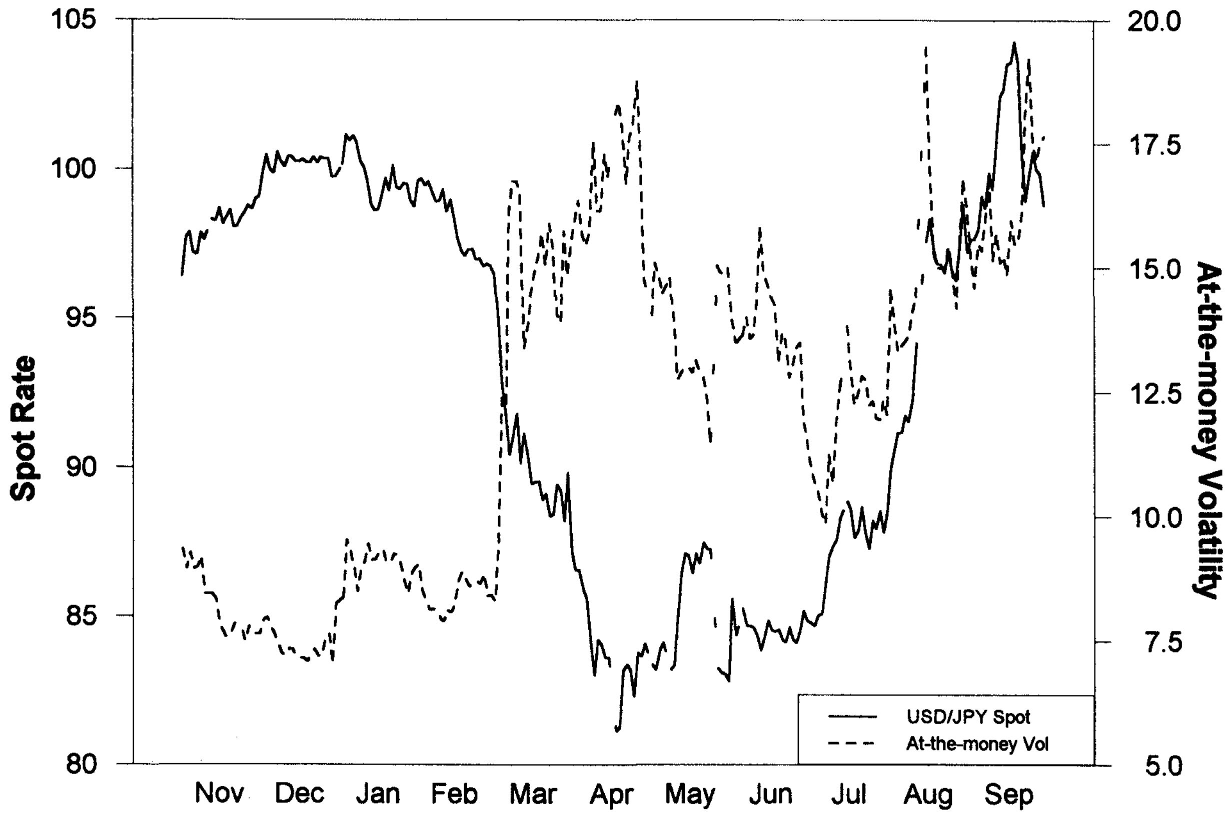

The balance portfolio theory states that the minimum variance portfolio shares are given by the correlation between the real return differential $\tilde{r}^* - \tilde{r}$ and the real returns denominated in the different currencies. Substituting (7) into (8-9), one finds after some elementary calculations that the home currency share λ' in the minimum variance portfolio is given by:

$$\frac{\operatorname{cov}(\tilde{r}^* - \tilde{r}, \tilde{r}^*)}{\operatorname{var}(\tilde{r}^* - \tilde{r})} = \lambda + \rho \frac{\lambda\varepsilon + (1-\lambda)\eta}{\sigma(1+\gamma^2(1-\rho^2))} \equiv \lambda'.$$

Finally, a standard mean-variance argument implies that, for a risk-aversion parameter θ and a supply share of the home currency of X, the *home* currency risk premium is given by $-RP = \theta \operatorname{var}(\tilde{r}^* - \tilde{r})(X - \lambda')$. Equating the home currency risk premium with the foreign currency's forward bias, one finds:

$$fb_t = b\sigma^2 + c(\sigma\rho), \text{ with} \tag{10}$$

$$b \;= \theta(1+\gamma^2(1-\rho^2))(X-\lambda) \tag{11}$$

$$c \;= -\,\theta(\lambda\varepsilon + (1-\lambda)\eta). \tag{12}$$

Thus, the forward bias has two components. The first is the standard Kouri-Dornbusch risk premium. It is related to the relative asset supply share of the home currency with respect to the home consumption share λ. An excess supply of the home (resp. foreign) currency raises (resp. lowers) the forward bias. The second is a term which depends on the correlation between the spot rate and its volatility, as was hoped for. The coefficient c depends both on the risk-aversion parameter θ and on the sensitivity of the consumption-weighted interest rate to volatility. Hence the magnitude and direction of the effect essentially depends on the intensity and sign of the weighted interest rates/volatility correlation.

The sample correlations of nominal interest rates with volatility are reported in Table 6. Because consumption shares are not known, consumption-weighted interest rates were simply taken to be the mean interest rates for the twin currencies. In the period under review, the correlation is negative for the dollar/Mark and the dollar/yen rates, and positive for the Mark/franc rate. Thus, the portfolio balance approach is broadly consistent with the empirical results of Section 1 for the US dollar vis-à-vis the Mark and the yen, but not for the Deutsche Mark vis-à-vis the franc.

Table 6

Sample correlation between volatility and one-month interest rates

Currency	i	i^*	$\frac{i+i^*}{2}$
DM/French franc	0.90	- 0.47	0.78
US$/DM	- 0.60	0.28	- 0.35
US$/yen	- 0.68	0.14	- 0.60

Conclusion and limitations

The empirical results reported in this paper purport that, over the period under review, forward biases were positively correlated with risk reversals. A tentative explanation runs as follows.

First, risk reversals may capture directional biases in the exchange rate for a large class of models, as shown in the appendix, but only with respect to the risk-neutral probability and conditionally on large deviations. This is perhaps of limited interest, because economists are primarily interested in deviations of the exchange rate from its expected value under the *true* probability distribution, i.e. in foreign exchange risk premia. Risk reversals do, however, capture the *skewness* of the distribution. A plausible interpretation, much investigated in the modern financial literature, is that volatility itself is stochastic, and that its time-varying correlation with the exchange rate induces skewness in the distribution. But this is precisely what makes risk reversals valuable. Hence, the price changes of risk reversals reflect the time-varying correlation between exchange rate risk and volatility risk. Moreover, this correlation is invariant with respect to a change in probability, so that it is the same whether computed under the true or the risk-neutral probability.

Second, some of the forward bias may be ascribed to a risk premium required by international investors. Assume for instance that the exchange rate and its volatility move counter-cyclically. This was seen to be the case in the first part of both the dollar/Mark and the dollar/yen samples. Because average interest rates appeared to be negatively correlated with volatility (at least according to sample correlations), a shock to volatility had the effect of lowering both the exchange rate and the average interest rate. In this case, the foreign currency was a reverse hedge because it depreciated when investors' aggregate consumption was low (and, conversely, appreciated when it was high). Consequently, the dollar/Mark and dollar/yen forward biases had to be relatively low. When changes in the value of the dollar and volatility became procyclical, as in the second part of the sample, the volatility-induced changes in the dollar and average interest rates started to move in opposite directions. This made the dollar a shelter currency, and its forward rate had to rise.

The theoretical result seems fairly intuitive: when world interest rates are negatively correlated with volatility, international investors will look at currencies which are more correlated with volatility as a way to hedge against volatility risk. Consequently, the forward biases of those currencies will rise. But this in turn may explain the positive relation found between the forward bias and risk reversals, because the price of risk reversals is all the higher, the more correlated the currency is with volatility.

Whether or not the above interpretation appears palatable, all the results so far should be considered as very preliminary and incomplete. One first task should be to ascertain whether there is an empirically robust relation between risk reversals and some derived measures of the correlation between exchange risk and volatility risk. Some simulations under a standard version of the stochastic volatility model would also help clarify the issue. Moreover, the econometric results presented here leave much room for improvement. The regressions are still unstable, indicating the possibility of specification errors, and a deeper investigation should be carried out in order to obtain more stable relations. Errors in variables are likely, and this calls for the use of instrumental variables. Also, one may note that, according to the simplest portfolio balance model, the variance coefficient may itself depend on the spot/volatility correlation parameter. Finally, the model's predictions remain at odds with the empirical evidence in the case of the Mark/franc market, and this may be due to the special exchange rate arrangements which prevail under the European Monetary Union. All these queries are part of the author's work in progress.

APPENDIX

Option prices on over-the-counter markets are quoted in volatilities. These volatilities are derived from the Garman-Kohlhagen (1983) formula, which is equivalent to a version of the Black-Scholes formula for options on a stock paying a continuous stream of dividends, with a rate given by the foreign interest rate. Let S_t be the spot exchange rate, which gives units of the leading currency per unit of the base currency. According to standard arbitrage arguments, its dynamic can be specified directly in terms of the risk-neutral probability. It is assumed to follow a lognormal process:

$$\frac{dS_t}{S_t} = (i - i^*)dt + \sigma dw_t,$$

where i-i^* is the interest rate differential between the home and the foreign countries, σ is the standard deviation of the instantaneous rate of change, and w is a standard Brownian motion. In probabilistic terms, the Garman-Kohlhagen value of a European call option on foreign exchange can be expressed as:

$$C(S,K,\tau,i,i^*,\sigma) = e^{-i^*\tau} S\, N_1(\log S_T \geq \log K) - e^{-i\tau} K\, N_2(\log S_T \geq \log K), \tag{13}$$

where $F = e^{(i-i^*)\tau}S$ is the forward rate, K the strike price of the option, τ the time to expiry, and N_1 and N_2 the cumulative normal distributions with mean $\log F+(\sigma^2/2)\tau$ and $\log F - (\sigma^2/2)\tau$, respectively, and standard deviation $\sigma\sqrt{\tau}$ (N_2 is the risk-neutral distribution). Similarly the value of a European put is given by:

$$P(S,K,\tau,i,i^*,\sigma) = e^{-i\tau} K\, N_2(\log S_T \leq \log K) - e^{-i^*\tau} S\, N_2(\log S_T \leq \log K). \tag{14}$$

Market participants use the formulas above as a convenient way to express options prices in terms of implied volatilities. By convention, the price of a European option, in vols, is the value of σ which makes the Garman-Kohlhagen value equal to its market value. At-the-money volatilities, in particular, are obtained for options whose strike prices are closest to the current forward rate at the time they are quoted.

In this simple setting an explicit formula is obtained by setting:

$$d = \frac{\log(F/K) + (\sigma^2/2)\tau}{\sigma\sqrt{\tau}},$$

so that $N_1(\log S_T \geq \log K) = \Phi(d)$ and $N_2(\log S_T \geq \log K) = \Phi(d-\sigma\sqrt{\tau})$, where Φ is the standard cumulative normal distribution. Recall that the *delta* of an option is defined as its sensitivity with respect to the current exchange rate, $\delta = \partial C/\partial S$ (or $-\partial P/\partial S$). For call and put options, moneyness and delta are positively related, i.e. the more out-the-money the options, the lower their deltas. The percentage difference between the exercise price K and the forward rate F can be found by equating $e^{-i^*\tau}\,\Phi(d)$ for a call (or $e^{-i^*\tau}\,\Phi(-d)$ for a put) with δ.

Risk reversals consist of a joint long (resp. short) position in an out-the-money call and short (resp. long) position in an out-the-money put on the same currencies, having the same δ and expiring at the same date. Hence, risk reversals are designed to be locally insensitive to the current exchange rate at the time they are issued. Usually, δ is chosen to be 0.25, as in the sample considered in this paper, but some 0.10-δ risk reversals are also traded. Let $-\alpha$ be the $e^{i^*\tau}\,\delta$ quantile of the cumulative normal distribution. (For $i^* = 0$ and $\delta = 0.25$, the value of α is approximately 0.67.) The strike prices $K_p < F < K_c$ of the put and call components are respectively:

$$K_p = Fe^{(\sigma_p^2/2)\tau - \alpha\sigma_p\sqrt{\tau}}, \tag{15}$$

$$K_c = Fe^{(\sigma_c^2/2)\tau - \alpha\sigma_c\sqrt{\tau}}. \tag{16}$$

It is clear that the price of risk reversals $\sigma_c - \sigma_p$, expressed in volatilities, is not sufficient to determine the actual pay-off of the option, since the payment $C(K_c(\sigma_c),\sigma_c) - P(K_p(\sigma_p),\sigma_p)$ depends on σ_c and σ_p separately, and not on their difference. It turns out, however, that $\sigma_c - \sigma_p$ is a good approximation of the value of the option,[6] so that in practice traders agree on this difference first before setting each component separately. The precise assignment varies over time and over the currencies traded. (A common convention is to set the put volatility to at-the-money volatility σ_{atm}, thus making the call volatility equal to $\sigma_{atm} + (\sigma_c - \sigma_p)$.)

If the real world behaved as the Black-Scholes model would imply, the price of risk reversals (in vols) would be identically zero. Because $\sigma_c - \sigma_p$ is the only quantity recorded in the data set, it is important to relate it to parameters of more general models in which non-zero risk reversals can be accounted for. In most cases, it is possible to express the price of a European option as the difference between the present value of the spot asset conditional upon optimal exercise and that of the strike price. Hence the call value of the risk reversal can be written as:

$$C = Se^{-i^*\tau} P_1(S_T \geq K_c) - K_c e^{-i\tau} P_2(S_T \geq K_c),$$

where P_1 and P_2 are two probabilities which depend on the model chosen. (The latter is the risk-neutral probability; it is remarkable that $P_i(S_T \geq K_c)$, $(i = 1,2)$, is independent of the particular level of the forward rate at which the option price is computed, at least for standard stochastic volatility models.) This in turn implies, given (16):

$$\frac{C}{Fe^{-i\tau}} = P_1(S_T \geq K_c) - e^{(\sigma_c^2/2)\tau + \alpha\sigma_c\sqrt{\tau}} P_2(S_T \geq K_c).$$

Equating this with (13) yields:

$$N_1(S_T \geq K_c) - P_1(S_T \geq K_c) = e^{e^{(\sigma_c^2/2)\tau+\alpha\sigma_c\sqrt{\tau}}} (N_2(S_T \geq K_c) - P_2(S_T \geq K_c)).$$

A similar calculation for the put gives:

$$N_1(S_T \leq K_p) - P_1(S_T \leq K_p) = e^{e^{(\sigma_p^2/2)\tau-\alpha\sigma_p\sqrt{\tau}}} (N_2(S_T \leq K_p) - P_2(S_T \leq K_p)).$$

Subtracting the latter equation from the former, one gets:

$$\begin{aligned}\Delta\breve{N} - \alpha\sigma_c\sqrt{\tau}N_2(S_T \geq K_c) - \alpha\sigma_p\sqrt{\tau}N_2(S_T \leq K_p) = \\ \Delta\breve{P} - \alpha\sigma_c\sqrt{\tau}P_2(S_T \geq K_c) - \alpha\sigma_p\sqrt{\tau}P_2(S_T \leq K_p),\end{aligned}$$

where $\Delta X = X_1 - X_2$ and $\breve{X} = X(S_T \geq K_c) - X(S_T \leq K_p)$, $X = N$ or P, and second order terms in σ are neglected. After some rearrangements this can be written as:

$$\alpha(\sigma_c - \sigma_p)\sqrt{\tau}(P_2(S_T \leq K_p) - N_2(S_T \leq K_c)) = \Delta\breve{N} - \Delta\breve{P} + \alpha\sigma_c\sqrt{\tau}(\breve{P}_2 - \breve{N}_2).$$

Since under the Black-Scholes assumptions the spot returns distribution is symmetric, the quantity $\Delta\breve{N}$ is small. In fact $\breve{N}_1 = 0$ and $\breve{N}_2 = -(\sigma_c + \sigma_p)\sqrt{\tau}\phi(\alpha)$. On the other hand $\breve{P}_1$ and $\breve{P}_2$ reflect by construction the skewness of spot returns. However, the *difference* $\Delta\breve{P}$ *is* likely to be much less sensitive to variations in the skewness. Collecting terms that are approximately constant, one obtains:

6 One can check that up to terms in $\sqrt{\tau}$:

$$\frac{C-P}{Fe^{-i\tau}} = (\sigma_c - \sigma_p)\sqrt{\tau}(\phi(-\alpha) - \alpha\Phi(-\alpha)) + \sigma_c(\sigma_c + \sigma_p)\sqrt{\tau}\alpha\phi(\alpha).$$

For $\delta = 0.25$, $\phi - \alpha\Phi \approx 0.15$ and $\alpha\phi \approx 0.21$.

$$\frac{\sigma_c - \sigma_p}{\sigma_c} = \frac{P_2(S_T \geq K_c) - P_2(S_T \leq K_p) + k}{P_2(S_T \leq K_p) - N_2(S_T \leq K_p)}. \tag{17}$$

Risk reversals thus provide a direct measure of the skewness of spot returns. The denominator is positive because the kurtosis of the true risk-neutral distribution P_2 is larger than that of the normal distribution. All quantities in (17) can be computed under any particular version of the stochastic volatility model: closed-form solutions as in Heston [10] or Leblanc [13] can be readily evaluated by using numerical simulations.

References

[1] Bates, David S.: "The crash premium: option pricing under asymmetric processes, with applications to options on Deutschemark futures", Wharton School, University of Pennsylvania.

[2] Bates, David S. (1993): "Jumps and stochastic volatility: exchange rate processes implicit in PHLX Deutsche options", NBER Working Paper No. 4596.

[3] Black, Fischer and Myron Scholes (1973): "The pricing of options and corporate liabilities", *Journal of Political Economy*, 81: pp. 637-659.

[4] Bodurtha, James and Georges Courtadon (1987): "Tests of an American option pricing model on the foreign currency options market", *Journal of Financial and Quantitative Analysis*, 22: pp. 153-167.

[5] Borensztein, Eduardo and Michael Dooley (1987): "Options on foreign exchange and exchange rate expectations", *IMF staff papers*, pp. 643-680.

[6] Dornbusch, Rudiger (1983): "Exchange rate risk and the macroeconomics of exchange rate determination", in R. Hawkins *et al.*, eds., *Research in International Business and Finance*, 3, Greenwich, Conn: JAI Press.

[7] Frankel, Jeffrey (1982): "In search of the exchange risk premium: a six currency test assuming mean-variance optimization", *Journal of International Money and Finance*, pp. 255-274.

[8] Garman, Mark and Steven Kohlhagen (1983): "Foreign currency option values", *Journal of International Money and Finance*, pp. 231-238.

[9] Hansen, Lars P. and Robert J. Hodrick (1980): "Forward exchange rates as optimal predictors of future spot rates: an econometric analysis", *Journal of Political Economy*, 88: pp. 829-853.

[10] Heston, Steven L. (1993): "A closed-form solution for options with stochastic volatility with applications to bond and currency options", *The Review of Financial Studies*, 6: pp. 327-343.

[11] Hull, John C. and A. White (1987): "The pricing of options on assets with stochastic volatilities", *Journal of Finance*, 42: pp. 281-300.

[12] Kouri, P. (1977): "International investment and interest rate linkages under flexible exchange rates", in R. Aliber, ed., *The Political Economy of Monetary Reform*, London: Macmillan Press.

[13] Leblanc, Boris (1994): "Processus de Bessel en finance: utilisation dans les modèles à volatilité stochastique", Conférence internationale en économie mathématique et finance mathématique, Tunis.

[14] Lyons, Richard K. (1988): "Tests of the foreign exchange risk premium using the expected second moments implied by option pricing", *Journal of International Money and Finance*, 7: pp. 91-108.

[15] Malz, Allan M. (1994): "Using option prices to estimate ex ante realignment probabilities in the European monetary system", Federal Reserve Bank of New York Working Paper.

[16] Melino, Angelo and Stuart Turnbull (1991): "The pricing of foreign currency options", *Canadian Journal of Economics*, 24: pp. 251-281.

[17] Nelson, Daniel B. and Dean P. Foster (1994): "Asymptotic filtering theory for univariate ARCH models", *Econometrica*, 62: pp. 1-41.

[18] Pham, Huyên and Nizar Touzi (1994): "Intertemporal equilibrium risk premia in a stochastic volatility model", Conférence internationale en économie mathématique et finance mathématique, Tunis.

[19] Scott, Louis O. (1987): "Option pricing when the variance changes randomly: theory, estimation and an application", *Journal of Financial and Quantitative Analysis*, 22: pp. 419-438.

[20] Wei, Shang-Jin and Jeffrey A. Frankel (1991): "Are option-implied forecasts of exchange rate volatility excessively variable?", NBER Working Paper No. 3910.

[21] West, Kenneth D. (1988): "Asymptotic normality, when regressors have a unit root", *Econometrica*, 56: pp. 1397-1418.

[22] Westerfield, Janice M. (1977): "An examination of foreign exchange risk under fixed and floating rate regimes", *Journal of International Economics*, 7: pp. 181-200.

[23] Wiggins, J.B. (1987): "Options values under stochastic volatilities", *Journal of Financial Economics*, 19: pp. 351-372.