

Implied Volatility Skews in the Foreign Exchange Market

Empirical Evidence from JPY and GBP: 1997-2002

The Leonard N. Stern School of Business
Glucksman Institute for Research in Securities Markets
Faculty Advisor: Stephen Figlewski
April 1, 2003

By Keith Gudhus*

*MBA 2003 candidate, Stern School of Business, New York University, 44 West 4th Street, New York, NY 10012, email: krig223@stern.nyu.edu, tel: (646) 246-8676. I would like to thank Stephen Figlewski and William Silber for their invaluable comments and suggestions, and Dan Silber and the Goldman Sachs foreign exchange desk for graciously providing the data which made this paper possible.

I. INTRODUCTION

The aim of this paper is to study the implied volatility skew (which we represent as the implied volatility of the 25 delta call minus the implied volatility of the 25 delta put) within the foreign exchange market. Specifically, we examine the skew for both JPY (quoted in Japanese yen per dollar) and GBP (quoted in dollars per British pound) across a variety of maturities, ranging from one week to one year.

For purposes of definition, a volatility smile refers to the variation of implied volatility with respect to strike price; a volatility skew exists when this smile is nonsymmetrical. Given that 3 month options are usually the most liquid and actively traded maturity, the main focus of our analysis is on the 3 month implied volatility skew for JPY and GBP.

We begin our analysis in Section II by surveying recent research into the implied volatility skew. We then describe the level and movement of the skew for JPY and GBP between November 14, 1997 and September 19, 2002 in Section III. At this point, we hypothesize that both the level of the underlying currency and the recent trend in the currency will be positively correlated with the skew, which we find support for in the next two sections. In Section IV we discover a positive correlation between the level of the underlying currency and its respective skew, and correct for autocorrelation problems inherent in the data. We also find a positive correlation between the recent trend of the underlying currency and its respective skew in Section V. In Section VI we combine the results from the prior two sections to complete our skew models. In the final section, we refer back to our initial hypotheses in order to further explain our results.

II. PREVIOUS WORK

Hull (2000) notes that the volatility smile in the foreign exchange market graphically corresponds to an upward-facing parabola, with out-of-the-money options possessing greater implied volatilities than at-the-money options. This smile corresponds to an implied probability distribution which exhibits more kurtosis (e.g. fatter tails) than a lognormal distribution. Hull notes that this smile is “consistent with empirical data showing that extreme movements in exchange rates happen more often than the lognormal distribution would predict.” Within options literature, these extreme moves are explained by two effects—nonconstant volatility and jumps in the price movement of the underlying currency.

Some of the most interesting literature regarding volatility skews relates to the equity options market, in which implied volatilities generally increase as the strike price decreases (Poon and Granger 2002). One explanation argues that the skew is caused by a *leverage effect*. Specifically, a decreasing stock price increases a firm’s leverage, which makes the firm’s equity riskier. Thus, implied volatility increases as the stock price decreases. A second explanation posits that the skew is caused by “*crash-o-phobia*.” (Rubinstein 1994). It argues that traders are constantly concerned about another stock market crash, and hence bid up the implied volatilities of out-of-the-money puts relative to out-of-the-money calls. Rubinstein’s theory, by relating observed skews to traders’ behaviors, offers an extremely important springboard for our investigation in Section III.

III. DATA

The data sample consists of approximately five years of daily data for JPY and GBP (from 11/14/97 to 9/19/02), which was obtained from the Goldman Sachs foreign exchange desk. For each currency, we have daily spot closes and daily implied volatility closes (for 1 week, 1 month, 3 month, 6 month, and 1 year maturities). The sample contains three separate implied volatilities for each maturity—the 25 delta put, the 50 delta option, and the 25 delta call—which are all expressed in annual terms and are for European options. The implied volatilities are those actually quoted by Goldman Sachs market makers. If we wanted to price the options, we would simply plug these implied volatilities into the Garman-Kohlhagen model, which is essentially the Black-Scholes formula with a foreign riskless interest rate as the payout on the underlying asset. This is the standard pricing convention in the foreign exchange market.

From this data set, the volatility skew is calculated for each maturity. For the purposes of this paper, we represent the skew by the following equation: **volatility skew = implied volatility of the 25Δ call - implied volatility of the 25Δ put**. Below we present a descriptive summary of our data in Tables 1 and 2:

Table 1: Summary Statistics for JPY Implied Volatility Skew

	<u>1 week</u>	<u>1 month</u>	<u>3 month</u>	<u>6 month</u>	<u>1 year</u>
N:	1219	1219	1219	1219	1219
Mean:	-1.03%	-0.92%	-0.54%	-0.35%	-0.23%
Standard Deviation:	1.25%	1.16%	0.94%	0.86%	0.81%
Minimum:	-5.00%	-4.01%	-2.78%	-2.19%	-1.80%
Maximum:	2.05%	2.25%	1.82%	1.30%	1.35%

Table 2: Summary Statistics for GBP Implied Volatility Skew

	<u>1 week</u>	<u>1 month</u>	<u>3 month</u>	<u>6 month</u>	<u>1 year</u>
N:	1219	1219	1219	1219	1219
Mean:	-0.09%	-0.10%	-0.11%	-0.12%	-0.13%
Standard Deviation:	0.46%	0.39%	0.30%	0.25%	0.22%
Minimum:	-1.50%	-1.50%	-1.07%	-0.83%	-0.67%
Maximum:	1.20%	1.10%	0.84%	0.59%	0.48%

It is important to note that in contrast to stocks, defining the skew in the foreign exchange market is arbitrary (e.g. a dollar call is a yen put, and vice versa). Given the quotation conventions of the foreign exchange market, the JPY skew is for dollar calls and dollar puts, and the GBP skew is for British pound calls and British pound puts. We plot below (in Figures 1 and 2) the 3 month skews over time to better illustrate the differing skew behaviors of JPY and GBP.

Figure 1: 3 month JPY Implied Volatility Skew





Figure 2: 3 month GBP Implied Volatility Skew



We can make a number of observations from the above graphs regarding the volatility skews for JPY and GBP. First, the JPY skew is negative the majority of the time. That is, the implied volatilities of the 25Δ puts are higher than those for the 25Δ calls. Furthermore, the magnitude of the skew is negatively biased, as the skew ranges in value from roughly -3 to +2.

In contrast to the negative bias of the JPY skew, the GBP skew is relatively symmetrical around zero. Furthermore, as opposed to the wide skew swings for JPY, the GBP skew rarely exceeds ± 1 .

In the next two sections, we aim to identify which variables explain the skew for JPY and GBP. As a starting point, we recall Section II, in which we referred to Mark Rubinstein's "crash-o-phobia" hypothesis, in which traders, fearful of stock market crashes, bid up the implied volatilities of out-of-the-money puts relative to out-of-the-money calls. We find this sort of behavioral analysis extremely insightful. Extending this concept a bit further, we expect that traders in the foreign exchange market price the

skew to reflect their assessment of future risks. In particular, two variables come to mind that could potentially explain the skew—the level of the underlying currency and the recent trend in that currency. Specifically, we hypothesize that traders expect the future risks of the underlying spot market to be in the direction of the recent currency trend and the recent currency level. Thus, we expect the skew to be positively correlated with both variables.

IV. SKEW AND SPOT CURRENCY LEVELS

In this section, we aim to test the first part of our hypothesis—that is, that the skew will be positively correlated with the underlying currency level. We begin our investigation by plotting the 3 month skew for JPY and GBP against the underlying spot level of the appropriate currency, which we display below in Figures 3 and 4:

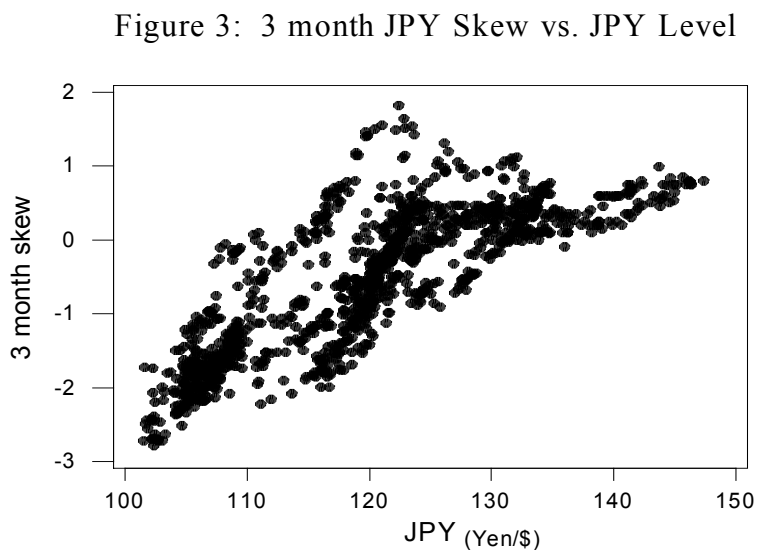
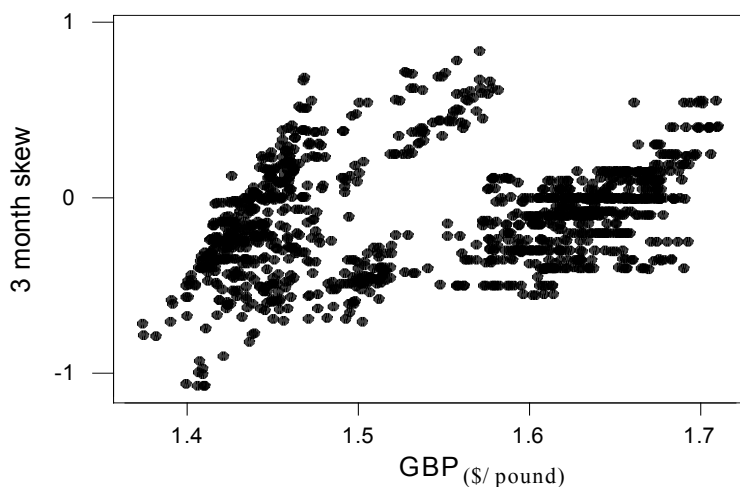


Figure 4: 3 month GBP Skew vs. GBP Level



As we can see, the JPY skew is positively correlated with the level of the dollar. All else being equal, at higher levels of the dollar (e.g. more yen per dollar), implied volatilities of dollar calls will increase relative to volatilities of dollar puts (for the same delta). The GBP skew also exhibits some positive correlation with the level of GBP (albeit less correlation than we saw with JPY). This is also evident by running regressions of the respective skews on their underlying spot currency levels. While the JPY skew regression yields an R-squared of 64.5%, the GBP skew regression yields an R-squared of merely 9.6%. The regression results are displayed on the following page in Tables 3 and 4:

Table 3: 3 month JPY Skew vs. JPY Level

The regression equation is: 3 month skew = - 9.33 + 0.0734 JPY

Predictor	Coef	SE Coef	T	P
Constant	-9.3349	0.1877	-49.74	0.000
JPY	0.073357	0.001560	47.03	0.000

S = 0.5626 R-Sq = 64.5% R-Sq(adj) = 64.5%

Durbin-Watson statistic = 0.05

Table 4: 3 month GBP Skew vs. GBP Level

The regression equation is: 3 month skew = - 1.61 + 0.966 GBP

Predictor	Coef	SE Coef	T	P
Constant	-1.6096	0.1344	-11.98	0.000
GBP	0.96617	0.08658	11.16	0.000

S = 0.2858 R-Sq = 9.3% R-Sq(adj) = 9.2%

Durbin-Watson statistic = 0.06

However, both the extremely low Durbin-Watson statistics (shown above in Tables 3 and 4) and the Residuals Versus the Order of the Data plots (shown on the following page in Figures 5 and 6), indicate the presence of autocorrelation.

Figure 5: Residuals Versus the Order of the Data (JPY)

(response is 3 month)





Figure 6: Residuals Versus the Order of the Data (GBP)

(response is 3 month)



In order to address the autocorrelation, we use the *Cochrane-Orcutt procedure* below:

1. We determine an estimate of ρ from the lag 1 entry in the ACF plot of the standardized residuals from our initial regressions. This value is 0.98 for JPY and 0.97 for GBP.

2. We create transformed variables $y_i^* = y_i - py_{i-1}$ and $x_i^* = x_i - px_{i-1}$
3. We perform a new regression of y_i^* on the x_i^* 's

We present the results for our new regressions in Tables 5 and 6 below:

Table 5: 3 month JPY Skew* vs. JPY Level*

The regression equation is: 3 month skew* = - 0.267 + 0.106 JPY*

Predictor	Coef	SE Coef	T	P
Constant	-0.266705	0.008378	-31.83	0.000
JPY	0.106436	0.003217	33.09	0.000

S = 0.1151 R-Sq = 47.4% R-Sq(adj) = 47.3%

Durbin-Watson statistic = 1.83

Table 6: 3 month GBP Skew* vs. GBP Level*

The regression equation is: 3 month skew* = - 0.174 + 3.69 GBP*

Predictor	Coef	SE Coef	T	P
Constant	-0.17448	0.01082	-16.13	0.000
GBP (p=0	3.6945	0.2300	16.06	0.000

S = 0.06405 R-Sq = 17.5% R-Sq(adj) = 17.4%

Durbin-Watson statistic = 2.19

Figure 7: Residuals Versus the Order of the Data (JPY)

(response is 3 month)





Figure 8: Residuals Versus the Order of the Data (GBP)

(response is 3 month)



As we can see, the Residuals Versus the Order of the Data plots (in Figures 7 and 8) and the much higher Durbin-Watson statistics (which are now above their critical values) indicate that our autocorrelation problems have been addressed. In addition, other residual plots indicate that the new regressions satisfy the standard normality and homoscedasticity assumptions.

However, even after correcting for autocorrelation, we still obtain extremely significant t-statistics (e.g. both p-values are 0) for the underlying currency level in both the JPY and GBP regressions. Thus, our conclusions remain the same—the skew is positively correlated with the underlying currency level for both JPY and GBP.

V. SKEW AND SPOT CURRENCY TRENDS

In this section, we aim to test the second part of our hypothesis—that is, that the skew will be positively correlated with the recent trend in the underlying currency. We continue our investigation by plotting the 3 month skew for JPY and GBP against the recent 100 day trend of the appropriate currency, which we display below in Figures 9 and 10:

Figure 9: 3 month JPY Skew vs. 100 day JPY Trend

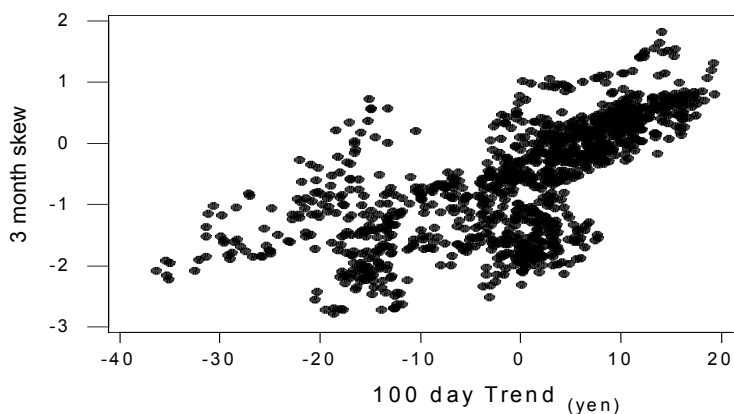
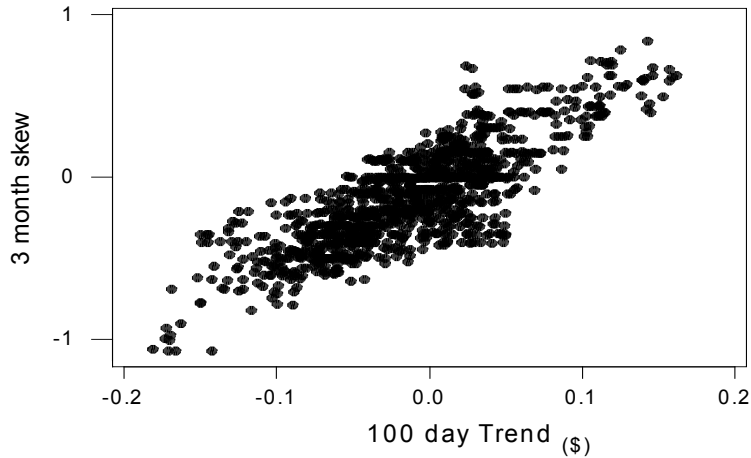


Figure 10: 3 month GBP Skew vs. 100 day GBP Trend



As we can see, both skews are positively correlated with the recent trend in their respective currencies (measured as the difference between the spot currency level today and that of 100 days ago). All else being equal, the more positive the recent trend in the underlying currency, implied volatilities of calls will increase relative to implied volatilities of puts (for the same delta). This is also evident by running regressions of the respective skews on the recent currency trends, which, after correcting for autocorrelation (using ρ estimates of 0.98 for JPY and 0.92 for GBP), yield extremely significant t-statistics for both trends. The regression results are displayed below in Tables 7 and 8:

Table 7: 3 month JPY Skew* vs. 100 day JPY Trend*

The regression equation is: $3 \text{ month skew}^* = -0.0116 + 0.0551 \text{ 100 day}^*$

Predictor	Coef	SE Coef	T	P
Constant	-0.011619	0.004122	-2.82	0.005
100 day	0.055069	0.002812	19.58	0.000

S = 0.1398 R-Sq = 25.0% R-Sq(adj) = 25.0%

Durbin-Watson statistic = 1.80

Table 8: 3 month GBP Skew* vs. 100 day GBP Trend*

The regression equation is: 3 month skew* = - 0.00689 + 2.56 100 day*

Predictor	Coef	SE Coef	T	P
Constant	-0.006891	0.001990	-3.46	0.001
100 day	2.5559	0.1800	14.20	0.000

S = 0.06726 R-Sq = 14.9% R-Sq(adj) = 14.9%

Durbin-Watson statistic = 2.05

VI. SKEW MODELS

From the above analysis, we see that the JPY skew is positively related to the level of JPY and the recent JPY trend, and the GBP skew is positively related to the level of GBP and the recent GBP trend. Now we look to synthesize these observations to create complete models for JPY and GBP skews. To fully encompass the trends in the underlying spot market, we decide to include both a short term trend (e.g. 20 days) and a long term trend (e.g. 100 days) as explanatory variables. In addition, we include the underlying level of the appropriate currency in each model. All models are corrected for autocorrelation, using ρ estimates of 0.96 for JPY and 0.90 for GBP. We display the regression results on the following page in Tables 9 and 10:

Table 9: 3 month JPY Skew* vs. JPY Level*, 20 day JPY Trend*, 100 day JPY

Trend*

The regression equation is:

$$3 \text{ month skew*} = - 0.480 + 0.0956 \text{ JPY*} + 0.00397 \text{ 20 day*} + 0.00627 \text{ 100 day*}$$

Predictor	Coef	SE Coef	T	P
Constant	-0.48034	0.02479	-19.38	0.000
JPY*	0.095619	0.005135	18.62	0.000
20 day*	0.003968	0.003245	1.22	0.222
100 day*	0.006272	0.003191	1.97	0.050

S = 0.1173 R-Sq = 48.3% R-Sq(adj) = 48.2%

Durbin-Watson statistic = 1.77

Table 10: 3 month GBP Skew* vs. GBP Level*, 20 day GBP Trend*, 100 day GBP

Trend*

The regression equation is:

$$3 \text{ month skew*} = - 0.163 + 0.998 \text{ GBP*} + 1.22 \text{ 20 day*} + 1.74 \text{ 100 day*}$$

Predictor	Coef	SE Coef	T	P
Constant	-0.16337	0.02985	-5.47	0.000
GBP*	0.9982	0.1927	5.18	0.000
20 day*	1.2243	0.2189	5.59	0.000
100 day*	1.7381	0.2011	8.64	0.000

S = 0.06546 R-Sq = 23.1% R-Sq(adj) = 22.9%

Durbin-Watson statistic = 2.03

Our results are quite encouraging, as they contain extremely significant t-statistics for most regression coefficients, and high R-squared values. However, an interesting phenomenon occurs in our JPY skew model—our regression coefficient for the 20 day JPY trend is insignificant. To address this problem, we drop it and rerun the regression, whose results we present on the below in Table 11:

Table 11: 3 month JPY Skew* vs. JPY Level*, 100 day JPY Trend*

The regression equation is:

$$3 \text{ month skew}^* = - 0.495 + 0.0988 \text{ JPY}^* + 0.00668 \text{ 100 day}^*$$

Predictor	Coef	SE Coef	T	P
Constant	-0.49550	0.02147	-23.08	0.000
JPY*	0.098796	0.004430	22.30	0.000
100 day*	0.006675	0.003174	2.10	0.036

S = 0.1174 R-Sq = 48.3% R-Sq(adj) = 48.2%

Durbin-Watson statistic = 1.76

As we can see, all coefficients are now statistically significant at the 95% significance level for our JPY model. As noted earlier, this is also the case for our GBP model as well (as seen in Table 10 above). Thus, we can say with a high degree of statistical confidence that the volatility skews for both JPY and GBP are positively correlated with the level of the underlying currency and the recent trend in that currency.

Given that the volatility skews for JPY and GBP were highly correlated with longer term trends in their underlying currencies, it makes sense that *daily changes* in these skews might be explained by shorter term currency trends. However, this testing of

first differences would have most likely resulted in similar results as above, so we don't continue along this line.

VIII. SUMMARY

The conclusions of our skew models are quite interesting. Our *models of skew levels* indicate that the higher the level of JPY and the stronger the JPY uptrend, the more positive the JPY skew; and the higher the level of GBP and the stronger the recent GBP uptrend, the more positive the GBP skew. In addition, we must note that our skew models for JPY offer significantly higher explanatory power than those for GBP.

Now we look to explain the relationships described above. These arguments follow from our initial thoughts regarding traders' behaviors described in Section III. We begin our discussion by offering two hypotheses to explain the effect of the underlying currency trend on the skew. The arguments we make relate to uptrends in either the JPY or GBP spot markets, but apply analogously to downtrends as well.

Our first explanation relates to buyers of option premium. We argue that as JPY (or GBP) trades up in the spot market, speculative (e.g. hedge fund and bank) players in the market expect the trend to continue and/or hedgers are forced to purchase additional upside protection. The net result means that there is greater demand for calls relative to puts (for the same level of delta). The second explanation relates to sellers of option premium. In essence, sellers of calls most likely have lost a considerable amount of money during a recent move up in the underlying spot market, and thus demand higher implied volatilities to continue selling more premium. In either situation, implied volatilities for calls increase relative to those for puts (for the same delta); thus, the skew

increases in value. These hypotheses are consistent with the belief by market players in the existence of continuing trends in JPY and GBP movements.

Now we look to explain the positive relationship between the underlying currency level and its respective skew. As we observed earlier, this relationship was much stronger for JPY than for GBP (e.g. t-stats of 22.30 for JPY and 5.18 for GBP). Thus, our explanation must address why this relationship is stronger for JPY.

One possible explanation revolves around central bank intervention in the foreign exchange markets. It is widely known that the Bank of Japan actively and consistently intervenes in the market, while the Bank of England intervenes much less frequently. Thus, all else being equal, we suspect that it is signals sent by the Bank of Japan (through its intervention) at certain JPY spot levels that places a greater influence on the volatility skew.

REFERENCES

Bates, D. 1996. Jumps in stochastic volatility: Exchange rate processes implicit in deutsche mark options. *The Review of Financial Studies* 9. pp. 69-79.

Hull, J. 2000. Options, Futures, & Other Derivatives. pp. 435-440.

Mayhew, S. 1995. Implied Volatility. *Financial Analysts Journal/July-August*. pp. 8-20.

Poon, S. and C. Granger. 2003. Forecasting Volatility in Financial Markets. *Journal of Economic Literature*. to be published summer 2003.

Rubinstein, M. 1994. Implied Binomial Trees. *Journal of Finance* 49. pp. 781-791.