Realized Volatility Fixings: Why They are Different

Xiaoquan Liu*

Shiu-yan Pong†

10 January, 2007

*Department of Accounting, Finance and Management, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK. Phone: +44 (0)1206 873849. E-mail: liux@essex.ac.uk.

†Credit Agricole Asset Management, 122 Leaden Hall Street, London EC3V 4QH, UK. Phone: +44 (0)20 70749402. E-mail: eddie.pong@ca-assetmanagement.co.uk.

Abstract

Banks and data vendors record data at different times of the day. This is known as 'fixing'. With daily exchange data, we first show that different fixings give rise to significant different estimation of realized volatility, and that some series of realized volatility are consistently higher/lower than others, which has important implications for trading volatility derivatives. We propose that the pattern is due to the intraday seasonality of volatility and hence use 30-minute high frequency exchange rate data to explore the issue. Empirical and simulation results support our conjecture that the fixings are different mainly due to the intraday variation in volatility.

When data is needed for empirical finance research, we go to a data provider such as the DataStream, only to find a number of different sources for the same data to choose from. To make the matter worse, the data from these sources are slightly but evidently different. What is less obvious however is how we should choose from them and in doing so what the economic implications are. In this article, we are going to explore this issue in the context of foreign exchange rate market.

Banks and data vendors record data at different times of the day. This is known as 'fixing'. For example, the Bloomberg London close records the exchange rate between the US dollar and the pound sterling at 18hr00 Greenwich Mean Time (GMT), whereas the Federal Reserve Bank at New York records the buying rate between the same currencies at 12hr00 Eastern Standard Time (EST), which is 17hr00 GMT.[1] Because the exchange rate fluctuates over time, the same exchange rate recorded by Bloomberg in London will differ from that recorded by the Fed in New York.

We focus on eight pairs of frequently traded exchange rates from six different sources. The exchange rates include:

- Canadian dollar against US dollar (CAD/USD)
- US dollar against pound sterling (USD/GBP)
- Japanese Yen against US dollar (JPY/USD)
- Swiss franc against US dollar (CHF/USD)
- US dollar against Euro (USD/EUR)

[1] European countries change into summer time on the last Sunday of March, whereas northern American countries do so on the first Sunday of April. All countries change back from summer time on the last Sunday of October. As a result, except the first week of April, the time lag between Europe and North America is 5 hours.

- US dollar against Australian dollar (USD/AUD)
- Norwegian krone against US dollar (NOK/USD)
- Swedish krone against US dollar (SEK/USD)

They are recorded by different banks and data vendors, including:

- Bloomberg London close (BLN) at 18hr00 GMT
- Bloomberg New York close (BNY) at 17hr00 EST, which is 22hr00 GMT
- European Central Bank mid rate (ECB) at 14hr15 Central European Time, which is 13hr15 GMT
- Federal Reserve Bank at New York buying rate (FNN) at 12hr00 EST, which is 17hr00 GMT
- Federal Reserve Bank at New York spot rate (FAM) at 10hr00 EST, which is 15hr00 GMT
- WMR rate (WMR) at 16hr00 GMT

We download daily data from the DataStream and the website for the Fed and ECB for the period from January 1999 to December 2005 inclusive.

Using a simple measure of realized volatility, we first show that for the same exchange rate, different sources give rise to very different levels of realized volatility. We find that the Bloomberg London close and the Bloomberg New York close always provide the highest realized volatility among the fixings, while WMR often leads to the lowest realized volatility. We also make pairwise comparisons between different fixings for each exchange rate and find that many of the differences are statistically significant.

The issue of volatility fixing is important for many volatility derivatives.

For example, in a variance swap, the parties involved trade between a pre-specified variance and a measure of realized variance over the life of the contract. The payoff at expiry is $N(\sigma_R^2 - K^2)$, in which σ_R is the realized volatility, K is the volatility strike price at the start of the contract, and N is the notional amount of the swap per unit of volatility. If one volatility fixing is always higher/lower than another, it will have significant impact on the final payoff of the variance swap.

This naturally leads to the question: Why are they different? We conjecture that this is mainly due to the intraday seasonality of volatility. Empirical evidence of distinct intraday pattern is first identified by Wood et al (1985) and Harris (1986) in the cash market, and Muller et all (1990) and Bailie and Bollerslev (1991) in the foreign exchange market. Since the release of data by Olsen & Associates (O&A), studies in this area have been burgeoning. They have documented a strikingly regular intraday periodicity for realized volatility, in that the intraday volatility is high at opening and closing hours of major exchanges and low during the middle of the day.[2] As different fixings are recorded at different times of a trading day, it is only natural that different fixings are systematically different.

In order to explore this possibility, we want to filter out the intraday seasonality from the volatility process and see if the remaining GARCH processes are still different.[3] With 30-minute high frequency data from 3 January 2005 to 3 March 2006 for the eight exchange rates under scrutiny, we first estimate the intraday seasonality using flexible Fourier form (FFF), following

[2]Taylor (2005) provides an excellent summary and literature review.

[3]In the literature of modelling and forecasting high frequency volatility, there is a consensus that the dynamics of the total volatility can be described as the product of a seasonality component and the remaining ARCH process. See Martens et al (2002).

Anderson and Bollerslev (1997) and Martens et al (2002).[4] Seasonality can then be filtered out from the intraday returns to give de-seasoned 30-minute returns.

We then fit GARCH (1,1)- MA(1) to the de-seasoned high frequency returns to estimate the GARCH parameters, which we will use to simulate the 30-minute return process without the seasonality component. Using these pure GARCH process, we define different daily volatility fixings, and make pair-wise comparisons between them, as we did with market data. We find that the volatility fixings from simulated processes without seasonality are no longer statistically different from each other: there are still small differences but they are not significant anymore. In this way, we demonstrate that volatility fixings are different mainly because of the intraday seasonality. Once we remove the seasonality from the processes, the resulting volatility fixings are no longer statistically different.

The rest of the article is organized as follows. In the next section, we present the statistical differences between volatility fixings with daily data. This is followed by contrasting the payoffs of a variance swap to demonstrate the economic significance. We then use high frequency data to explore the possibility that these differences arise from the intraday seasonality pattern. The last section concludes.

Daily Data and Volatility Fixing Patterns

Daily data for eight exchange rates are downloaded from the DataStream and the website of ECB and Fed for the period from January 1999, when

[4] Martens et al (2002) have performed a horse race between a number of measures of seasonality and recommend FFF as accurate and computationally efficient.

the Euro was formally launched, to December 2005 inclusive. We calculate the annualized monthly realized volatility as follows,

$$vol = \sqrt{\frac{1}{n-1} \times \sum_{t=1}^{n} (r_t - \bar{r})^2} \times \sqrt{255} \tag{1}$$

where

$$r_t = \ln \frac{S_t}{S_{t-1}}$$

and $\bar{r}$ is the sample mean.

We choose a simple measure to test whether the difference between any two fixings is statistically different from zero: the ratio between average realized volatility and its standard deviation, called z-statistic, follows standard normal distribution, with the null hypothesis that the different is zero.

Table 1 reports the average of annualized volatilities for the currencies. There is some pattern among the results. Bloomberg London (BLN) and Bloomberg New York (BNY) always produce higher realized volatility. Out of the eight currencies, Bloomberg London fixing has the highest realized volatility for three currencies: USD/GBP, JPY/USD, and USD/AUD, while Bloomberg New York fixing has the highest realized volatility for four currencies: CHF/USD, NOK/USD, SEK/USD, and USD/EUR. On the other hand, WMR consistently produces lower realized volatility. Out of the eight currencies, it has the lowest for five currencies: CHF/USD, USD/AUD, NOK/USD, SEK/USD, and USD/EUR. The Fed am fixing has the lowest realized volatility for two currencies: USD/GBP and JPY/USD.

With this general picture in mind, we want to know exactly whether the differences are significant or not. We make pairwise comparisons between the currencies for the 15 pairs of fixings. The results are reported in Table 2. The first column of Table 2 shows the two fixings that we compare and the difference is calculated as the fixing before the slash minus the one behind it. For example, for USD/GBP, we know from Table 1 that BLN has the highest realized volatility. Therefore, we are not surprised to see that the difference between BLN and all the other fixings are positive. Among them, the difference is significant between BLN and FNN (0.0026), between BLN and FAM (0.0039), and between BLN and WMR (0.0035).

There are some interesting patterns in Table 2. Across currencies, because NOK/USD and SEK/USD are relatively less frequently traded exchange rates, they have the largest number of significantly different fixings. Out of the 15 pairs of fixings, 11 are significantly different for SEK/USD, and 9 are significantly different for NOK/USD. Across fixings, because BLN and BNY are always high, and WMR is always low, the differences between realized volatilities from BLN/WMR and BNY/WMR are always significant: out of 8 currencies, 6 of them are significantly different realized volatilities for both pairs.

Furthermore, we want to examine whether the differences between fixings are significant *en masse*. We employ a multiple comparison test [Scheffe (1959) pp55-59] with the null hypothesis test that the the average volatilities are all equal for different fixings. The test statistic follows the F distribution, and the results are reported in Table 3. Consistently with the results in Table 2, the two northern European currencies NOK/USD and SEK/USD comprehensively reject the null at high statistical level, indicating again that the differences are significant. The exchange rate between Swiss franc and US dollar rejects the null at 90% level, giving reason ground that the differences between fixings are significant.

Payoffs from Trading Variance Swaps

We now put the results in Table 1 to 3 into perspective by looking at the payoffs of a variance swap. Given our discussion of volatility swap previously, the holder of a volatility swap, who profits from higher realized volatility, might be tempted to go with either BLN fixing or BNY fixing, while the seller may want to choose the WMR fixing.

In a variance swap contract, the terms and conditions would typically specify the following[5],

- Trading date, observation start date and end date
- Variance buyer and seller: The buyer profits from a higher realized volatility, while a seller profits from a lower realized volatility
- Vega amount and variance amount: Typically, the vega amount is the nominal amount of money for the contract, say \$100,000. The variance amount is defined as

  $$\text{variance amount} = \frac{\text{vega amount}}{\text{strike} \times 2}. \tag{2}$$

  It is a market practice to define the variance notional in volatility terms, so that if the realized volatility is 1 'vega' (volatility point) above the strike at maturity, the payoff is approximately equal to the variance notional.
- The underlying and the strike price: The strike is the annualized percentage of volatility multiply by 100. Therefore a strike of 10 refers to 10% annualized volatility. The fixing will also be specified, say Bloomberg New York ticker.

[5] These are based on a sample used by JP Morgan Securities, Ltd.

- Equity amount: This is the payoff of the variance swap calculated as

$$\text{equity amount} = \text{variance amount} \times \Big(\text{final reazlied volatility}^2 - \text{strike}^2\Big). \tag{3}$$

Based on the information above, Table 4 tabulates the payoffs for BNY fixing and WMR fixing as reported in Table 1. Possible strikes are evenly spaced between 6.50 and 10.50. The table summarizes the payoffs and the percentage differences between the fixings, the latter calculated as the ratio between the payoffs divided by the greater absolute payoff from the two fixings.

Variance swaps are attractive to investors who want to take on volatility risk and take directional bets on volatility. As we can see from Table 4, slight differences between the realized volatility and the strike are geared up. Therefore, a slight difference in realized volatility fixings can give rise to distinct payoffs. For example, for the exchange rate between USD/GBP when strike is 7.50, the BNY fixing provides a profit of 27,382 while the WMR fixing leads to a loss of 2,097. These lead to a percentage difference of 108%. For the exchange rate between NOK/USD when the strike is 9.50, the BNY fixing provides a profit of 68,827 while the WMR fixing leads to a loss of 22,917, with a percentage difference of 133%. It is easy to see significant economic benefits in choosing the 'right' fixing. Generally speaking, a buyer would favor higher fixings like the BNY and the BLN, while a seller would prefer lower fixings like the WMR.

So the natural question is: What are the reasons behind the differences? As fixings are recorded at different times during the day, we conjecture that the pattern is due to the intraday seasonality of volatility. This intraday seasonality sees a clear U-shape for volatility during a trading day, possibly as a result of the subsequent opening and closing of the three major trading venues in Asia, Europe, and North America. This stylized fact has been documented extensively in the volatility literature, especially with the advent of high frequency data.

In the next section, we are going to investigate whether the intraday seasonality can explain the volatility fixing pattern.

High Frequency Data, De-seasonalization, and ARCH simulations

We use 30-minute exchange rates from 3rd January 2005 to 3 March 2006. The exchange rates are mid-rate between the Euro and the US dollar, the Japanese Yen, the Australian dollar, the Canadian dollar, the Swiss franc, the pound sterling, the Norwegian krone, and the Swedish krone. Cross-rates between the currencies can be derived[6].

To model the intraday seasonality, we follow Andersen and Bollerslev (1997) and Martens et al (2002) and use the flexible Fourier form (FFF). Define

$$x_{d,n} = \ln s_{d,n}^2 + \ln Z_{d,n}^2 \tag{4}$$

where $s_{d,n}$ is the return on day d and period n, and $Z_{d,n}$ is an i.i.d process with mean zero and variance one. The seasonal pattern is estimated by OLS,

$$\hat{\hat{x}}_{d,n} = c + \sum_{i=1}^{p} \left(\gamma_i \cos \frac{2\pi i n}{N} + \delta_i \sin \frac{2\pi i n}{N} \right) \tag{5}$$

where c is a constant, p is set equal to 4,[7] and N is the total number of

[6]Summary statistics of the realized volatilities are available from the authors.

[7]Different specifications have been used but all of them lead to similar results.

intraday periods. The intraday seasonal index is then defined as,

$$s_{d,n} = \exp\left(\frac{\hat{\hat{x}}_{d,n}}{2}\right). \tag{6}$$

To filter out seasonality, we divide the 30-minute log returns $r_{d,n}$ by the seasonality index to obtained the de-seasoned returns $\tilde{r}_{d,n}$,

$$\tilde{r}_{d,n} = \frac{r_{d,n}}{s_{d,n}}. \tag{7}$$

Figure 1 presents the seasonality index for USD/GBP and USD/EUR. There is a clear pattern of highs and lows on a weekday and the pattern is strikingly similar over the week except Friday. Friday is more volatile than the other four weekdays, as more macroeconomics news are announced on this day [see Ederington and Lee (1993, 1995)].

Following Anderson and Bollerslev (1997) and Martens et al (2002), we fit a GARCH (1,1)-MA(1) model with t-distribution to the de-seasoned returns. The MA component is intended to control for the weak first-order autocorrelation in currency returns. We use t-distribution for the asset returns as currency returns are known to have significant non-normality in the high-frequency domain. The dynamics of the conditional mean μ and the conditional variance h_t of the 30-minute de-seasoned returns r_t are specified as follows,

$$h_t = \omega + \alpha\varepsilon_{t-1}^2 + \beta h_{t-1} \tag{8}$$

$$\mu_t = \xi + \theta\varepsilon_{t-1}.$$

For GARCH model having conditional scaled t-distribution with ν degrees of freedom, the density of $D(0,1)$ is

$$f(z) = \pi^{-\frac{1}{2}}(\nu-2)^{-\frac{1}{2}}\left(\frac{\Gamma[(\nu+1)/2]}{\Gamma(\nu/2)}\right)\left(1+\frac{z^2}{\nu-2}\right)^{-\frac{\nu+1}{2}}, \nu > 2. \tag{9}$$

With the estimated GARCH parameters[8], we then simulate the 30-minute return process without the impact of seasonality. For each exchange rate and for each fixing, we carry out Monte Carlo simulations, and then define fixings in the same way as the market data. For example, for the BLN fixing, we start at the 36th return period (corresponding to 18hr00), add up 48 returns to obtain a one-day return. We do this for a period of 30 days. We then estimate annualized monthly volatility based on these daily returns. We repeat this for 100,000 times and then make pairwise comparisons between different exchange rates and fixings. The results are reported in Table 5.

Table 5 can be viewed as the de-seasoned counterpart of Table 2. In this table, we simulate volatility fixings with parameters from the de-seasoned returns to see if there are still significant differences between them. Results indicate that except for three pairs of fixings for SEK/USD and two pairs for USD/EUR, all the differences are no longer significant. Therefore, all the realized volatilities are statistically the same. This confirms our conjecture that the volatility fixings from market data are different mainly because of the intraday seasonality pattern of the volatility.

Conclusion

In this study, we are interested in the questions of how volatility fixings are different from each other and more importantly why.

We first use 7 years of daily market data to establish the fact that volatility fixings are not necessarily the same as each other, although they all come from the same return process. In particular, the BLN and the BNY are always higher than the others and the WMR fixing is always the lowest. Less

[8]The parameter estimates are available from the authors.

frequently traded exchange rates, including NOK/USD and SEK/USD, have more pairs of statistically different fixings. These have significant bearing on the payoffs of volatility derivatives.

We then examine the economic significance of the results by contrasting the payoffs to a variance swap with a series of possible strike prices. We see that subtle distinction between fixings can lead to huge differences in the final payoffs.

We explore the reasons behind the phenomenon. We conjecture that the intraday volatility seasonality is the main reason. To test this proposition, we use 30-minute high frequency data, first estimate the seasonal pattern with fast fourier form (FFF) and then de-seasoned the returns by the seasonality index. We fit a GARCH (1,1)-MA(1) model to the de-seasoned data to obtain GARCH parameters, and use the parameters to simulate the return processes that are free from the intraday volatility seasonality. Results show that, when the volatility process is a pure GARCH process, volatility fixings are no longer different from each other, thus confirming our conjecture.

References

Andersen, T. G. and T. Bollerslev, 1997, Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance*, 4, 115-158.

Andersen, T. G. and T. Bollerslev, 1998, DM-dollar volatility: intraday activity patterns, macroeconomic announcements and long run dependencies, *Journal of Finance*, 53, 219-265.

Areal, N. M. P. C. and S. J. Taylor, 2002, The realized volatility of FTSE-100 futures prices, *Journal of Futures Markets*, 22, 627-648.

Ederington, L. H. and J. H. Lee, 1993, How markets process information, *Journal of Finance*, 48, 1161-1191.

Ederington, L. H. and J. H. Lee, 1995, The short run dynamics of the price adjustment to new information, *Journal of Financial and Quantitative Analysis*, 30, 117-134.

Martens, M., Y-C Chang, and S. J. Taylor, 2002, Intraday volatility forecasts using different seasonality adjustment methods, *Journal of Financial Research*, 25, 283-297.

Sanciaume, A. and A. Sandilya, 2004, Realized volatility and fixing times, memo, Lehman Brothers Plc.

Scheffe, H., 1959, *The analysis of variance*, New York: Wiley.

Taylor, S. J., 2005, *Asset price dynamics, volatility, and prediction*, Princeton University Press.

Taylor, S. J. and X. Xu, 1997, The incremental volatility information in one million foreign exchange quotations, *Journal of Empirical Finance*, 4, 317-340.

Zhou, B., 1996, High-frequency data and volatility in foreign-exchange rates, *Journal of Business and Economic Statistics*, 14, 45-52.

Table 1. The average of annualized volatility from January 1999 to December 2005 for different fixings

This table reports the average of annualized volatility for seven years from January 1999 to December 2005 for eight currency pairs and six different fixings. BLN stands for Bloomberg London fixing, BNY Bloomberg New York fixing, ECB European Central Bank fixing, FNN Fed noon fixing, and FAM Fed fixing at 10hr00 (EST). Blue font indicates the maximum and red font indicates the minimum among the fixings. All volatilities are annualized and expressed in percentages.

	CAD/USD	USD/GBP	JPY/USD	CHF/USD	USD/AUD	NOK/USD	SEK/USD	USD/EUR
BLN	6.582	7.833	9.601	10.383	10.688	9.753	10.125	9.697
BNY	6.539	7.769	9.475	10.560	10.500	10.165	10.408	9.921
ECB	6.895	7.728	9.333	10.316	10.445	9.884	9.932	9.920
FNN	6.620	7.575	9.545	10.273	10.279	9.610	9.927	9.649
FAM	6.769	7.447	9.090	10.163		9.330	9.582	9.591
WMR	6.603	7.479	9.380	10.080	10.076	9.268	9.574	9.429

Table 2. Differences between the fixings of eight exchange rates from January 1999 to December 2005

This table reports the pairwise differences between the average annualized monthly volatility for eight exchange rates and six fixings, using daily exchange rate between January 1999 to December 2005. All numbers are in percentage. Blue font indicates significance at 95% level, and red font indicates significance at 90% level.

	CAD/USD	USD/GBP	JPY/USD	CHF/USD	USD/AUD	NOK/USD	SEK/USD	USD/EUR
BLN-BNY	0.04	0.06	0.13	-0.18	0.19	-0.41	-0.28	-0.22
BLN-ECB	-0.31	0.11	0.27	0.07	0.24	-0.13	0.19	-0.22
BLN-FNN	-0.04	0.26	0.06	0.11	0.41	0.14	0.20	0.05
BLN-FAM	-0.19	0.39	0.51	0.22		0.42	0.54	0.11
BLN-WMR	-0.02	0.35	0.22	0.30	0.61	0.49	0.55	0.27
BNY-ECB	-0.36	0.04	0.14	0.24	0.06	0.28	0.48	0.00
BNY-FNN	-0.08	0.19	-0.07	0.29	0.22	0.55	0.48	0.27
BNY-FAM	-0.23	0.32	0.38	0.40		0.83	0.83	0.33
BNY-WMR	-0.06	0.29	0.09	0.48	0.42	0.90	0.83	0.49
ECB-FNN	0.27	0.15	-0.21	0.04	0.17	0.27	0.00	0.27
ECB-FAM	0.13	0.28	0.24	0.15		0.55	0.35	0.33
ECB-WMR	0.29	0.25	-0.05	0.24	0.37	0.62	0.36	0.49
FNN-FAM	-0.15	0.13	0.46	0.11		0.28	0.35	0.06
FNN-WMR	0.02	0.10	0.17	0.19	-0.20	0.34	0.35	0.22
FAM-WMR	-0.17	0.03	0.29	-0.08		-0.06	-0.01	0.16

Table 3. Test statistic for simultaneous comparison

This table reports the test statistic for the simultaneous comparison between all fixings for a particular currency. The null hypothesis is that the average volatility is equal for all fixings. Blue font indicates significance at 95% level and red font indicates significance at 90% level.

currency	F test stat
CAD/USD	1.33
USD/GBP	1.58
JPY/USD	1.56
CHF/USD	1.85
USD/AUD	1.48
NOK/USD	5.44
SEK/USD	4.47
USD/EUR	1.63

Table 4. Payoffs from trading variance swap under different strike prices

This table reports the payoffs from trading variance swaps with the BNY volatility fixing and WMR volatility fixing as shown in Table 1. The strike price is evenly spaced between 6.50 and 10.50. The vega amount is 100,000. "Var amt" refers to variance amount according to equation (3), payoff are calculated according to equation (4), and the percentage difference are the ratio between the two payoffs divided by the bigger absolute payoff.

strike	BNY	WMR	Diff	BNY	WMR	Diff	BNY	WMR	Diff	BNY	WMR	Diff
	CAD/USD			USD/GBP			JPY/USD			CHF/USD		
6.50	3,912	10,382	62%	139,287	105,273	24%	365,582	351,803	4%	532,797	456,588	14%
7.50	-89,943	-84,336	6%	27,382	-2,097	108%	223,504	211,563	5%	368,424	302,376	18%
8.50	-173,479	-168,532	3%	-69,957	-95,968	27%	103,092	92,555	10%	230,962	172,685	25%
9.50	-249,955	-245,528	2%	-157,330	-180,603	13%	-2,497	-11,924	79%	111,914	59,771	47%
10.50	-321,388	-317,383	1%	-237,584	-258,641	8%	-97,497	-106,027	8%	6,017	-41,160	115%
	USD/AUD			NOK/USD			SEK/USD			USD/EUR		
6.50	523,077	455,968	13%	469,825	335,737	29%	508,280	380,088	25%	432,125	358,893	17%
7.50	360,000	301,839	16%	313,848	197,639	37%	347,176	236,077	32%	281,175	217,707	23%
8.50	223,529	172,210	23%	182,807	80,270	56%	212,214	114,185	46%	153,978	97,977	36%
9.50	105,263	59,346	44%	68,827	-22,917	133%	95,139	7,429	92%	43,033	-7,073	116%
10.50	0	-41,544	100%	-32,966	-115,972	72%	-9,160	-88,517	90%	-56,304	-101,638	45%

Table 5. Simulated differences between volatility fixings without seasonality

This table reports the simulated differences between the de-seasoned average annualized monthly volatility for eight exchange rate pairs and six fixings, using 100,000 simulations. All numbers are in percentage. Blue font indicates significance at 95% level.

	CAD/USD	USD/GBP	JPY/USD	CHF/USD	USD/AUD	NOK/USD	SEK/USD	USD/EUR
BLN-BNY	0.02	0.002	0.002	-0.001	0.002	-0.002	0.00002	-0.007
BLN-ECB	0.01	0.003	-0.009	-0.001	0.003	-0.003	0.009	-0.0003
BLN-FNN	0.01	-0.002	-0.001	-0.0006	0.0003	-0.004	0.003	0.004
BLN-FAM	0.02	-0.002	-0.006	0.004		-0.003	0.004	0.003
BLN-WMR	0.02	0.004	-0.001	-0.002	-0.0008	0.0008	-0.001	-0.0008
BNY-ECB	-0.003	0.001	-0.008	0.0001	0.0005	-0.002	0.009	0.007
BNY-FNN	-0.004	-0.004	0.0008	0.0008	-0.002	-0.002	0.003	0.01
BNY-FAM	0.003	-0.002	-0.004	0.006		-0.001	0.004	0.01
BNY-WMR	0.009	-0.001	0.0007	-0.0003	-0.003	0.002	-0.001	0.006
ECB-FNN	-0.001	-0.005	0.008	0.0007	-0.003	-0.0004	-0.006	0.004
ECB-FAM	0.006	-0.003	0.003	0.006		0.0007	-0.006	0.003
ECB-WMR	0.01	-0.003	0.009	-0.0004	-0.004	0.004	-0.001	-0.0005
FNN-FAM	0.007	0.002	-0.005	0.005		0.001	0.0004	-0.0008
FNN-WMR	0.01	0.002	-0.0001	-0.001	-0.001	0.005	-0.004	-0.004
FAM-WMR	0.006	0.0006	0.005	-0.006		0.004	-0.005	-0.004

Figure 1. Daily seasonality for USD/GBP and USD/EUR, estimated using 30-minute high frequency data from 3 January 2005 to 3 March 2006