UNIVERSITY OF EXETER
XFi Centre for Finance and Investment

Day-of-the-Month Effects in the Performance of Momentum Trading Strategies in the Foreign Exchange Market

Richard D.F. Harris
University of Exeter

Evarist Stoja
University of Bristol

Fatih Yilmaz
Bank of America

Paper Number: 08/07

October 2008

Abstract

In this paper, we document a very strong day-of-the-month effect in the performance of momentum strategies in the foreign exchange market. We show that this seasonality in trading strategy performance is attributable to seasonality in the conditional volatility of foreign exchange returns, and in the volatility of conditional volatility. Indeed a two-factor model employing conditional volatility and the volatility of conditional volatility explains as much as 70 percent of the intra-month variation in the Sharpe ratio. We further show that the seasonality in volatility is in turn closely linked to the pattern of US macroeconomic news announcements, which tend to be clustered around certain days of the month.

Keywords: Momentum; Moving average rules; Seasonality; Conditional volatility.

Address for correspondence: Professor Richard D. F. Harris, Xfi Centre for Finance and Investment, University of Exeter, Exeter EX4 4PU, UK. Tel: +44 (0) 1392 263215, Fax: +44 (0) 1392 263242 Email: R.D.F.Harris@exeter.ac.uk.

Momentum trading strategies are widely used in the foreign exchange market by currency fund managers and commodity trading advisors. Indeed, for short horizons, foreign exchange dealers are more likely to use technical indicators, such as those that exploit momentum, than to base their forecasts on economic fundamentals.[1] The most common momentum trading strategy is the moving average rule, which generates directional trading signals on the basis of the intersection of moving averages of the current and lagged exchange rate measured over different horizons. The moving average rule is widely used in the currency fund management industry, either directly as a trading indicator or indirectly as a benchmark for other, more sophisticated momentum trading strategies. A number of studies have demonstrated that momentum trading strategies, and in particular moving average rules, have in the past produced excess returns over buy-and-hold strategies. Moreover, while there remains some contention about the source of these excess returns, the evidence suggests that they are not simply compensation for transaction costs or risk.[2]

In this paper, we document a very pronounced day-of-the-month effect in the performance of moving average rules in the foreign exchange. In particular, using a simple two-lag moving average strategy applied to an equally weighted portfolio of G10 currencies over the period 1998-2008, we show that the Sharpe ratio is close to zero in the first half of the month, but rises sharply in the second half of the month, particularly during the third week. A similar day-of-the-month pattern exists in the maximum drawdown of the moving average strategy. We show that this seasonality in the performance of momentum trading strategies is largely explained by corresponding seasonality in the conditional volatility of foreign exchange returns, and in the volatility of conditional volatility. Indeed a two-factor model that incorporates the conditional volatility of returns and the volatility of conditional volatility explains almost 70 percent of the intra-month variation in the Sharpe ratio. We further show that seasonality in the volatility dynamics of foreign exchange returns is itself potentially explained by the pattern of US macroeconomic announcements, which tend to be clustered around certain days of the month.

The Day-of-the-Month Effect in Momentum Trading Strategies

[1] See, for example, the survey by Taylor and Allen (1992).
[2] See, for example, Okunev and White (2003).

Momentum trading strategies rely on successfully extracting information about the short-term trend in the exchange rate, and are commonly used by currency fund managers and commodity trading advisors to provide directional trading signals. The most widely used momentum trading strategy is the moving average rule, which compares a short-run moving average of the current and lagged exchange rate with a long-run moving average, where ‘short-run’ and ‘long-run’ are arbitrarily defined. Here, we consider the simplest, but also widely used moving average rule, the MA(1, 2), which amounts to comparing the current exchange rate with the one-period lagged exchange rate. In particular, the MA(1, 2) rule infers a positive trend if today’s exchange rate is higher than yesterday’s exchange rate, and a negative trend if it is lower.

We apply the MA(1, 2) rule to an equally weighted portfolio of G10 ex-USD currencies measured against the USD. Daily data were obtained from DataStream for the period 01/05/1997 to 31/05/2007. We use the MA(1, 2) strategy to generate buy and sell signals and calculate the Sharpe ratio and maximum drawdown for each day of the month traded, averaged over the 10 year period. We include the carry interest earned on the domestic or foreign currency while it is held, but ignore trading costs, which are negligible for monthly trading in the foreign exchange market.

Figure 1 reports the average Sharpe ratio by day of the month, while Figure 2 reports the maximum drawdown by day of the month. The pattern of performance of the MA(1, 2) rule over different days of the month is striking. The Sharpe ratio is close to zero during the first half of the month until day 12, when it starts to rise sharply, peaking initially around days 13-15, and then again at around days 20-22. A similar pattern is observed with the maximum drawdown of the MA(1, 2) strategy, with its lowest value occurring during days 21-22 of the month.

[Figures 1 and 2]

The Importance of Volatility Dynamics

Christoffersen and Diebold (2006) have shown that directional predictability is inextricably linked to the dynamics of conditional volatility. In particular, they show that even if the conditional mean return is constant (i.e. even if the mean return is unpredictable), predictability in the conditional volatility of returns implies predictability in the sign of returns, as long as the conditional mean return is non-zero. This is because the probability that the return is negative depends on the size of the tail of the return distribution to the left of zero. If the dispersion of returns around the (non-zero) mean is predictable, then so too will be the size of the tail of the distribution below zero, and hence also the probability of a negative return. The Appendix gives more details of this result.

Christoffersen and Diebold (2006) note a number of interesting consequences of this result. Firstly, the sensitivity of sign predictability to changes in the variance is always negative, i.e. increasing volatility always reduces sign predictability, but the relationship is non-monotonic. In particular, the sensitivity approaches zero either as the information ratio (the ratio of the mean return to the standard deviation of returns) approaches zero (because as volatility becomes large, or as the mean approaches zero, the probability of a positive return approaches its unconditional value of 0.5) or as the information ratio rises very high (because as volatility approaches zero for a given positive mean, the probability of a positive return approaches its unconditional value of unity). The absolute sensitivity is maximised when the information ratio is about 1.4. As noted by Christoffersen and Diebold (2006), this is rather high for an information ratio, and is unlikely to be obtained systematically in practice. However, the frequency with which this information ratio is achieved depends on the volatility of volatility. Secondly, the correlation between the forecast sign and the realized sign is determined by the standard deviation of the forecast sign, which again depends on the volatility of volatility. Thus sign predictability, and hence trading strategy performance, should be negatively associated with volatility and positively associated with the volatility of volatility.

To investigate the importance of these effects, we estimate the conditional variance of returns for the equally-weighted portfolio using an exponentially weighted moving average (EWMA) estimator with a decay factor of 0.95. Figure 3 reports the conditional standard deviation of returns for each day of the month, averaged over the 10-year sample. The seasonal pattern of volatility over the sample period matches the seasonal pattern of the Sharpe ratio quite closely. In particular, volatility is notably lower in the second half of the month, particularly during days 13-15 and days 20-23.

[Figure 3]

Figure 4 plots the volatility of conditional volatility for each day of the month, measured by the standard deviation (across the 120 months in the sample) of the EWMA standard deviation of returns, for each day of the month. Again, there is a clear seasonal pattern to the volatility of volatility, which tends to be higher in the second half of the month than the first half, peaking at around days 21-22. Thus it would appear that the performance of the MA(1, 2) trading strategy is highest when volatility is low and the volatility of volatility is high. This is consistent with the theoretical results of Chistoffersen and Diebold (2006) described above.

[Figure 4]

To further explore the importance of volatility dynamics, we estimate a regression of the average Sharpe ratio on each day of the month on conditional volatility and the volatility of conditional volatility. The results of this regression are reported in Table 1. The two-factor model explains almost 70 percent of the intra-month variation in the performance of the MA(1, 2) trading strategy applied to the equally-weighted G10 portfolio. Both variables are highly significant at conventional significance levels and with the expected signs. Moreover, the model appears to be well specified in terms of the diagnostic tests.

[Table 1]

Figure 5 plots the actual average Sharpe ratio across days of the month together with the fitted values from the regression model. Clearly the model gives a good overall fit to the actual Sharpe ratio, particularly around the low volatility days when the Sharpe ratio tends to be highest.

[Figure 5]

The Influence of News

An obvious potential explanation of the seasonality in volatility dynamics – and hence of the seasonality in the performance of moving average trading strategies – is the timing of news announcements. Such a hypothesis is inevitably difficult to test directly since (a) it is impossible to gather information on all news announcements that might affect exchange rates and (b) it is difficult to assess the relative importance of different types of news announcement. Here we take a simplistic approach and examine the distribution of 36 different types of US macroeconomic data releases across days of the month. Figure 6 plots the number of data releases for each day of each month from 2002-2007. The seasonality in the volatility dynamics shown in Figure 3 and Figure 4 is evidently at least partially explained by the impact of news. The low volatility (and high volatility of volatility) period, which is around days 21-23 of the month, is the period with the lowest concentration of US data releases. This macroeconomic news announcement effect on volatility is well documented in the literature.[3]

[Figure 6]

Summary and Implications

There is now considerable evidence that momentum trading strategies – and in particular, moving average rules – are able to generate excess returns that are not simply compensation for transaction costs or risk. In this paper, we investigate the drivers of momentum strategies excess returns employing the directional predictability framework of Christoffersen and Diebold (2006). We first document a very pronounced calendar effect in both the Sharpe ratio and maximum drawdown of a simple moving average rule applied to an equally weighted portfolio of G10 currencies over the period 1997 to 2007. We show that this calendar effect in the

[3] See, for example, Ederington and Lee (1993), Kim (1999), Andersen, Bollerslev, Diebold and Vega (2003, 2007).

performance of momentum trading strategies is closely associated with a corresponding calendar effect in the conditional volatility of foreign exchange returns, and in the volatility of conditional volatility. Indeed, a simple two-factor model including conditional volatility and the volatility of conditional volatility explains almost 70 percent of the intra-month variation in the Sharpe ratio of the equally-weighted portfolio. We further show that that the seasonality in volatility dynamics may be driven by the pattern of US macroeconomic news announcements, which tend to be clustered around certain days of the month. Our results clearly have important implications for foreign exchange traders. In particular, they suggest that momentum based trading strategies should be tailored to exploit the pattern of news announcements, thus helping to avoid periods of higher volatility (and low volatility of volatility) when trading strategy performance is lowest.

References

Andersen, T.G., T. Bollerslev, F.X. Diebold and C. Vega, 2003, "Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange", *American Economic Review* 93.1, 38-62.

Andersen, T.G., T. Bollerslev, F.X. Diebold and C. Vega, 2007, "Real-time price discovery in global stock, bond and foreign exchange markets", *Journal of International Economics* 73, 251–277.

Christoffersen, P.F. and F.X. Diebold, 2006, "Financial Asset Returns, Direction-of-Change Forecasting and Volatility Dynamics", *Management Science*, 52, 1273-1287.

Ederington , L. H. and J. H. Lee, 1993, "How Markets Process Information: News Releases and Volatility", *Journal of Finance*, Vol. 48 (4), 1161-1191.

Kim, S.J., 1999 "Do macro-economic news announcements affect the volatility of foreign exchange rates? Some evidence from Australia", *Applied Economics* 31:12, 1511 – 1521.

Okunev, J. and D. White, 2003, "Do momentum-based strategies still work in foreign currency markets?" *Journal of Financial and Quantitative Analysis* 38, 425-447.

Taylor, M.P. and H. Allen, 1992, "The Use of Technical Analysis in the Foreign Exchange Market", *Journal of International Money and Finance*, 11, 304-314.

Appendix: Sign Predictability and Volatility Dynamics

Suppose that returns are drawn from a normal distribution as follows:

$$r_t \mid \Omega_{t-1} \sim N(\mu, \sigma_t^2) \tag{A1}$$

where r_t is the return between period t–1 and period t, Ω_{t-1} is the time t–1 information set, μ is the constant conditional mean and σ_t^2 is the volatility of r_t conditional on the time t–1 information set. As shown by Christoffersen and Diebold (2006), as long as $\mu \neq 0$, equation (A1) implies that the sign of returns is predictable. To see this, we can write the probability of the return being positive, conditional on Ω_{t-1}, as

$$\begin{aligned}\Pr_t(r_t > 0) &= 1 - \Pr_t(r_t < 0) \\ &= 1 - \Pr_t\left(\frac{r_t - \mu}{\sigma_t} < \frac{-\mu}{\sigma_t}\right) \\ &= 1 - F\left(\frac{-\mu}{\sigma_t}\right)\end{aligned} \tag{A2}$$

where $F(.)$ is the cumulative standard normal density function. From (A2) it is clear that if σ_t^2 is predictable, then so too is $\Pr_t(r_t > 0)$. If $\mu > 0$, the unconditional probability of the return being positive, $\Pr(r_t > 0)$, will be greater than 0.5 but predictable variations in the conditional probability, $\Pr_t(r_t > 0)$, will be determined by predictable variations in σ_t^2.

Differentiating (A2) with respect to σ_t, the 'sensitivity' of sign predictability to changes in volatility is given by

$$\frac{\partial \Pr_t(r_t > 0)}{\partial \sigma_t} = -f\left(\frac{-\mu}{\sigma_t}\right)\left(\frac{\mu}{\sigma_t^2}\right) \tag{A3}$$

where $f(.)$ is the standard normal density of returns. The sensitivity of sign predictability to changes in the variance is always negative, i.e. increasing volatility always reduces sign predictability, but the relationship is non-monotonic. In particular, the sensitivity approaches zero either as the information ratio, μ/σ_t, approaches zero (because as volatility becomes large, or as the mean approaches zero, $\Pr_t(r_t > 0)$ approaches its unconditional value of 0.5) or as the information ratio rises very high (because as volatility approaches zero for a given positive mean, $\Pr_t(r_t > 0)$ approaches its unconditional value of unity). By differentiating (A3) with respect to σ_t and solving for μ/σ_t, we find that the absolute sensitivity is maximised when $\mu/\sigma_t = \sqrt{2}$.

Figure 1 Sharpe Ratio

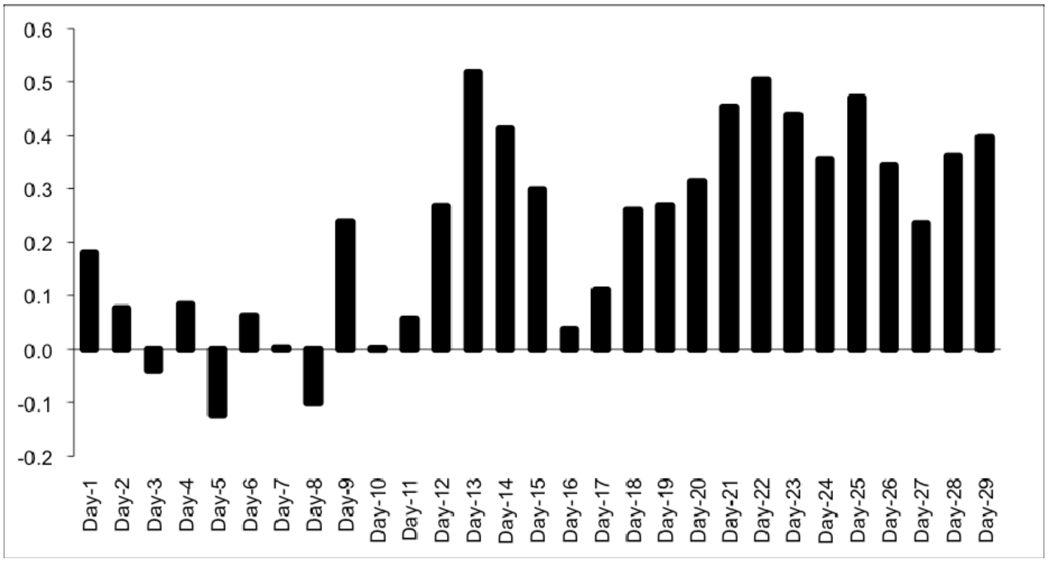

Notes: The figure reports the Sharpe ratio for each day of the month traded over the period 01/05/1997 to 31/05/2007.

Figure 2 Maximum Drawdown

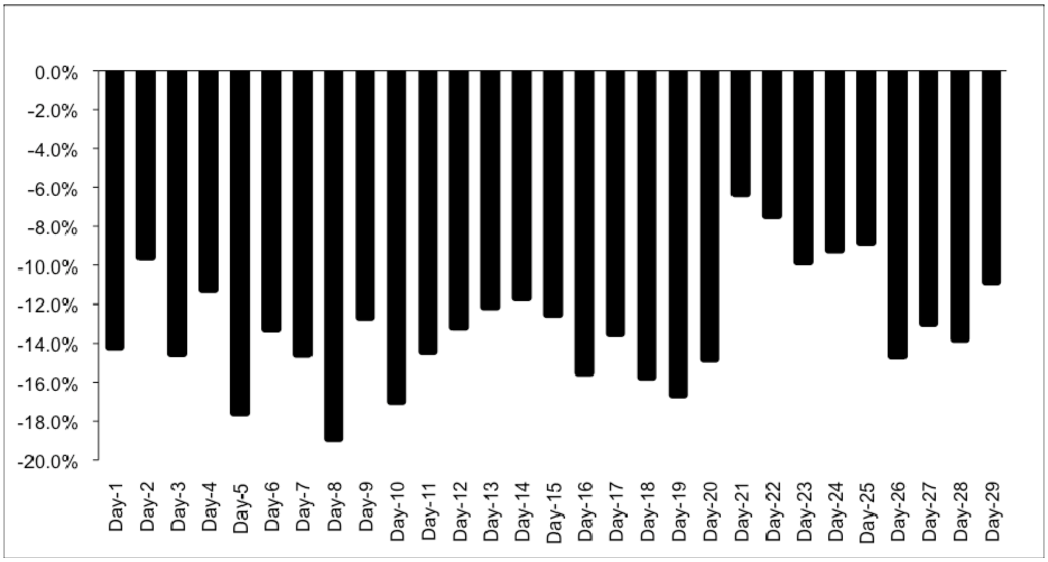

Notes: The figure reports the maximum drawdown for each day of the month traded over the period 01/05/1997 to 31/05/2007.

Figure 3 Conditional Volatility

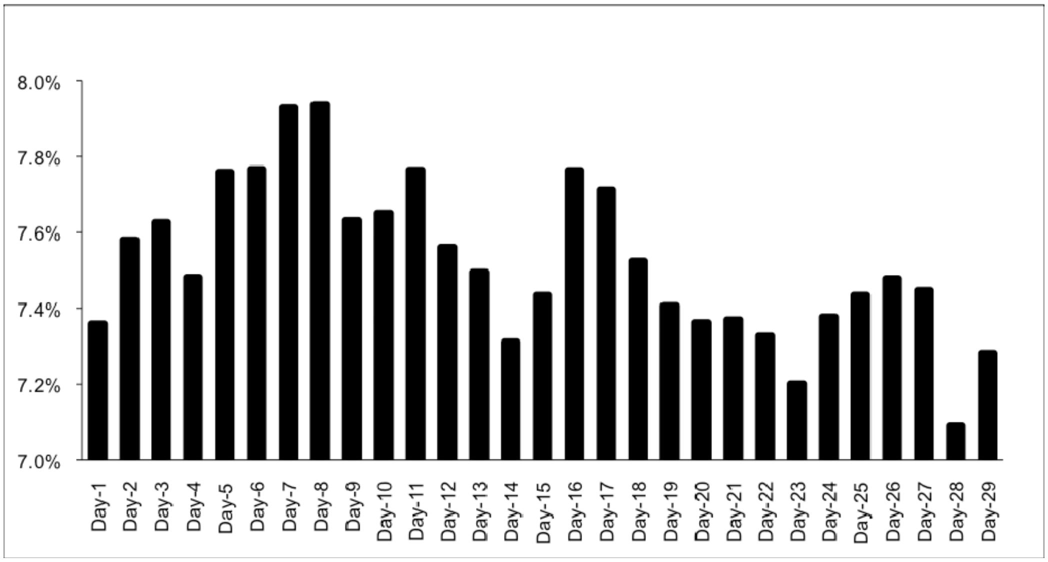

Notes: The figure reports the average conditional standard deviation for each day of the month traded over the period 01/05/1997 to 31/05/2007.

Figure 4 Volatility of Conditional Volatility

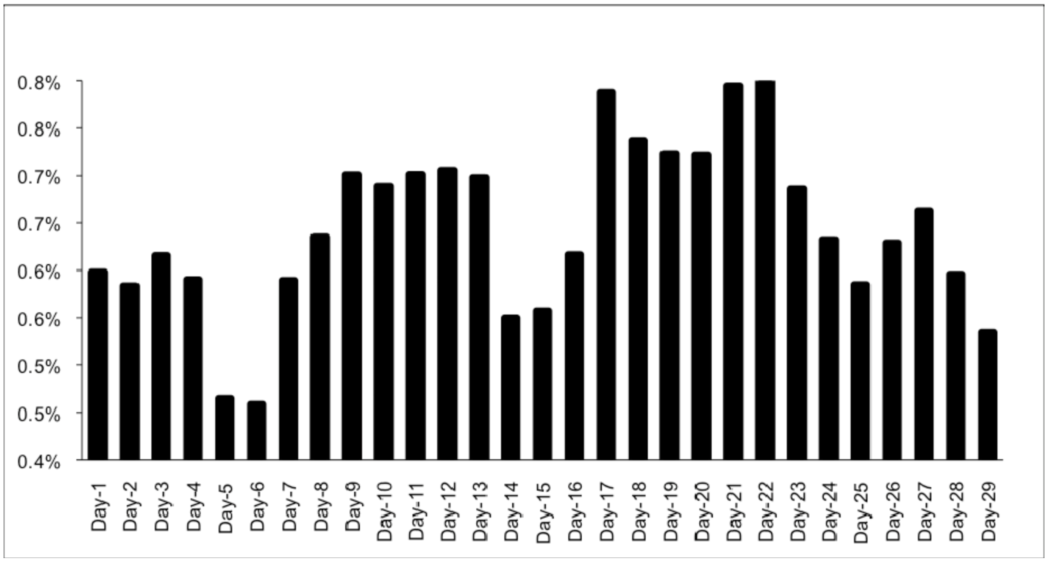

Notes: The figure reports the standard deviation across months of the conditional standard deviation for each day of the month traded over the period 01/05/1997 to 31/05/2007.

Figure 5 Sharpe Ratio Variations Across Days of the Month and Model Fit

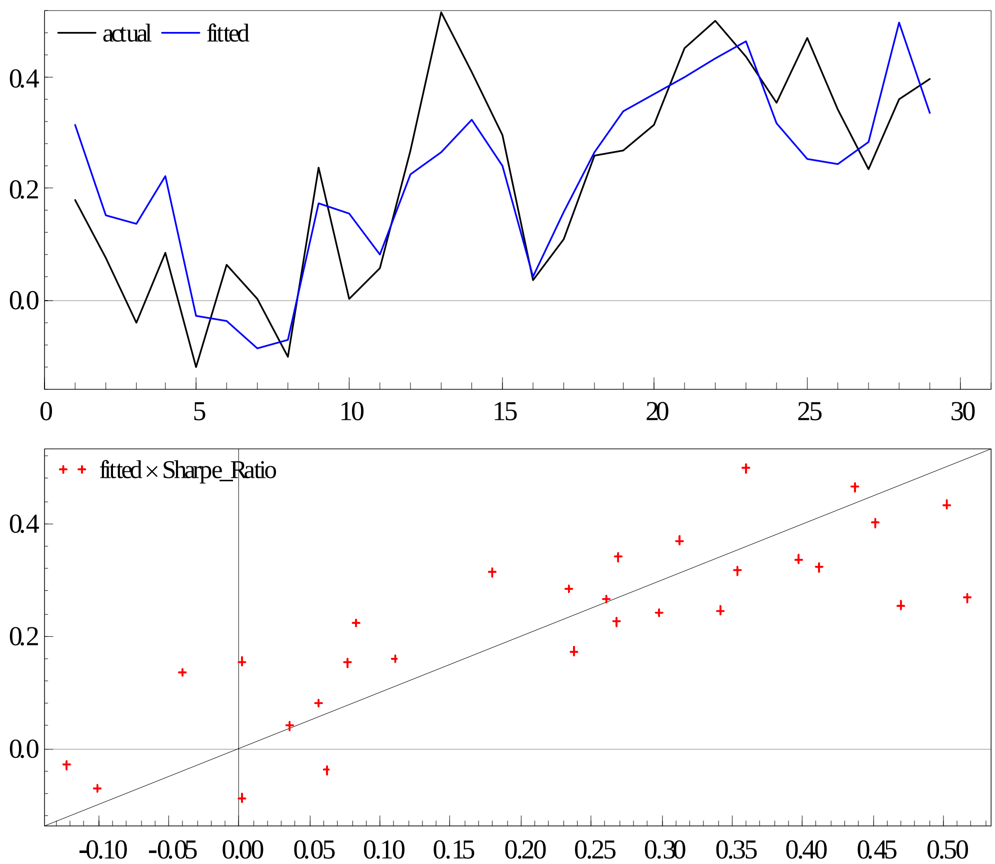

Notes: the figure reports the actual Sharpe ratio variations across days of the month and the fitted values from a regression of the Sharpe ration on conditional volatility and the volatility of conditional volatility over the period 01/05/1997 to 31/05/2007.

Figure 5 The Pattern of US Macroeconomic News Announcements

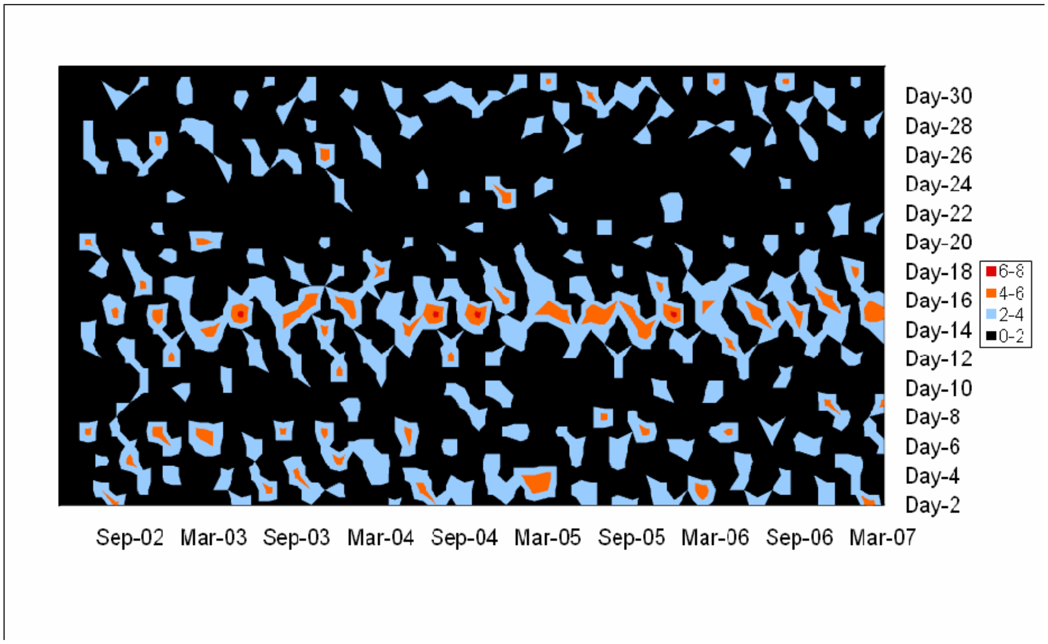

Notes: the figure reports the frequency of US macroeconomic news announcements by day of the month over the period 01/05/2002 to 31/05/2007.

Table 1 Two-Factor Regression Results

	Coefficient	t-statistic	P-value
α_0	5.14	6.57	0.00
α_1	-0.70	-6.99	0.00
α_2	0.48	2.06	0.05
Adjusted R^2	0.69		

Diagnostic Tests	p-value
No serial correlation	0.27
Homoscedasticity	0.67
Normality	0.62

Notes: The table reports the estimated coefficients of the two-factor regression for the period 01/05/1997 to 31/05/2007. Standard errors for the estimated parameters are reported in parentheses. The table also reports the p-values for tests of the null hypotheses of no autocorrelation, no heteroscedasticity and normality, respectively.