High Frequency Equity Pairs Trading: Transaction Costs, Speed of Execution and Patterns in Returns

David Bowen[a]
Centre for Investment Research, UCC

Mark C. Hutchinson[b]
Department of Accounting, Finance and Information Systems, and Centre for Investment Research, UCC

Niall O'Sullivan[c]
Department of Economics, and Centre for Investment Research, UCC

March 2010

Abstract

In this paper we examine the characteristics of high frequency pairs trading using a sample of FTSE100 constituent stocks for the period January to December 2007. We show that the excess returns of the strategy are extremely sensitive both to transaction costs and speed of execution. When we specify a moderate level of transaction costs (15 basis points) the excess returns of the strategy are reduced by more than 50%. Likewise, when we implement a wait one period restriction on execution the returns of the strategy are eliminated. When we further examine the time series properties of pairs trading returns we see that the majority of returns occur in the first hour of trading. Finally, we find that the excess returns bear little exposure to traditional risk factors but are weakly related to market and reversal risk factors.

Keywords: Intra Day, Pairs trading, Reversal strategies.

[a] David Bowen, Centre for Investment Research, University College Cork, College Road, Cork, Ireland Tel + 353 21 4903783 Email: d.bowen@student.ucc.ie

[b] Corresponding Author: Dr Mark Hutchinson, Department of Accounting, Finance and Information Systems, O'Rahilly Building, University College Cork, College Road, Cork, Ireland. Tel: +353 21 4902597 Email: m.hutchinson@ucc.ie

[c] Niall O'Sullivan, Department of Economics, Aras na Laoi, University College Cork, College Road, Cork, Ireland Tel + 353 21 4902765 Email: niall.osullivan@ucc.ie

Relative value trading strategies are widely used by hedge fund managers and proprietary trading desks. The most common approach to relative value trading in the single stock market is pairs trading, where two stocks which move together are identified and long/short positions are taken when they diverge abnormally, profiting when the prices converge. There is much evidence on the profitability of relative value strategies, such as pairs trading, at lower frequencies and the evidence suggests that the excess returns from following the strategies are not simply compensation for transaction costs or risk.[1] Likewise, there is plenty of evidence on the characteristics of intra day data, including the, well documented, negative serial correlation in returns and U shaped pattern of intra day volume and volatility. However, to date there is little research examining the attributes of intra-day relative value trading strategies.

In this paper, our focus is on high frequency intra-day pairs trading. We document several interesting characteristics of the returns from following such a strategy. First, we show that the returns are extremely sensitive to the magnitude of transaction costs and speed of execution. When we specify a range of transaction costs from zero to 15 basis points the excess returns of the strategy vary from 15.2% to 7.0%. Further, when we impose a wait one period for execution restriction excess returns are eliminated. Second, we find evidence that the returns are quite unrelated to traditional risk factors. We specify a range of factors including broad market, momentum and reversal measures. These models explain less than 2% of the portfolio returns with only weak sensitivity to broad market and reversal factors observed. Third, we identify the periods in the day when the strategy generates the majority of its returns. Over half of the returns are generated in the first hour of the trading day, with 75% generated in the first and last hours of trading.

Dataset

Pairs trading relies on successfully identifying pairs of stocks which move together and profiting when the prices diverge and then converge. Historically this methodology was developed using daily data, but more recently funds have moved into the intra day domain. We apply the pairs trading methodology to a sample of UK listed stocks for the period January 1st to December 31st 2007. Data come from the London Stock Exchange Tick by Tick database. As they are the most liquid stocks, we limit our study to the FTSE100 index constituents. Following Brownlees and Gallo [2006] and Andersen et al [2006] we extend the trading day to 16:35:00 to ensure that delayed closing prices are captured. Cancelled, overnight trades and auction trades are removed from the database. Simultaneous ticks are replaced with the median value of the simultaneous observations (Brownlees and Gallo [2006]).

[1] See for example Gatev, Goetzmann and Rouwenhorst [2006].

A feature of tick data is the irregular temporal spacing between trades. Since the LSE Tick database reports every trade that occurs on the London Stock Exchange, the inter-trade duration, which is measured as the time between consecutive trades, can range from seconds to hours, depending on the stocks. To limit the number of time periods without any observations, we use a 60 minute return period interval for this study.[2]

Pairs Trading

For our pairs trading strategy we divide our sample into 30 overlapping subsample periods of 396 hour duration. Each subsample period begins 22 hours after the start of the previous period and is divided into a 264 hour formation period and a 132 hour trading period. The formation period is used to identify pairs of stocks which move together with the actual transactions occurring in the trading period.

To identify pairs of stocks which move together we rank every pair based on the sum of squared deviations of the normalised price series estimated in the formation period. The top 5 and top 20 pairs from the formation period are selected for the trading period. In the trading period pairs are opened when the normalised prices diverge by more than a multiple, j, of the historical standard deviation metric. For larger values of j the portfolio will trade less and capture larger divergence in price. Pairs are opened by going long the lower normalised priced security and short the higher normalised priced security. Pairs are closed when the normalized price series converge, or at the end of the trading period. Returns are calculated by going long one GBP in the lower priced equity and short one GBP in the higher priced equity.

[Exhibit 1 and 2]

To illustrate the methodology Exhibit 1 and Exhibit 2 display the normalized price series of both Barclays and Lloyds Banking Group during the formation period and trading period respectively which spans January 2, 2007 to March 12, 2007. Barclays and Lloyds Banking Group both operate in the financial sector. It is apparent from Exhibit 1 that the two stocks trade in a similar pattern. In Exhibit 2 over the trading period ($j=2$), we can see the pairs open four times. The first opening occurs at 10:00 am on February 20, 2007. The pairs remain open for 9 trading hours and close at 10:00 am on February 21, 2007. In this instance, Lloyds Banking Group is the winner, and the pairs converge when Barclays' normalized price increases over the 9 hour period. The pairs open and close three more times over the

[2] To ensure our main findings our not sensitive to this assumption we repeated that analyses using 15, 30 and 120 minute intervals. These results are available from the authors on request.

trading period. In the last occurrence, the spread of the normalized price series once again increases to a margin greater than two times the historical standard deviation but the trading period ends before the pairs can converge; so the pairs are closed on the last trading hour of the period.

We use a committed capital measure of excess returns, scaling returns by the number of pairs that are matched in the formation period. Since the overlapping 132 hour trading periods begin at an interval of 22 hours, we are left with six concurrent time series of pairs trading returns beginning February 19th. We then average these returns to calculate our final portfolio.

The Importance of Transaction Costs

Much research has linked the profitability of trading strategies to transaction costs. As the strategies involve far more trades than a traditional long only approach transaction costs can eliminate excess returns (see Knez and Ready [1996] for an example). In particular we would expect that the returns to high frequency trading are particularly sensitive.

[Exhibit 3]

For this study we specify two alternate metrics as our trade trigger, using both $j=2$ and $j=3$ x historical standard deviations. We can see in Exhibit 3 the effect of using these different values for j. Using a $j=2$ trigger the average price deviation when a trade is opened ranges from 174 to 229 basis points, whereas using a $j=3$ trigger, leads to an opening range of 261 to 344 basis points. When we include more pairs in our portfolio the mean deviation widens but the frequency of trades decreases.

Given we must trade a pair of stocks twice to capture this deviation it is clear that profitability will be closely linked to transaction costs. To investigate the importance of transaction costs we estimate the returns to the strategy using a range from 0 to 15 basis points for both the $j=2$ and $j=3$ triggers.

[Exhibit 4]

Exhibit 4 displays these results for the $j=2$ portfolio. The annualised mean return for zero transaction costs is 19.8% dropping to 5.5% for 15 basis points transaction costs. The worst daily returns is -0.9% and the best is +1.2%. The results are quite similar when we include more pairs in the portfolio, albeit with marginally lower returns and standard deviation. Looking at the skewness and kurtosis we can see that the distributional characteristics are close to normal.

[Exhibit 5]

Next in Exhibit 5 we report the results of a similar analysis, this time with $j=3$. With the wider spread on trade entry and less frequent trading these results are less sensitive to transaction costs with annual returns ranging from 15.2% to 7% for the top 5 pairs and 12.3% to 4.3% for the top 20 pairs. Minimums and maximums are similar to those reported in Exhibit 4.

Speed of Execution

The literature on negative serial correlation in intraday stock returns (see for example Engle and Russell [2006]) implies that pairs trading profitability may be quickly eroded by the passage of time. To test the sensitivity of the strategy to speed of execution we implement a simple restriction. Following a trading signal we wait one full interval before executing.

[Exhibit 6]

We can see from Exhibit 6 the dramatic effect this has on profitability. The mean returns decrease to about 4% per annum before transaction costs and negative net of transaction costs. Clearly, the profitability of the strategy is closely linked to speed of execution.

Risk Factors

We are also interested in understanding the exposure of the strategy to traditional market risk factors. To gauge the sensitivity of the returns of the strategy to market risk factors we estimate a simple five factor model including risk factors which proxy for broad equity market, size, value, momentum and reversal excess returns. The size factor is the total daily return of the Hoare-Govett Small Companies (HGSC) index minus the total daily return of the FTSE 100 Index while the value factor is simply the difference between the total daily returns of the MSCI UK value index and the MSCI UK growth index. The momentum factor is the difference between equally weighted portfolios formed from the highest and lowest 30% of returns over the past 250 days for all UK listed stocks. To construct the reversal factor deciles are formed on the prior twenty one days returns and excess returns are calculated as the difference between top three deciles returns minus the returns in the bottom three deciles.

[Exhibit 7]

The results from estimating this model are reported in Exhibit 7 for the Top 5 and Top 20, *j=3* portfolios.[3] There is little evidence of exposure to the factors. The Top 5 portfolio reversal coefficient is statistically significant, as is the Top 20 broad market coefficients but the explanatory power of the models is very weak. With adjusted R^2 of only 2% very little of the returns are explained by traditional risk factors.

Patterns in Returns

It is well known that there are various patterns in trading volume and the behaviour of security prices (see for example Admati and Pfleiderer [1988]). For example, over a trading day both volume and volatility tend to be U shaped with heavy trading, accompanied by volatility at the start and end of the day. Given this observed phenomenon we are interested in investigating whether there is a relationship between these patterns in the data and the returns to pairs trading.

[Exhibit 8 and 9]

Initially in Exhibit 8 and 9 we plot the cumulative returns to the portfolio by Month and also by day of the week. The monthly returns are interesting in that the UK did not seem to experience the negative returns to quantitative strategies caused by the unwinding of positions during the credit crisis observed in the US market (Khandani and Lo [2007]). In fact one of the highest returning months, August is when the losses were highest in US equity market neutral funds.

When we look at returns by day of the week in Exhibit 9 we can see that Monday is the highest returning day and Thursday the lowest. However statistical testing indicates zero significant difference in returns. The picture is quite different when we examine hour by hour returns. Exhibit 10 displays the cumulative returns for each hour of the trading day beginning 08.00 to 09.00 and finishing 15.00 to 16.35.

[Exhibit 10]

We can see that the vast majority of the returns are earned in the first hour of the day, followed by the final hour of the day. It would seem the increased volume and variance of price changes in the first and last hour of the day creates profitable for arbitrageurs.

[3] Results, available from the authors on request, are similar for the *j*=2 portfolio.

Summary and Implications

There is considerable evidence that relative value trading strategies, implemented with daily data, generate excess returns, which are not simply compensation for bearing risk. In this paper we examine the performance of a relative value strategy, pairs trading, using high frequency data, highlighting three key issues for the strategy. We first document the very high sensitivity of the returns to transaction costs, particularly for the trading trigger which results in more frequent trading. We further show that the returns are extremely sensitive to speed of execution, with all excess returns being eliminated if you wait one period to execute. Our factor model analysis indicates that the returns are not related to traditional risk factors. Finally we identify a clear pattern to the returns with the majority occurring in the first and last hour of trading.

The implications of these findings for hedge funds and proprietary trading desks are clear. The success of implementing an intra day pairs trading strategy will depend upon minimizing the level of transaction costs and, depending upon this level, striking a balance between the deviation trigger and the amount of trading. Likewise speed of execution is essential as the opportunity to capture divergence and convergence is fleeting. Further, outside the control of the fund will be the level of intra day price volatility and liquidity. With a short sample period we cannot draw to firm conclusions other to highlight that for our study the majority of returns occurred in the first hour of the day, where typically we would anticipate seeing the relatively largest volatility and liquidity.

References

Admati, Anat. R. and P. Pfleiderer, "A Theory of Intraday Patterns: Volume and Price Variability", Review of Financial Studies, Vol. 1, No. 1 (1988), pp. 3-40.

Andersen, T., Bollerslev T., Christoffersen, P., and F. Diebold, "Volatility and Correlation Forecasting", In Elliott, G., Granger, C. and A. Timmermann, eds., Handbook of Economic Forecasting, Amsterdam: North-Holland, 2006, pp. 777-878.

Brownlees, Christian T. and G.M. Gallo, "Financial Econometric Analysis at Ultra-High Frequency: Data Handling Concerns", Computational Statistics & Data Analysis, Vol. 51, No. 4 (2006), pp. 2232-2245.

Engle, Robert F. and J.R. Russell, "Analysis of High Frequency Data", In Ait Sahalia,Y. and L.P. Hansen, eds., Handbook of Financial Econometrics, Amsterdam: North-Holland, 2006.

Gatev, Evan, Goetzmann, W.N., and K.G. Rouwenhorst, "Pairs Trading: Performance of a Relative-Value Arbitrage Rule", Review of Financial Studies, Vol. 19, No. 3 (2006), pp. 797-827.

Khandani, Amir. E. and A.W. Lo, "What Happened to the Quants in August 2007?" Working paper, MIT, (2007).

Knez, Peter .J. and M.J. Ready, "Estimating the Profits from Trading Strategies", Review of Financial Studies, Vol. 9, No. 4 (1996), pp. 1121-1163.

Exhibit 1: Formation Period

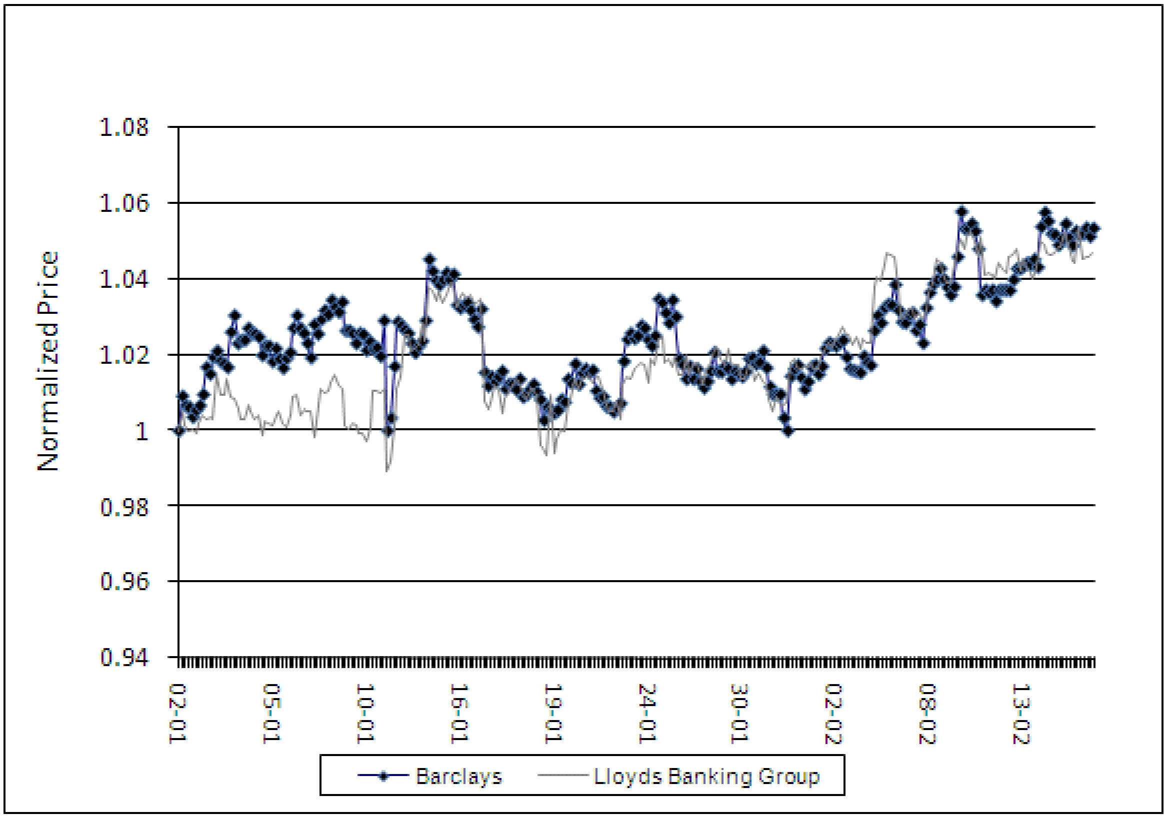

Exhibit 2: Trading Period

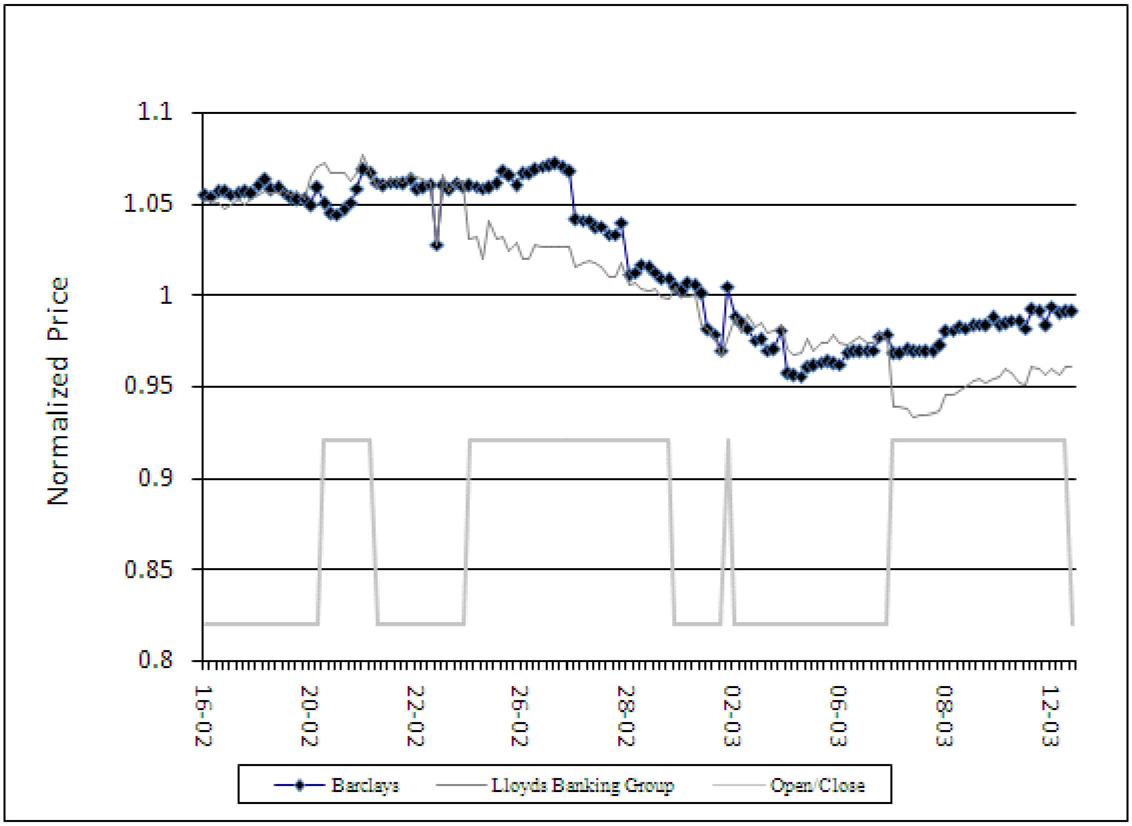

Exhibit 3: Trading Statistics

	Top 5	Top 20
2 Standard Deviations		
Average price deviation trigger for opening pairs	1.74%	2.29%
Average round trips per pair	3.84	3.49
Average number of pairs traded in trading period	4.91	19.25
3 Standard Deviations		
Average price deviation trigger for opening pairs	2.61%	3.44%
Average round trips per pair	2.46	2.36
Average number of pairs traded in trading period	4.61	17.70

Notes: This table reports average statistics for the Top 5 and Top 20, 2 x and 3 x standard deviation triggers.

Exhibit 4: Summary Statistics 2 Standard Deviation Trigger

		Transaction Costs		
		5 bp	10 bp	15 bp
Panel A: Top 5 Pairs				
Annual Mean	19.801	15.014	10.229	5.444
Annual Std Dev	5.040	5.007	4.983	4.971
Skewness	0.233	0.190	0.146	0.102
Kurtosis	4.054	4.007	3.959	3.910
Minimum	-0.009	-0.009	-0.009	-0.009
Maximum	0.012	0.011	0.011	0.010
Panel B: Top 20 Pairs				
Annual Mean	15.029	10.696	6.364	2.032
Annual Std Dev	3.431	3.407	3.388	3.375
Skewness	0.279	0.259	0.240	0.219
Kurtosis	4.039	4.062	4.082	4.098
Minimum	-0.007	-0.007	-0.007	-0.007
Maximum	0.009	0.009	0.008	0.008

Notes: This table reports key statistics for the Top 5 (Panel A) and Top 20 (Panel B) portfolios using the 2 x standard deviation trigger at different levels of transaction costs.

Exhibit 5: Summary Statistics 3 Standard Deviation Trigger

		Transaction Costs		
		5 bp	10 bp	15 bp
Panel A: Top 5 Pairs				
Annual Mean	15.240	12.498	9.755	7.013
Annual Std Dev	4.609	4.575	4.546	4.522
Skewness	0.010	-0.009	-0.029	-0.049
Kurtosis	4.199	4.182	4.161	4.139
Minimum	-0.009	-0.009	-0.009	-0.009
Maximum	0.009	0.009	0.009	0.009
Panel B: Top 20 Pairs				
Annual Mean	12.348	9.686	7.023	4.361
Annual Std Dev	3.159	3.142	3.129	3.120
Skewness	0.120	0.124	0.129	0.134
Kurtosis	4.808	4.830	4.840	4.840
Minimum	-0.008	-0.008	-0.008	-0.008
Maximum	0.007	0.007	0.007	0.007

Notes: This table reports key statistics for the Top 5 (Panel A) and Top 20 (Panel B) portfolios using the 3 x standard deviation trigger at different levels of transaction costs.

Exhibit 6: Summary Statistics of Wait 1 Period 3 Standard Deviation Trigger

		Transaction Costs		
		5 bp	10 bp	15 bp
Panel A: Top 5 Pairs				
Annual Mean	3.645	0.904	-1.838	-4.580
Annual Std Dev	4.328	4.318	4.314	4.315
Skewness	-0.262	-0.264	-0.265	-0.265
Kurtosis	5.542	5.499	5.442	5.373
Minimum	-0.012	-0.012	-0.012	-0.012
Maximum	0.008	0.008	0.008	0.007
Panel B: Top 20 Pairs				
Annual Mean	5.098	3.773	0.873	-2.027
Annual Std Dev	3.069	3.779	3.772	3.769
Skewness	0.630	0.228	0.213	0.197
Kurtosis	7.830	6.427	6.354	6.265
Minimum	-0.007	-0.010	-0.010	-0.010
Maximum	0.011	0.011	0.010	0.010

Notes: This table reports key statistics for the Top 5 (Panel A) and Top 20 (Panel B) portfolios using the 3 x standard deviation trigger at different levels of transaction costs with the wait one period to execute restriction.

Exhibit 7: Five Factor Model 3 x Historical Standard Deviation triggers

	Top 5	Top 20
Intercept	0.044	0.054
	(1.97)**	(3.48) ***
Market	0.00	0.04
	(0.18)	(2.56) **
Size	0.00	0.03
	(-0.03)	(1.10)
Value	-0.05	0.00
	(-1.33)	(-0.15)
Momentum	0.01	0.00
	(0.13)	(0.01)
Reversal	0.15	-0.02
	(2.23) **	(-0.53)
Adjusted R^2	0.02	0.02
No Serial Correlation	0.13	0.13

Notes: This table reports the estimated coefficients of the five-factor regression for the Top 5 and Top 20 pairs portfolios for the period 16th February to 31st December 2007. T-statistics are reported in parenthesis. ***, **, * indicate significance at the 1%, 5% and 10% level respectively.

Exhibit 8: Returns by Month

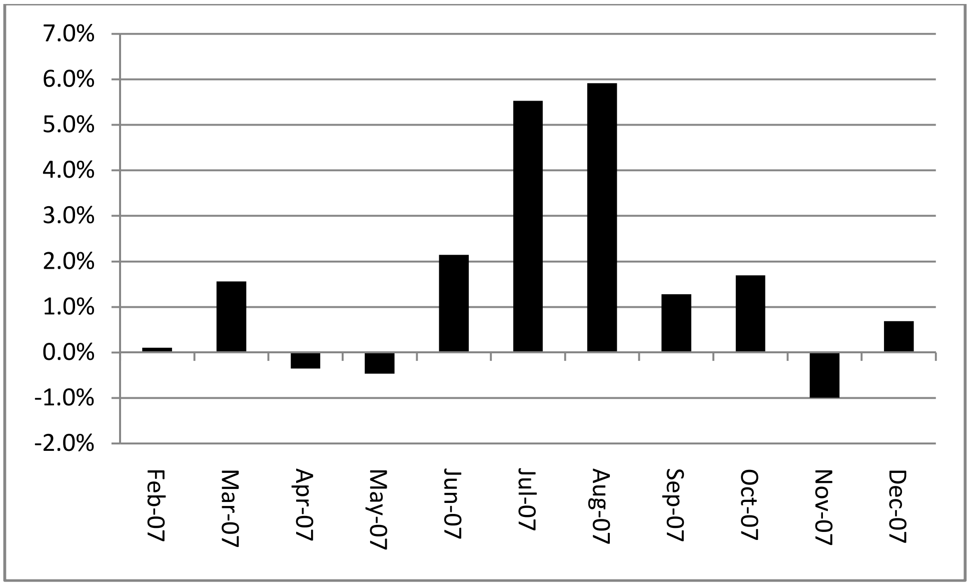

Exhibit 9: Returns by Day of the Week

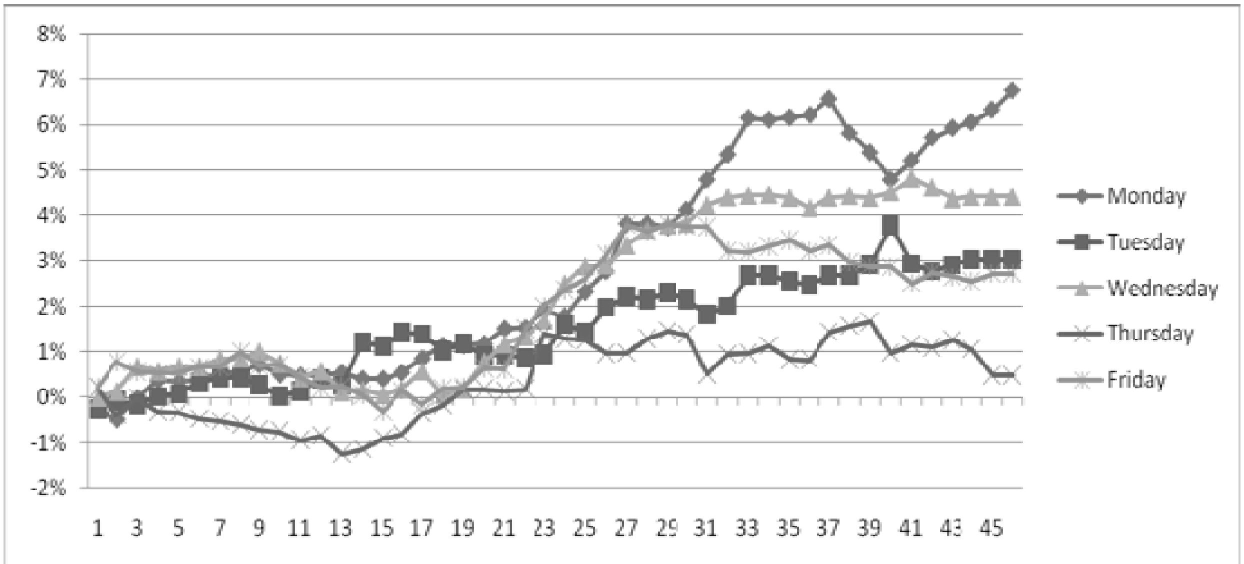

Exhibit 10: Returns by Hour of the Day