Stock market predictability: Non-synchronous trading or inefficient markets? Evidence from the National Stock Exchange of India

Camilleri, Silvio John and Green, Christopher J.

2014

This paper was published as follows:
Camilleri, S.J. and C.J. Green, (2014),"Stock market predictability: Non-synchronous trading or inefficient markets? Evidence from the National Stock Exchange of India", *Studies in Economics and Finance*, 31 (4), pp. 354 - 370.

Publication may be accessed at:
http://dx.doi.org/10.1108/SEF-06-2012-0070

STOCK MARKET PREDICTABILITY: NON-SYNCHRONOUS TRADING OR INEFFICIENT MARKETS? EVIDENCE FROM THE NATIONAL STOCK EXCHANGE OF INDIA

Silvio John Camilleri and Christopher J. Green

Camilleri: Banking & Finance Department, University of Malta
Green: Department of Economics, Loughborough University

Camilleri, S.J. and C.J. Green, (2014),"Stock market predictability: Non-synchronous trading or inefficient markets? Evidence from the National Stock Exchange of India", *Studies in Economics and Finance*, 31 (4), pp. 354 - 370.

ABSTRACT

Purpose
The main objective of this study is to obtain new empirical evidence on non-synchronous trading effects through modelling the predictability of market indices.

Design / Methodology / Approach
We test for lead-lag effects between the Indian Nifty and Nifty Junior indices using Pesaran-Timmermann tests and Granger-Causality. We then propose a simple test on overnight returns, in order to infer whether the observed predictability is mainly attributable to non-synchronous trading or some form of inefficiency.

Findings
The evidence suggests that non-synchronous trading is a better explanation for the observed predictability in the Indian stock market.

Research limitations / implications
The indication that non-synchronous trading effects become more pronounced in high-frequency data, suggests that prior studies using daily data may underestimate the impacts of non-synchronicity.

Originality / value
The originality of the paper rests on various important contributions: (a) we look at overnight returns to infer whether predictability is more attributable to non-synchronous trading or to some form of inefficiency, (b) we investigate the impacts of non-synchronicity in terms of lead-lag effects rather than serial correlation, and (c) we use high-frequency data which gauges the impacts of non-synchronicity during less active parts of the trading day.

JEL Classification: G12, G14
Keywords: Non-Synchronous Trading, Stock Markets, National Stock Exchange of India.

Camilleri, S.J. and C.J. Green, (2014),"Stock market predictability: Non-synchronous trading or inefficient markets? Evidence from the National Stock Exchange of India", *Studies in Economics and Finance*, 31 (4), pp. 354 - 370.

1. Introduction

Security pricing theory usually assumes that trading opportunities occur and can be exploited continuously, whereas empirical research invariably uses data which are sampled at discrete intervals. This distinction is particularly important when markets are characterized by thin and non-synchronous trading: possibly prolonged periods when securities trade intermittently and at different intervals. As new information becomes available, the last transaction price of a less-frequently traded share may not reflect the true, full-information value of the firm. At each point in time, the cross-section of last trade prices reflects different overlapping information, some of which may be out-of-date. A sequence of such cross-sections conveys the impression that particular stocks react more slowly than others to new information. However, this is exactly the pattern which may also be expected if the market is not information-efficient and traders react slowly to news. In contrast, if the cause is non-synchronous trading, these cross-sectional characteristics are due to institutionally-determined variations in the time interval since the last trade was made, and do not reflect inefficiencies in the market.

Non-synchronous trading also has an important impact on the calculation of any market index which is based on the prices of the most recent transactions in its component shares. Index values will be based on a cross-section of prices which may be anything from a few seconds to several days out-of-date depending on when the last trade occurred. Under these circumstances, the index components and therefore the index as a whole will not continuously reflect all currently available information (Fisher, 1966). This induces bias in a range of inferences concerning the time series properties of security returns[1]. In addition, security return series will be autocorrelated even if prices fully reflect all available information at the time of a trade (Roll, 1984). These too are characteristics which are consistent equally with market inefficiency or non-synchronicity. However, as Lo and MacKinlay (1990a) emphasize, many researchers avoid modelling non-synchronous trading explicitly, because the assumption of homogeneous sampling intervals is far more analytically tractable.

From the point of view of market practitioners, it is particularly important to understand the nature of the information contained in share returns. Much of the original impetus underlying the development of high-frequency and programme trading was derived from the relentless search for market inefficiencies in the form of mispricing. According to the efficient markets tradition, trading which profits from such mispricing necessarily leads to an improvement in market efficiency and to prices that more fully reflect all available information (Fama, 1991). More recently, Jarrow and Protter (2011) (and others) have questioned this widely-accepted argument and shown that high-frequency trading can itself generate mispricing especially in response to particular kinds of market signal. This makes it more important than ever that traders who are seeking to interpret market information must be able to distinguish between signals from an inefficient market on the one hand and noise from market microstructure such as non-synchronous trading, on the other.

The main objective of this study is to glean new empirical evidence on the information about non-synchronous trading and market efficiency provided by the predictability of market indices. Our approach to this issue is novel in four respects. First, we use new methods to identify the relationship between predictability and non-synchronicity. Previous studies of non-synchronous trading concentrated on the autocorrelation of returns, whereas we focus on lead-lag relationships between two indices containing more and less liquid shares respectively. We anticipate that the more liquid index will adjust to new market information more rapidly and therefore it will tend to lead the less liquid index. Predictability is tested using two separate methods: Pesaran-Timmermann tests and Granger-Causality tests using Vector Autoregressions (VARs). However, general lead-lag effects could be consistent either with non-synchronous trading or with some form of market inefficiency. Our second contribution therefore, is that we propose a new and simple test, based on price changes at overnight trading breaks, to infer which is more important in explaining the lead-lag relationships which we observe: non-synchronous trading or delayed price adjustment.

The third contribution of this paper is that we employ a high quality, high frequency dataset in combination with daily data. Trading activity typically varies systematically through the trading day (Wood *et. al.*, 1985); it tends to abate in the middle of the day and peak towards the end. This is also true of the National Stock Exchange of India (NSE), which is the object of our study (Shah and Sivakumar, 2000). It follows that non-synchronous trading effects are likely to be amplified in the middle of the day, and be less significant at the end of the day. Evidently, such effects will be difficult to detect using data sets based exclusively on closing prices where there is a trading peak. We would therefore argue that intra-day data is essential to obtain satisfactory evidence of the effects of non-synchronicity.

Our fourth contribution is that we use data from an emerging market: the NSE. Emerging markets tend to exhibit: more thin trading, higher serial correlation and slower adjustment of prices to news than developed ones (Bekaert and Harvey, 2002). The NSE includes a substantial proportion of less liquid securities, typical of an emerging market, and therefore provides an interesting new setting to gauge non-synchronous trading effects in daily and higher frequency data.

The rest of the paper is structured as follows: Section 2 reviews prior literature; section 3 describes the data sets; sections 4-5 test for predictability using Pesaran-Timmermann tests and Granger-Causality; section 6 uses the information in overnight trading breaks to investigate whether the observed predictability is more attributable to non-synchronous trading or delayed adjustments of traders' valuations; and section 7 concludes.

2. Non-Synchronous Trading: Research Background

Non-synchronous trading induces certain regularities in stock returns, especially serial correlation (Scholes and Williams, 1977; Cohen *et. al.*, 1979); and market indices tend to exhibit more serial correlation than individual stocks (Fisher, 1966). However, serial correlation may also be caused by delayed price adjustments and these could occur for several reasons (Atchison *et. al.*, 1987). First, when traders pick off limit orders whose prices go stale as new information becomes available, transactions occur at outdated prices but could still be consistent with market efficiency, since efficiency does not require *all* participants to price new information instantaneously. Second, if traders monitor less liquid stocks less intensively than more liquid ones, new information about the former stocks will generally take longer to be reflected in the price. Third, serial correlation can be induced as positive feedback traders buy stocks when prices rise and *vice-versa*.

The evidence suggests that non-synchronous trading does increase serial correlation, but also that serial correlation cannot be explained wholly by non-synchronous trading. Studies of US data include: Lo and MacKinlay (1990b), Boudoukh *et. al.*, (1994), Kadlec and Patterson (1999), and Li and Yung (2006). For the UK, Clare *et. al*. (2002) investigated samples of securities using monthly data, and noted that stocks which trade at the end of the sampling interval still exhibited serial correlation, emphasizing that non-trading cannot fully account for serial correlation.

The foregoing studies focus on the serial correlation structure of the return data. Given that changes in expectations may take longer to show up in share price fluctuations if securities trade infrequently, non-synchronous trading may also result in lead-lag effects between security prices and between indices. This induces predictability, although not necessarily profitable trading opportunities, as noted by Day and Wang (2002) in a study of the Dow Jones Industrial Average. Conrad *et. al.*, (1991) used GARCH models to study the transmission of volatility between different size-sorted portfolios of US stocks and found that volatility in larger stocks affects the volatility of smaller stocks.

Few papers seek to address directly the relationship between lead-lag effects and non-synchronous trading. Using weekly and monthly data, Chiao *et. al.*, (2004) reported a contemporaneous relationship between Taiwanese stocks of different market capitalisation, implying that smaller stocks do not take longer to adjust than larger ones. Poshakwale and Theobald (2004) analysed daily data from the NSE (India) and the Bombay Stock Exchange (BSE) and, in contrast to Chiao *et. al.* (2004), found traces of lead-lag effects from larger to smaller stocks. About 50% of the predictability was attributable to non-synchronous trading, whilst the rest was caused by a mixture of non-synchronous trading and different adjustment speeds. However, the former two studies do not use intra-day data (as we do in this paper) and given that non-synchronous trading is a short-term phenomenon, its effects might not be readily detectable using daily or lower-frequency data.

AlKhazali (2011) investigated the issue of thin trading and (perceived) inefficiencies in the context of the weekly index data from six Gulf Cooperation Council countries. Testing the Random Walk Hypothesis using different variance ratio tests, the null of a random walk was rejected in the original market data, yet it could no longer be rejected after adjusting for thin trading.

3. Empirical Setting and Data Characteristics

The NSE was established in 1994 and is one of two major Indian exchanges, together with the BSE. Equities traded on NSE increased from around 640 in 1999 to over 1600 in 2012 with most major stocks traded on both the NSE and the BSE. Daily volumes have risen sharply: 400,000 transactions per day were common in 1999 increasing to around 5 million in 2012.

The data for this study were extracted from the NSE's historical trades data CDs. The daily data set consists of closing observations of the NSE Nifty and Nifty Junior indices over the period 1st January 1999 to 31st December 2012: a total of 3500 observations. The Nifty consists of the top fifty traded stocks, whilst Nifty Junior is composed of the next fifty most frequently traded stocks. We also utilise intra-day data, consisting of the same two indices sampled at one minute intervals, over the period from 15th June 1999 to 25th June 1999. This consists of nine continuous trading days starting at 10am and ending at 3.30pm: a total of 2970 observations. The choice of the earlier part of the period to sample the high-frequency data was deliberate: in view of the increase in trading volumes throughout the years, one may expect to obtain clearer evidence of non-synchronicity using earlier data.

Given that non-synchronous trading effects depend on the level of liquidity, it is important to assess the relative liquidity of the sampled indices. When comparing the trading frequencies of the underlying shares, it emerges that the Nifty is unambiguously more liquid than the Nifty Junior. The average waiting time between trades across Nifty (Nifty Junior) stocks is around 16 (21) seconds (Table 1, panel A). The average waiting time of the ten least frequently traded shares is appreciably longer (Table 1, panel B).

Table 1 about here

Our basic hypothesis is that the less liquid Nifty Junior returns will be partly predictable using the Nifty returns. Conversely, the Nifty Junior should have little predictive content for the Nifty, particularly as it is likely that more active foreign traders concentrate their holdings in the more liquid stocks and thus contribute more actively to Nifty's efficiency [2]. Predictability may be the outcome of non-synchronous trading or of the actual delayed adjustments of traders' expectations. Therefore our analysis aims to infer the extent that predictability can be attributable to each of these possible causes. We also expect non-synchronous trading effects to be more pronounced in the high-frequency data set given that the daily observations coincide with the typical trading surge at the closing.

4. Pesaran-Timmermann Tests

We now turn to the first predictability investigation using tests proposed by Pesaran and Timmermann (1992) which measure the dependence between two time series in terms of whether they fluctuate in the same direction. The tests thus consider the direction of changes,

Camilleri, S.J. and C.J. Green, (2014),"Stock market predictability: Non-synchronous trading or inefficient markets? Evidence from the National Stock Exchange of India", *Studies in Economics and Finance*, 31 (4), pp. 354 - 370.

sidelining their magnitude. The null hypothesis is that the two series are independent, and the test statistic is asymptotically normally distributed. The test statistic (S) for assessing the relationship between two variables x_t and y_t is given by:

$$S = \frac{\hat{P} - \hat{P}_*}{\sqrt{\{\hat{V}(\hat{P}) - \hat{V}(\hat{P}_*)\}}} \xrightarrow{a} N(0,1) \qquad \ldots(1)$$

where:

$$\hat{P} = \frac{1}{n}\sum_{t=1}^{n} Sign(y_t x_t) \qquad \ldots(2)$$

$$\hat{P}_y = \frac{1}{n}\sum_{t=1}^{n} Sign(y_t) \qquad \ldots(3)$$

$$\hat{P}_x = \frac{1}{n}\sum_{t=1}^{n} Sign(x_t) \qquad \ldots(4)$$

$$\hat{P}_* = \hat{P}_y\hat{P}_x + (1-\hat{P}_y)(1-\hat{P}_x) \qquad \ldots(5)$$

$$\hat{V}(\hat{P}) = \frac{1}{n}\hat{P}_*(1-\hat{P}_*) \qquad \ldots(6)$$

$$\hat{V}(\hat{P}_*) = \left[\frac{1}{n}(2\hat{P}_y - 1)^2\hat{P}_x(1-\hat{P}_x)\right] + \left[\frac{1}{n}(2\hat{P}_x - 1)^2\hat{P}_y(1-\hat{P}_y)\right] + \left[\frac{4}{n^2}\hat{P}_y\hat{P}_x(1-\hat{P}_y)(1-\hat{P}_x)\right] \qquad \ldots(7)$$

The function *Sign(z_t)* takes a value of 1 when z_t is positive and zero otherwise. Thus, $\hat{P}$ measures the proportion of occurrences where both time series fluctuated in the same direction: $\hat{P}$ varies between unity and zero. $\hat{P}_x$ and $\hat{P}_y$ are the proportions of negative and positive changes in each separate series; these are used with $\hat{P}_*$, $\hat{V}(\hat{P})$ and $\hat{V}(\hat{P}_*)$ to rescale $\hat{P}$ and construct a normal distribution. This test is applied to the log returns on the two indices [3]. Since we are particularly interested in the lead-lag effects, we apply Pesaran-Timmermann tests on the relationships between x_t and y_{t-i} and between x_{t-i} and y_t.

Table 2 about here

Table 2 shows the Pesaran-Timmermann statistic up to 25 lags. As we would expect, there is strong evidence that the indices move in the same direction contemporaneously, at both the low and high data frequency. Examining the relationship between the current and lagged returns of the daily data, we see that the first lag is highly significant for both indices, but none of the other lags is highly significant [4]. This suggests that the difference in liquidity between the two indices does not result in significantly different predictability effects, also implying that the predictability effect at the first lag is not the result of non-synchronicity. One possible explanation might be "runs" in the data, since these can be a normal feature of any series that may be classified as a random walk with drift.

In the high-frequency data, we see that for both indices, the first and second lags are highly significant. This bidirectional predictability may again be attributed to runs in the data. An alternative explanation is that it may be unrealistic to expect abrupt price changes in the high-frequency data given that when new information becomes available, stale limit orders are "picked off", resulting in "outdated" transaction prices. However, the Nifty remains highly significant in predicting the direction of change of the Nifty Junior for four further lags. This is consistent with a larger non-synchronous trading effect in the Nifty Junior than the Nifty. The difference between the results obtained at low-frequency and high-frequency supports the argument that non-synchronous trading effects show up more clearly in higher frequency data.

Overall, these results indicate that the Nifty leads the Nifty Junior in high frequency data, although with some feedback at the shortest lags, while the two indices tend to move nearly contemporaneously in lower frequency daily data. This is in line with our prior expectations: predictability runs stronger from Nifty to Nifty Junior, and part of it is the result of non-synchronicity which becomes more pronounced in high frequency data.

5. Granger-Causality Tests: Theory

If shocks in one time series lead to movements in another time series, then the former series is said to Granger-cause the latter (Granger, 1969). To examine causal orderings between the Nifty and Nifty Junior, we set up the bivariate Vector Autoregression (VAR):

$$x_t = \sum_{i=1}^{n} \alpha_{i1} x_{t-i} + \sum_{i=1}^{n} \beta_{i1} y_{t-i} + u_{1t} \qquad \ldots(8)$$

$$y_t = \sum_{i=1}^{n} \alpha_{i2} x_{t-i} + \sum_{i=1}^{n} \beta_{i2} y_{t-i} + u_{2t} \qquad \ldots(9)$$

where: x_t and y_t are the Nifty and Nifty Junior returns respectively; $\alpha_{i,j}$, $\beta_{i,j}$ are coefficients; and $u_{i,t}$ are residuals. The significance of each block of parameters ($\alpha_{i,1}$, $\beta_{i,1}$, $\alpha_{i,2}$, $\beta_{i,2}$) can be tested using F or χ^2 tests. The $\alpha_{i,1}$ and $\beta_{i,2}$ describe the parts of x_t and y_t respectively explained by a univariate time series model; the other parameters describe the causal orderings: if the $\beta_{i,1}$ are significant then y_t is said to cause x_t; likewise, if the $\alpha_{i,2}$ are significant then x_t is said to cause y_t.

The relationship between Granger Causality and market efficiency has been hotly debated. It is now generally recognised that evidence of causality in stock returns may reflect other factors than market inefficiencies (Fama, 1991), including in the present context the effects of non-synchronous trading. However, these effects should disappear relatively rapidly. Therefore, in the daily dataset, if the Nifty and Nifty Junior are priced efficiently and there are no non-synchronous trading effects, we might expect their returns to be contemporaneously correlated but to exhibit few causal orderings. However, for the high-frequency data, greater evidence of causality may be expected, especially if there are non-synchronous trading effects arising from variations in liquidity across securities. In such case some of the prices used to calculate index values may be out-of-date, giving rise to Granger causality, particularly from the more liquid Nifty to the less liquid Nifty Junior [5].

5.1 Granger-Causality Tests: Daily Data

A preliminary 24-order VAR was estimated on both returns series so as to select the optimal lag-order of the model. The Akaike Information Criterion (AIC) selected a VAR(18) whilst the Schwarz Bayesian Criterion (SBC) selected a VAR(1). We estimated a VAR(1) and a VAR(18), and then chose the former on the basis of higher system log likelihood ratio, higher F-statistics, and more particularly parsimony. Checks on regression residuals for possible clustering of large errors on specific days of the week (especially Mondays), suggested that there was no evidence of such tendencies, in line with Choudhry (2000).

Table 3 shows summary statistics for the VAR(1) model. In both equations, the coefficient denoting any lead-lag relationship between indices is not significant at the 95% level. For completeness, Granger non-causality tests were also conducted and the null of non-causality could not be rejected. However, a test for contemporaneous covariance between the two series unambiguously rejected the null of no covariance.

Tables 3 and 4 about here

Evidently, these results are in line with the two key inferences from the Pesaran-Timmermann tests: there is only a weak lead-lag relationship (if at all) between the indices at daily intervals, and it might be more accurate to postulate a contemporaneous relationship at this frequency. This suggests that differing liquidity levels between indices do not lead to non-synchronous trading effects of sufficient length to be gauged through significant lead-lag effects in a daily data set, which is consistent with one of our basic hypotheses.

Camilleri, S.J. and C.J. Green, (2014),"Stock market predictability: Non-synchronous trading or inefficient markets? Evidence from the National Stock Exchange of India", *Studies in Economics and Finance*, 31 (4), pp. 354 - 370.

5.2 Granger-Causality Tests: One-Minute Data

A preliminary 24-order VAR was estimated on the high-frequency data, where the AIC and SBC selected a VAR(9) and a VAR(3) respectively. Log-likelihood ratio statistics rejected all VAR orders which were less than 7, and therefore a VAR(9) was selected. The diagnostics for the initial VAR(9) showed evidence of heteroskedasticity; large errors tended to occur at approximately equally-spaced intervals coinciding with the opening of the trading day, particularly at the first two observations. This is not necessarily surprising as a higher amount of news is priced during the first observation following the overnight interval. This implies that the first two minutes' prices probably provide a particularly weak forecast for the subsequent minutes' returns. To model this effect, we introduced a dummy variable with a value of unity for the first two observations of each trading day and zero thereafter. The dummy is highly significant in the Nifty equation and it reduced the non-normality and ARCH tendencies in the residuals (Table 4).

The hypothesis tests show that the null of non-causality can be confidently rejected for both series as can the null of no contemporaneous covariance. These results support the Pesaran-Timmermann tests at one-minute intervals: firm evidence that Nifty leads Nifty Junior, and that the two indices are contemporaneously correlated. Causality runs more strongly from Nifty to Nifty Junior as we would predict, with the feedback from the Nifty Junior being significant but smaller in magnitude and relatively short-lived.

The evidence so far is consistent with the argument that market-wide information is first reflected in the Nifty index, and some minutes later in the Nifty Junior, so that we obtain lead-lag effects in high frequency data. Due to the trading surge at the end of the day, closing observations are computed through reasonably recent prices and therefore any predictability effects cease to be significant at the 95% confidence level when using daily data. The feedback in the one-minute data from the Nifty Junior to Nifty is unlikely to be the result of non-synchronous trading, given that the Nifty is the more liquid index. Even so, this is in line with our prior expectations that non-synchronicity results in predictability, but not *all* predictability is caused by non-synchronicity.

6. The Information in Trading Breaks

We now analyse how trading break and post-trading break returns can provide information about whether delayed price adjustments in the data mainly emanate from traders' delays in adjusting their expectations ("inefficiency") or whether they are more attributable to non-synchronicity. Here, we concentrate on overnight breaks: the period from the cessation of trading activity at the end of the day until the subsequent morning (or until after the weekend). During an overnight trading break, market participants have enough time to adjust their judgements regarding fundamental values of firms, and any outdated limit orders will lapse or be cancelled. This implies that trades occurring immediately after a trading break will reflect the underlying value of shares and we may rule out "inefficient" delayed price adjustments by traders. Therefore, if the lead-lag effects between indices persist in the post-trading break data, they must be due mainly to non-synchronous trading rather than mispriced trades. Non-synchronous trading can coexist with trading breaks as an infrequently traded stock may trade after a longer delay immediately following the break. This produces a delayed adjustment of market price data, but one which is not attributable to delayed adjustment in expectations.

Our high-frequency VARs (Table 4) indicate that the first three and the sixth (one-minute) Nifty lags are significant in determining the value of the Nifty Junior. This suggests that it takes around six minutes for sufficient transactions to take place in the less liquid stocks, to achieve complete adjustment to news, irrespective of whether this is due to lagged adjustment of expectations or non-synchronous trading. Therefore we look at the Nifty Junior initial returns during the first six minutes of the day ($IR(M)_{t+1}$) and compare them to the Nifty overnight returns ($OR(N)_{t\rightarrow t+1}$) and the Nifty Junior overnight returns ($OR(M)_{t\rightarrow t+1}$). If the observed lead-lag effect up to six minutes consists of non-synchronous trading, we may expect predictability to persist following the trading break. Thus, if $OR(N)_{t\rightarrow t+1}$ is correlated with

$IR(M)_{t+1}$, this may be taken as an indication of non-synchronous trading. Conversely, if the observed predictability is attributable to inefficiency, we may expect the lead-lag effect to disappear during the overnight break. Provided traders have enough time to process news during this break, the overnight price adjustment in the Nifty Junior index should show up contemporaneously with that of the Nifty. Thus, if the $OR(N)_{t\to t+1}$ is correlated with the $OR(M)_{t\to t+1}$ (rather than $IR(M)_{t+1}$), this may be taken as an indication that predictability may be explained by some form of inefficiency.

We use two tests for the correlation between indices around the overnight breaks. The first is the simple correlation coefficient between the respective series (Table 5 Panel A). This reveals that $OR(N)_{t\to t+1}$ is more highly correlated with $IR(M)_{t+1}$ than with $OR(M)_{t\to t+1}$ suggesting that non-synchronous trading is the more important cause of observed predictability.

The second test involves two OLS regressions:

$$OR(M)_{t\to t+1} = \alpha_1 + \beta_1 OR(N)_{t\to t+1} + \varepsilon_1 \qquad \ldots(10)$$

$$IR(M)_{t+1} = \alpha_2 + \beta_2 OR(N)_{t\to t+1} + \varepsilon_2 \qquad \ldots(11)$$

where α_i and β_i are estimated coefficients, and ε_i is an error term. For this test, the sample was split into two to accommodate missing data, but the results are qualitatively the same across the two sub-samples (Table 5 Panel B) [6]. They again show that $OR(N)_{t\to t+1}$ is more closely related to $IR(M)_{t+1}$ than to $OR(M)_{t\to t+1}$. The key finding here is that both tests suggest that the predictability between the two indices persists across the overnight breaks. Such predictability cannot reasonably be attributed to traders delaying their adjustment of expectations, since they have ample time to do so during this break. Therefore, we can conclude that the causal relationship from the Nifty to Nifty Junior in the high-frequency data is at least in part attributable to non-synchronous trading effects.

Camilleri, S.J. and C.J. Green, (2014),"Stock market predictability: Non-synchronous trading or inefficient markets? Evidence from the National Stock Exchange of India", *Studies in Economics and Finance*, 31 (4), pp. 354 - 370.

Table 5 about here

The weaker relationship between $OR(N)_{t\rightarrow t+1}$ and $OR(M)_{t\rightarrow t+1}$ points at a contemporaneous co-movement in the indices during the overnight break, possibly as traders "catch up" with the day's information during this period. However, this effect is surprisingly weak, suggesting that the causal relationship from the Nifty to Nifty Junior in the high-frequency data is unlikely to be attributable to informational inefficiencies.

Finally, the causal relationship running in the reverse direction from Nifty Junior to Nifty clearly cannot be attributed to non-synchronous trading. This may be due in part to possible information spillover effects amongst stocks. However, our results are consistent with those of Atchison, *et. al.* (1987) and Lo and MacKinlay (1990a) who found that indices exhibit higher predictability than that which may be expected exclusively from non-synchronicity.

7. Conclusion

Market efficiency is one of the main research questions in finance, and a major symptom of inefficiency is predictable stock returns. However, predictability may be consistent either with some form of inefficiency, or with non-synchronous trading effects deriving from institutional characteristics of the securities and the trading setup. Since non-synchronicity induces some degree of predictability, it may lead to flawed inferences regarding market efficiency (in line with AlKhazali, 2011). From a practitioner's point of view it may lead to trading decisions which are individually incorrect or, more seriously, to mis-programming of electronic trading routines in such a way that they exacerbate rather than dampen market inefficiencies (Jarrow and Protter, 2011).

In this paper we used a new methodology to detect the difference between non-synchronous trading and market inefficiency insofar as they affect the predictability of stock returns. We investigated the causal orderings between the more liquid Nifty and less liquid Nifty Junior and their implications for non-synchronous trading and market efficiency. The basic finding is that

Nifty leads Nifty Junior, particularly in the high-frequency dataset. This tends to confirm our hypothesis that non-synchronous trading effects may be detected through the pattern of lead-lag relationships in returns. We also showed that non-synchronous trading effects are more pronounced in high-frequency data, which is consistent with the findings of Papachristou (1999).

In addition, we proposed a simple methodology for discriminating between two main sources of predictability in high-frequency data: non-synchronous trading or inefficiencies attributable to delayed price adjustments. This involved investigating the predictability of returns following the overnight trading breaks where traders have time to update their information. We find that Nifty Junior returns in the first six minutes on the NSE are remarkably predictable using the overnight Nifty returns, and this strongly suggests that the observed predictability is attributable mainly to non-synchronicity. However, the finding that the less liquid Nifty Junior has some, though weaker, power to predict the Nifty, points at some elements of inefficiency as well. Our finding that predictability emanates mainly from non-synchronous trading, reinforces the argument that a more useful criterion for market efficiency is the absence of profitable trading opportunities rather than the absence of predictability (Buckle *et. al.*, 1999). It also underlines the concerns expressed in the recent literature that programme trading may feed off apparent market signals (such as those generated by non-synchronous trading) to create inefficiencies where none previously existed, and increasing the probability of recurrent "flash crashes" (Sornette and Von der Becke, 2011).

This investigation suggests that previous studies may need reinterpretation, insofar as they were mostly based on daily data. The typical increase in trading activity at the end of the day tends to diminish non-synchronous trading effects, and leads researchers to underestimate the impact of non-synchronicity. Furthermore, since trading activity varies throughout the trading day, it would also be interesting to investigate how far non-synchronous trading effects become more pronounced during the middle of the day when trading tends to abate.

References

AlKhazali, O. (2011), "Does Infrequent Trading Make a Difference on Stock Market Efficiency? Evidence from the Gulf Cooperation Council (GCC) Countries", *Studies in Economics and Finance*, Vol. 28 No. 2, pp. 96-110.

Atchison, M.D., Butler, K.C. and Simonds, R.R. (1987), "Nonsynchronous Security Trading and Market Index Autocorrelation", *The Journal of Finance*, Vol. 42, pp. 111-118.

Bekaert, G. and Harvey, C.R. (2002), "Research in Emerging Markets: Looking to the Future", *Emerging Markets Review*, Vol. 3, pp. 429-448.

Boudoukh, J., Richardson, M. and Whitelaw, R. (1994), "A Tale of Three Schools: Insights on Autocorrelations of Short-Horizon Returns", *Review of Financial Studies*, Vol. 7, pp. 539-573.

Buckle, M.J., Clare, A.D. and Thomas, S.H. (1999), "Developing a Trading Rule from the FTSE-100 Stock Index Futures Contract: Evidence in Support of EMH", *Journal of Business Finance and Accounting*, Vol. 26, pp. 249-260.

Chiao, C., Hung, K. and Lee, C.F. (2004), "The Price Adjustment And Lead-Lag Relations Between Stock Returns: Microstructure Evidence From The Taiwan Stock Market", *Journal of Empirical Finance*, Vol. 11, pp. 709–731.

Choudhry, T. (2000), "Day of the Week Effect in Emerging Asian Stock Markets: Evidence from the GARCH Model", *Applied Financial Economics*, Vol. 10, pp. 235-242.

Clare, A, Morgan, G. and Thomas, S. (2002), "Direct Evidence of Non-Trading on the London Stock Exchange", *Journal of Business Finance and Accounting*, Vol. 29 Nos. 1&2, pp. 29-53.

Cohen, K., Maier, S., Schwartz, R. and Whitcomb, D. (1979), "On the Existence of Serial Correlation in an Efficient Securities Market", *TIMS Studies in the Management Sciences*, Vol. 11, pp. 151-168.

Conrad, J., Gultekin, M.N. and Kaul, G. (1991), "Asymmetric Predictability of Conditional Variances", *Review of Financial Studies*, Vol. 4 No. 4, pp. 597-622.

Day, T.E., and Wang, P. (2002), "Dividends, Nonsynchronous Prices, and the Returns from trading the Dow Jones Industrial Average", *Journal of Empirical Finance*, Vol. 9, pp. 431-454.

Engle, R.F. and Granger, C.W.J. (1987), "Cointegration and Error Correction: Representation, Estimation and Testing", *Econometrica*, Vol. 55, pp. 251-276.

Fama. E.F. (1991), "Efficient Capital Markets: II", *Journal of Finance*, Vol. 46 No. 5, pp.1575-1617.

Fisher, L. (1966), "Some New Stock Market Indexes", *Journal of Business*, Vol. 39, pp. 191-225.

Granger, C.W.J. (1969), "Investigating Causal Relations By Econometric Methods And Cross Spectral Methods", *Econometrica*, Vol. 37 No. 3, pp. 424-438.

Green, C.J. and Joujon, E. (2000), "Unified Tests of Causality and Cost of Carry: The Pricing of The French Stock Index Futures Contract", *International Journal of Finance and Economics*, Vol. 5 No. 2, pp. 121-140.

Jarrow, R.A. and Protter, P. (2011), "A Dysfunctional Role of High Frequency Trading in Electronic Markets", *Johnson School Research Paper Series*, #08-2011, Cornell University, http://ssrn.com/abstract=1781124, accessed January 8th 2013.

Kadlec, G.B. and Patterson, D.M. (1999), "A Transactions Data Analysis of Nonsynchronous Trading", *The Review of Financial Studies*, Vol. 12 No. 3, pp. 609-630.

Li, D.D. and Yung, K. (2006), "Stock Return Autocorrelation and Institutional Investors: The Case of American Depository Receipt", *Review of Accounting and Finance*, Vol. 5 No. 1, pp. 45-58.

Lo, A.W. and MacKinlay, A.C. (1990a), "An Econometric Analysis of Non-Synchronous Trading", *Journal of Econometrics*, Vol. 45, pp. 181-211.

Lo, A.W. and MacKinlay, A.C. (1990b), "When Are Contrarian Profits Due To Stock Market Overreaction", *The Review of Financial Studies*, Vol. 3 No. 2, pp. 175-205.

Niarchos, N.A. and Alexakis, C.A. (1998), "Stock Market Prices, 'Causality' And Efficiency: Evidence From The Athens Stock Exchange", *Applied Financial Economics*, Vol. 8, pp. 167-174.

Papachristou, G. (1999), "Stochastic behaviour of the Athens Stock Exchange: A Case Of Institutional Nonsynchronous Trading", *Applied Financial Economics*, Vol. 9, pp. 239-250.

Pesaran, M.H. and Timmermann, A. (1992), "A Simple Non-Parametric Test Of Predictive Performance", *Journal of Business and Economic Statistics*, Vol. 10 No. 4, pp. 461-465.

Poshakwale, S. and Theobald, M. (2004), "Market Capitalisation, Cross-Correlations, the Lead/Lag Structure and Microstructure Effects in the Indian Stock Market", *Journal of International Financial Markets, Institutions and Money*, Vol. 14, pp. 385-400.

Roll, R. (1984), "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market", *Journal of Finance*, Vol. 39, pp. 1127-1139.

Scholes, M., and Williams, J. (1977), "Estimating Betas From Nonsynchronous Data", *Journal of Financial Economics*, Vol. 5, pp. 309-327.

Shah, A. and Sivakumar, S. (2000), "Changing Liquidity in the Indian Equity Market", *Emerging Markets Quarterly*, Summer, pp. 62-71.

Sornette, D. and Von der Becke, S., (2011), "Crashes and High Frequency Trading", *Swiss Finance Institute Research Paper Series*, No 11 – 63, http://ssrn.com/abstract=1976249, accessed January 8th 2013.

Tian, G.G. and Wan, G.H. (2004), "Interaction Among China-Related Stocks: Evidence From A Causality Test With A New Procedure", *Applied Financial Economics*, Vol. 14, pp. 67-72.

Wood, R.A., McInish, T.H. and Ord, J.K. (1985), "An Investigation of Transactions Data for NYSE Stocks", *Journal of Finance*, Vol. 40 No. 3, pp. 723-739.

Camilleri, S.J. and C.J. Green, (2014),"Stock market predictability: Non-synchronous trading or inefficient markets? Evidence from the National Stock Exchange of India", *Studies in Economics and Finance*, 31 (4), pp. 354 - 370.

Tables:

Table 1. Average Trading Frequencies for Nifty and Nifty Junior Shares

Panel A shows the average number of transactions for all shares in the indices for five different trading days common to both the daily and one-minute data sets. The average waiting time is the average interval between trades for each share during the 5 sample trading days, averaged over all the shares in the index. Each trading day is 5½ hours long.

Panel B reports the same statistics for the ten least frequently traded shares in the indices.

Panel C reports the results of paired means tests on the null hypothesis of no difference in trading frequency between the indices. The test provides a day-by-day comparison, and the critical values for a two-tailed test (n=5) are 2.1318 (90%), 2.7765 (95%), 4.6041 (99%). The t-values reject the null hypothesis at the 95% level of confidence. Significance at the 99%, 95% and 90% levels of confidence is denoted by ***, ** and * respectively.

Trading Days	**Nifty**		**Nifty Junior**	
Panel A: All Shares				
	Av. no. Trades per Share	**Av. Waiting Time (seconds)**	**Av. no. Trades per Share**	**Av. Waiting Time (seconds)**
15-Jun-99	1162	17.0	1015	19.5
17-Jun-99	1240	16.0	900	22.0
21-Jun-99	1027	19.3	998	19.8
23-Jun-99	1386	14.3	956	20.7
25-Jun-99	1310	15.1	859	23.1
Panel B: 10 Least Traded Shares				
	Av. no. Trades per Share	**Av. Waiting Time (minutes)**	**Av. no. Trades per Share**	**Av. Waiting Time (minutes)**
15-Jun-99	41	8	29	12
17-Jun-99	25	13	21	16
21-Jun-99	30	11	24	14
23-Jun-99	41	8	40	8
25-Jun-99	31	11	24	14

Panel C:	
Null Hypothesis: No difference between Nifty and Nifty Junior	**t-statistic**
All Shares	
Av. no. Trades per Share	3.3879 **
Av. Waiting Time (seconds)	3.3976 **
10 Least Traded Shares	
Av. no. Trades per Share	3.3029 **
Av. Waiting Time (minutes)	3.8335 **

Camilleri, S.J. and C.J. Green, (2014),"Stock market predictability: Non-synchronous trading or inefficient markets? Evidence from the National Stock Exchange of India", *Studies in Economics and Finance*, 31 (4), pp. 354 - 370.

Table 2. Pesaran-Timmermann Tests

The table shows Pesaran-Timmermann Test Statistics for the relationship between the Nifty and Nifty Junior. We report statistics for the contemporaneous returns and between current and lagged returns using up to 25 lags. $S(M_t,N_{t-i})$ are tests between the Nifty Junior (M_t) and i lags of the Nifty (N_{t-i}); $S(N_t,M_{t-i})$ are tests between the Nifty (N_t) and i lags of the Nifty Junior (M_{t-i}). The S statistics are asymptotically normally distributed, and therefore the critical values for the test are: 2.58, 1.96 and 1.65 for the 99%, 95% and 90% level of confidence respectively. Significance at these levels of confidence is denoted by ***, ** and * respectively.

	Daily Data		One-Minute Data	
Lag(i)	$S(M_t,N_{t-i})$	$S(N_t,M_{t-i})$	$S(M_t,N_{t-i})$	$S(N_t,M_{t-i})$
0	37.14 ***	37.14 ***	18.83 ***	18.83 ***
1	6.30 ***	3.61 ***	13.41 ***	11.61 ***
2	1.23	0.52	8.18 ***	3.37 ***
3	1.73 *	0.95	5.96 ***	0.34
4	1.03	0.25	4.47 ***	0.87
5	0.81	0.37	5.19 ***	0.52
6	-0.74	-2.24 **	3.92 ***	-1.08
7	0.13	-0.31	1.95 *	0.48
8	-0.39	0.59	0.35	-1.74
9	1.15	0.95	-0.36	-2.42
10	0.93	1.38	0.90	-1.56
11	-0.04	-0.42	0.67	-2.49
12	-0.34	-0.16	-0.31	-0.89
13	0.50	1.12	0.23	-1.10
14	1.64	0.01	-0.71	-0.79
15	0.98	-0.32	-0.95	-0.51
16	0.24	0.93	-1.82	-1.63
17	2.04 **	1.70 *	-0.11	-0.99
18	-0.23	-0.23	-1.52	-1.56
19	0.33	-0.62	-0.14	-1.25
20	-0.72	-0.33	-0.71	-1.75
21	-0.59	-0.35	0.82	0.11
22	0.93	-0.95	1.06	-0.68
23	-0.35	0.17	0.27	-0.66
24	0.35	1.55	1.32	0.72
25	-0.66	0.10	1.00	1.59

Camilleri, S.J. and C.J. Green, (2014),"Stock market predictability: Non-synchronous trading or inefficient markets? Evidence from the National Stock Exchange of India", *Studies in Economics and Finance*, 31 (4), pp. 354 - 370.

Table 3. VAR Results and Tests for Daily Data

The table shows the final VAR for the Nifty and Nifty Junior log returns using daily data. LRN and LRM are the Nifty and Nifty Junior log returns respectively.

Diagnostics: SE is the standard error of the regression; **RSS** is the residual sum of squares. **F(2,3495)** is the F test for zero slope coefficients; **ARCH** is the LM test of Engel (1982) for ARCH residuals based on the regression of squared residuals on squared fitted values, and distributed as $\chi^2(1)$.

Granger: These show Granger Non-Causality Tests for the significance of the LRM variables in the LRN equation and *vice-versa*. In each case, the test cannot reject the null of non-causality.

Log Likelihood: LL is the log-likelihood for each equation estimated separately by OLS, and for the two equation system estimated by SUR. **The LR test** tests the null hypothesis that there is no contemporaneous covariance between the two series. This is calculated as twice the difference between the system Log-Likelihood and the sum of the individual equation log likelihoods, and is distributed as $\chi^2(2)$. We reject the null hypothesis that shocks are contemporaneously uncorrelated.

Significance at the 99%, 95% and 90% levels of confidence is denoted by ***, ** and * respectively.

3498 observations	**Nifty Regression (LRN)**		**Nifty Junior Regression (LRM)**	
Regressor	**Coefficient**	**T-Ratio**	**Coefficient**	**T-Ratio**
LRN(t-1)	0.085 ***	(2.660)	0.065 *	(1.771)
LRM(t-1)	-0.023	(-0.817)	0.102***	(3.210)
Intercept	0.001*	(1.806)	0.001	(1.546)
Diagnostics				
S.E. and RSS	0.017	0.958	0.019	1.259
R-Bar-Squared	0.004		0.0228	
F(2,3495)	7.326***	$F_{0.01} = 4.63$	41.864 ***	$F_{0.01} = 4.63$
ARCH: $\chi^2(1)$	85.54***	$\chi^2_{0.01} = 6.635$	342.5 ***	$\chi^2_{0.01} = 6.635$
Tests				
Granger: $\chi^2(1)$	0.667	$\chi^2_{0.05} = 3.841$	3.138	$\chi^2_{0.05} = 3.841$
LL: equation and system	9383.6	20578.1	8905.0	20578.1
LR test: $\chi^2(2)$	4579***		$\chi^2_{0.01} = 9.21$	

Table 4. VAR Results and Tests for One-Minute Data

The table shows the final VAR for the Nifty and Nifty Junior log returns using one-minute data. LRN and LRM are the Nifty and Nifty Junior log returns respectively; Dummy takes a value of unity for the first 2 observations of the trading day and zero otherwise.

Diagnostics: SE is the standard error of the regression; **RSS** is the residual sum of squares. **F(2,2028)** is the F test for zero slope coefficients; **ARCH** is the LM test of Engel (1982) for ARCH residuals based on the regression of squared residuals on squared fitted values, and distributed as $\chi^2(1)$.

Granger: These show Granger Non-Causality Tests for the significance of the LRM variables in the LRN equation and *vice-versa*. In each case, the test rejects the null hypothesis of non-causality.

Log Likelihood: LL and LR tests are described in Table 3; the LR test is distributed as $\chi^2(19)$. We reject the null hypothesis that shocks are contemporaneously uncorrelated.

Significance at the 99%, 95% and 90% levels of confidence is denoted by ***, ** and * respectively.

2960 observations	**Nifty Regression (LRN)**		**Nifty Junior Regression (LRM)**	
Regressor	**Coefficient**	**T-Ratio**	**Coefficient**	**T-Ratio**
LRN(t-1)	0.203***	(10.51)	0.241***	(16.26)
LRN(t-2)	0.012	(0.61)	0.077***	(4.96)
LRN(t-3)	0.010	(0.50)	0.075***	(4.84)
LRN(t-4)	-0.010	(0.49)	0.019	(1.21)
LRN(t-5)	0.013	(0.64)	0.023	(1.47)
LRN(t-6)	-0.030	(1.50)	0.034**	(2.18)
LRN(t-7)	0.048***	(2.37)	0.011	(0.72)
LRN(t-8)	0.000	(0.01)	-0.045***	(2.90)
LRN(t-9)	0.015	(0.77)	0.039***	(2.59)
LRM(t-1)	0.162***	(6.06)	0.027	(1.31)
LRM(t-2)	-0.036	(1.35)	-0.080***	(3.92)
LRM(t-3)	0.041	(1.53)	0.031	(1.52)
LRM(t-4)	-0.028	(1.06)	-0.011	(0.53)
LRM(t-5)	0.033	(1.26)	-0.010	(0.50)
LRM(t-6)	-0.028	(1.06)	0.031	(1.51)
LRM(t-7)	-0.078***	(2.97)	-0.042**	(2.07)
LRM(t-8)	-0.002	(0.06)	0.036*	(1.78)
LRM(t-9)	-0.032	(1.28)	0.011	(0.55)
Dummy	0.003***	(19.18)	0.000	(1.29)
Intercept	0.000	(0.45)	0.000	(0.69)
Diagnostics				
S.E. and RSS	0.0006	0.0010	0.0004	0.0006
R-Bar-Squared	0.185		0.162	
F(19,2940)	36.34***	$F_{0.01} = 1.96$	31.02***	$F_{0.01} = 1.96$
ARCH: $\chi^2(1)$	975.1***	$\chi^2_{0.01} = 6.635$	65.61***	$\chi^2_{0.01} = 6.635$
Tests				
Granger: $\chi^2(9)$	51.6***	$\chi^2_{0.01} = 21.67$	328.1***	$\chi^2_{0.01} = 21.67$
LL: equation and system	17915	36954	18705	36954
LR test: $\chi^2(19)$	666***		$\chi^2_{0.01} = 36.19$	

Camilleri, S.J. and C.J. Green, (2014),"Stock market predictability: Non-synchronous trading or inefficient markets? Evidence from the National Stock Exchange of India", *Studies in Economics and Finance*, 31 (4), pp. 354 - 370.

Table 5. Correlations between Indices around Overnight Trading Breaks

Table 5 shows the results of two tests of the correlation between the index returns around overnight trading breaks. The sample period is 11th June 1999 to 16th November 1999 (112 observations). These dates were chosen so as to obtain a sample in which the NSE commenced the session through continuous trading rather than through an initial call auction.

Panel A gives the simple correlation coefficients between the series which show that the Overnight Nifty Return [$OR(N)_{t\to t+1}$] is more correlated with the Initial Nifty Junior Return [$IR(M)_{t+1}$] than with the Overnight Nifty Junior Return [$OR(M)_{t\to t+1}$]. As a robustness check, we also calculated the correlation coefficients within each of six sub periods, with 18/19 observations each. These confirm the finding of higher correlation levels between $OR(N)_{t\to t+1}$ and $IR(M)_{t+1}$. A paired t-test on these data indicated that the difference between the sets of correlations is significant at the 95% confidence level.

Panel B shows regression estimates which confirm that $OR(N)_{t\to t+1}$ performs better in explaining $IR(M)_{t+1}$ rather than $OR(M)_{t\to t+1}$. In estimating the models in Panel B, the sample period was split into two due to a missing intra-day observation for the 22nd September 1999.

Significance at the 99%, 95% and 90% levels of confidence is denoted by ***, ** and * respectively.

Panel A: Correlations						
Sample: 11th June 1999 to 16th November 1999 (112 observations)						
Correlation between $OR(N)_{t\to t+1}$ and $OR(M)_{t\to t+1}$: 0.2916						
Correlation between $OR(N)_{t\to t+1}$ and $IR(M)_{t+1}$: 0.5153						
Correlations within 6 sub-periods						
$OR(N)_{t\to t+1}$ and $OR(M)_{t\to t+1}$	0.0279	0.0515	0.6930	0.2752	0.0681	0.2090
$OR(N)_{t\to t+1}$ and $IR(M)_{t+1}$	0.7688	0.1021	0.7269	0.5987	0.3655	0.5848

Panel B: Regressions				
Sample Period:	**11-Jun-1999 to 21-Sep-1999**		**23-Sep-1999 to 16-Nov-1999**	
Observations	72		39	
	Coefficient	**T-Ratio**	**Coefficient**	**T-Ratio**
First Model: Dependent Variable is OR(M)t→t+1				
α	0.00012	0.963	0.00003	0.177
$OR(N)_{t\to t+1}$	0.75604	3.333***	0.28411	1.352
R-bar-squared	0.1246		0.0213	
Second Model: Dependent Variable is IR(M)t+1				
α	0.00257	3.394***	0.00213	0.947
$OR(N)_{t\to t+1}$	7.5782	5.690***	9.5373	3.420***
R-bar-squared	0.3065		0.2196	

Camilleri, S.J. and C.J. Green, (2014),"Stock market predictability: Non-synchronous trading or inefficient markets? Evidence from the National Stock Exchange of India", *Studies in Economics and Finance*, 31 (4), pp. 354 - 370.

Footnotes

1 Non-synchronous trading also affects the information-revelation properties of the market index and suggests reasons for trading index futures rather than index tracking funds (Green and Joujon, 2000).

2 Studies which suggest that foreign investors restrict holdings to relatively liquid shares, contributing to their higher efficiency, include Niarchos and Alexakis (1998) and Tian and Wan (2004).

3 Augmented Dickey Fuller tests showed that log prices are I(1) while log returns are I(0).

4 Other coefficients at lags 6 and 17 are significant at the 95% level; we can consider these to be rogue occurrences, on the grounds that there seems no plausible practical explanation for predictability at these specific lag lengths.

5 Before estimating the VARs we established that the series are not cointegrated, since VARs on differenced data of cointegrated series are mis-specified (Engle and Granger, 1987).

6 The intra-day file for 22nd September 1999 was unavailable.