An Analysis of the Impacts of Non-Synchronous Trading On Predictability: Evidence from the National Stock Exchange, India

by

Silvio John Camilleri and Christopher J. Green

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

Green: **Department of Economics, Loughborough University**

Correspondence to:
Christopher J. Green, Department of Economics, Loughborough University, Loughborough, Leicestershire, LE11 3TU, United Kingdom
Tel: +44 (0) 1509 222711; *Fax:* +44 (0) 1509 223910;
E-mail: C.J.Green@lboro.ac.uk; Silvio.J.Camilleri@um.edu.mt

November 2004

Preliminary Draft

AN ANALYSIS OF THE IMPACTS OF NON-SYNCHRONOUS TRADING ON PREDICTABILITY: EVIDENCE FROM THE NATIONAL STOCK EXCHANGE, INDIA

by

Silvio John Camilleri and Christopher J. Green

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

Green: **Department of Economics, Loughborough University**

Preliminary Draft

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

ABSTRACT

The serial correlation effects which non-synchronous trading can induce in financial data have been documented by various researchers. In this paper we investigate non-synchronous trading effects in terms of the predictability that may be induced in the values of stock indices. This analysis is applied to emerging-market data, on the grounds that such markets might be less liquid and thus prone to a higher degree of non-synchronous trading. We use both a daily data set and a higher frequency one, since the latter is a prerequisite for capturing intra-day variations in trading activity. When considering one-minute interval data, we obtain clear evidence of predictability between indices with different degrees of non-synchronous trading. We then propose a simple test to infer whether such predictability is mainly attributable to non-synchronous trading or an actual delayed adjustment on part of traders. The results obtained from an intra-day analysis suggest that the former cause seems a better explanation for the observed predictability. Future research in this area is needed to shed light on the degree of data predictability which may be exclusively attributed to non-synchronous trading, and how empirical results may be influenced by the chosen data frequency.

1. Introduction

Theoretical models of market microstructure are typically formulated in continuous time, yet empirical studies of securities market data necessarily use discrete time data sets. This discrepancy may constitute a challenge to researchers aiming to investigate the validity of theoretical models by analysing empirical data. This challenge becomes more pronounced when the data set includes non-synchronous trading effects, meaning that some particular stocks do not trade for prolonged periods – yet researchers typically assume that the prices were sampled simultaneously. Such assumptions may be vital for the sake of research manageability, yet they might also affect the validity of any inferences. For instance, the degree of efficiency in the market may be underestimated given that non-synchronous trading might give the impression that stock prices are not adjusting immediately to news.

For this reason, the market microstructure discipline stands to gain from a deeper understanding of non-synchronous trading effects as well as the increased availability of empirical evidence of these effects, gleaned through various methodologies. The main objective of this study is to provide new empirical evidence on this question. The paper breaks new ground in three important respects: first, we use a new test method to identify the relation between predictability and non-synchronicity; second, we employ a high quality, high frequency dataset; and third, we use data from an emerging stock market: the National Stock Exchange of India.

More specifically, we investigate the effects of non-synchronous trading on financial data; mainly in terms of inducing traces of predictability in the values of stock indices. As discussed below, prior studies which specifically tackled non-synchronous trading effects tended to focus on the autocorrelation structures of stock market data. This investigation takes an alternative approach, in that it focuses on lead-lag effects in the values of two indices. The latter predictability effects are tested for through three different techniques; namely Pesaran-Timmermann tests, Vector Autoregressions and Granger-Causality, and Impulse Response Functions. We then propose a simple test in order to infer whether any predictability effects may be attributed to non-synchronous trading or whether they constitute actual delayed adjustments of traders' expectations.

The second important aspect of this empirical investigation lies in the use of a high-frequency data set, in combination with daily data. Following the notion that trading activity varies throughout the trading day, one may deduce that non-synchronous trading effects become more amplified during those periods when trading activity abates – usually towards the middle of the day. This implies that an intra-day data set is a prerequisite in order to obtain unbiased empirical evidence of non-synchronous trading effects.

Our third innovation is in using data from an emerging market: The National Stock Exchange of India (NSE). We would expect differences between the microstructure effects of stock market trading in emerging economies and those in the industrial economies, especially given the macro evidence on market differences. As compared with developed stock markets, emerging markets tend to exhibit: higher serial correlation, less frequent trading, slower adjustment of prices to news, and indirect evidence of more insider trading (Bekaert. and Harvey, 2002). The NSE provides an interesting setting for our analysis: it has been established for just over 10 years and has had full electronic trading and recording of data from the outset. However, it still includes a significant proportion of less liquid securities, such as we might find in any emerging market. This enables us to investigate thoroughly the issues of efficiency and non-synchronous trading using daily and higher frequency data.

The rest of the paper is structured as follows: Section 2 reviews the existing literature relating to non-synchronous trading and formulates the expected results for this study whilst Section 3 describes the data sets. Sections 4-6 test for predictability effects in the data set through different techniques, namely Pesaran-Timmermann Tests, Vector Autoregression and Granger Causality, and Impulse Response Functions. Section 7 investigates whether the observed predictability effects are more attributable to non-synchronous trading as opposed to actual delayed adjustments of expectations on part of traders. Section 8 concludes.

2. Non-Synchronous Trading: Literature Review and Expected Results

This Section provides an outline of previous research in the area, and then formulates the expected results from the study.

2.1 Research Background

Non-synchronous trading effects take place when transactions in particular securities occur infrequently. In such cases, the last transaction price quotations might cease to reflect the fundamental value of the firm as new information becomes available. At face value, this gives the impression that the stock price delays in adjusting to new information; yet the underlying cause of the apparent inefficiency is that the most recent trading price relates to a past transaction, and is therefore outdated. Last transaction prices of infrequently traded securities might be used for calculating the value of a portfolio of stocks or the value of a market index. At times, the validity of this methodology is undermined since such calculations might be based on a partly outdated data set due to non-synchronous trading. The latter data does not imply that there still exist market participants who are willing to trade at those prices.

One problem with non-synchronous trading effects is that it is usually cumbersome for researchers to inquire the last transaction time for each and every stock and it is often assumed that securities prices were sampled simultaneously. Yet this might not necessarily be the case,

say when working with a cross-section of security closing prices, and it might therefore amount to a limitation in the research methodology.

Non-synchronous trading also induces particular characteristics in stock price data. For instance, stock price indices tend to exhibit higher levels of serial correlation than individual stocks, as discussed by Fisher (1966). Cohen, Maier, Schwartz and Whitcomb (1979) showed that non-synchronous trading induces serial correlations in market returns.

Atchison, Butler and Simonds (1987) compared the observed serial correlation of a portfolio of NYSE stocks to that predicted by a model of non-synchronous trading as proposed by Scholes and Williams (1977). They found that the actual serial correlation was higher than that predicted by the non-synchronous trading model, and they attributed this to other sources of delayed price adjustment. There might be various reasons why prices take longer to adjust to new information. For instance, market participants who submit limit orders do not necessarily monitor these orders continuously. As new information becomes available, such orders may become mispriced and some other participants might "pick off" these orders and trade profitably on the basis of superior information. In this way a transaction which occurs at an outdated price might still be consistent with market efficiency, since an efficient market does not require *all* market participants to price in the new information instantaneously. A further reason why a delayed price adjustment may occur is that participants might not devote enough time in monitoring less liquid stocks, as they do with the most liquid ones. Thus, new information relating specifically to the former stocks might take longer to get priced in. Therefore, not all of the pricing delays which are evident in stock price data are the result of non-synchronous trading.

Various other authors have investigated the effects of non-synchronous trading on the autocorrelation of stock returns, and these include Lo and MacKinlay (1990), Boudoukh, Richardson and Whitelaw (1994) and Kadlec and Patterson (1999). Whilst all these studies conclude that non-synchronous trading increases return serial correlation, they disagree as to what is the specific autocorrelation level which emanates from non-synchronous trading. Part of the discrepancy in between the studies may be attributed to the differing assumptions as regards the non-trading intervals of securities.

An "explicit" case of non-synchronous trading was analysed by Papachristou (1999) in the context of the Athens Stock Exchange. Prior to 1989 the trading day on this exchange consisted of successive trading sessions, and shares of a particular industry traded in each of the sessions. Thus, the returns reported at the end of the day – and therefore the closing index value – constituted partly "outdated" information since they related to transactions which took place earlier on during the day. The non-trading periods for the particular stocks in this case were deterministic, and the author called this "deterministic non-synchronicity". The author

showed that the overall effects are similar to those of stochastic non-synchronicity, including serial correlation in market index returns and cross-correlation in between stocks.

Given that changes in expectations may take longer to show up in share price fluctuations if the latter trade infrequently, non-synchronous trading may result in lead-lag effects in between the prices of various stocks. This induces predictability elements in the data. Yet, this degree of predictability does not necessarily translate into abnormally profitable trading opportunities, as shown by Day and Wang (2002) after simulating a trading strategy in the Dow Jones Industrial Average index which was adjusted for non-synchronous trading effects.

The aim of this study is to investigate the lead-lag effects in between two indices which differ in their degree of non-synchronous trading. Intuitively, we may expect that the index which features the less liquid securities will "take longer to adjust to new market information" and therefore the more liquid index will lead the less liquid one. Overall, this seems to be the first study that analyses lead-lag relationships between two indices, in the context of non-synchronous trading effects as the central issue. As noted above, previous non-synchronous trading studies tended to focus on the serial correlation structure of the return data. Conversely, most studies of lead-lag relationships in between indices or stock portfolios do not place their principal emphasis on non-synchronous trading. Yet, some of the latter lead-lag relationship studies still propose relevant conclusions to the issue at hand, and these include Lo and MacKinlay (1990) and Mills and Jordanov (2000).

Lo and MacKinlcay (1990) analysed US stock price data which was sampled at different inter-day frequencies. They found that large-capitalization stocks tend to lead the stocks of smaller companies. One important factor which results in such "causality" is the cross-correlation between stocks over time. The latter may well be a by-product of non-synchronous trading, as shown by various authors such as Cohen, Maier, Schwartz and Whitcomb (1979). Mills and Jordanov (2000) reported similar lead-lag evidence for a number of UK stocks which were sampled at monthly intervals. The authors constructed 10 different size portfolios of stocks which were selected from the FTSE-Actuaries All Share Index. Their methodology comprised the creation of Generalised IRFs and thus went beyond the analysis of the correlation structures of the respective portfolios.

Our analysis is one of the first studies that test for lead-lag relationships between two indices, using high-frequency data with the specific aim of gleaning evidence of non-synchronous trading effects. Using a high-frequency data set is important given that intra-day effects are probably more relevant to analysing such effects from a market microstructure point of view as discussed in Section 2.2.

2.2 Expected Results

The novel approach in this study is that the effects of non-synchronous trading will be investigated in terms of leads and lags in between two market indices which feature differing degrees of liquidity. We may expect the more liquid index to lead the less liquid one. Any lead-lag effects might be consistent with the fact that market participants take longer to adjust their judgement regarding the fundamental value of the firms which trade less frequently. Yet, such lead-lag effects may simply be the result of non-synchronous trading effects in the data, as discussed above. In line with the above literature, we may expect both of the former factors to be contributing to lead-lag effects. In this way, Section 7 proposes a simple methodology in order to infer which of these two predictability causes is most relevant in explaining the lead-lag relationships.

As outlined above, another novel contribution in this analysis emanates from the use of a high-frequency data set, in addition to a daily data set. One may expect non-synchronous trading effects to be more visible in a high-frequency data set. This rests on the fact that trading activity typically varies throughout the trading day, as shown by various authors such as Wood, McInish and Ord (1985). Non-trading periods for less liquid stocks might be more likely to occur during particular periods of the day, and such effects require a high-frequency data set to detect. Empirical studies tend to show a rise in trading activity at the end of the day. In our empirical setting of NSE India, trading activity peaks at the end of the day, as discussed by Shah and Sivakumar (2000). This implies that non-synchronous trading effects become less significant at the end of the day, and thus more difficult to detect when using a data set which is based on closing prices.

Summing up, we expect the more liquid index to lead the less liquid one and that such an effect becomes more pronounced when analysing the high frequency data set. Based on the inferences of previous studies, we expect such predictability elements to be partly attributable to actual delays in price adjustments as well as due to non-synchronous trading. At the end of the paper we investigate which of the latter causes is more relevant in explaining the observed lead-lag effects.

3. Empirical Setting and Data Set Characteristics

This Section provides a brief description of our empirical setting and the data sets.

The National Stock Exchange of India (NSE) was established in 1994 and is one of two major Indian exchanges, together with the Bombay (Mumbai) Stock Exchange (BSE). During 2000, around 1,300 equities traded on NSE, through 960 brokerage firms. [1] Most major stocks are quoted on both NSE and BSE and these exchanges compete both for listings and order flow. As at 1999, the volume of a typical trading day on the NSE was around 400,000 transactions.

[1] Shah and Sivakumar (2000).

The NSE was set up with on-line, continuous, screen-based, nationwide electronic trading. Subsequently, the exchange introduced as an experiment a pre-opening and post-closing call auction. This followed the basic rule that the resulting price should maximise the total traded quantity, implying that the gap between demand and supply should be ideally zero. For the rest of the trading day the system continued to function as a continuous pure limit order book market, with time and price priorities applied to incoming orders. There are no official market-makers on the NSE, though one should not ignore the possibility that some traders opt to perform such functions.

The data was extracted from the NSE's historical trades data CDs. [2] These include data on index values, and the volume and price of all trades carried out on the exchange on a trade-by-trade basis.

The analysis is applied both to daily data as well as higher frequency data sampled at one minute intervals. The daily data set constitutes of the closing observations of the NSE Nifty and Midcap indices – the main index and the smaller capitalisation index respectively. Each index comprises 50 stocks and no security may be included in both of the indices. The daily data period ranges from 1st January 1999 to 31st December 2003 – a total of 1257 observations.

The high frequency data set constitutes of the values of both indices, sampled at one minute intervals, over the period ranging from 15th June 1999 to 25th June 1999. This period includes nine continuous trading sessions starting from 10 a.m. and ending at 3.30 p.m., yielding a total of 2970 observations.

Figures 1 and 2 show the logarithmic series and log return plots for both indices, sampled over daily intervals. Figures 3 and 4 show logarithmic series and log return plots for both indices, sampled over one minute intervals. Summary statistics and ADF tests for these series are shown in Tables 1 and 2.

Both in the case of the daily and one minute frequency data sets, the Augmented Dickey Fuller tests do not reject the null hypothesis of a unit root when considering the natural logs of the original price series. The null hypothesis of a unit root is rejected when considering the logarithmic returns. This implies that the logarithmic prices may be classified as I(1) since they are stationary in the first differences.

An informal procedure which may be used to cross-check this test is to look at the autocorrelation of the series at different lags. High autocorrelation coefficients which decline slowly may be taken as an indication of a unit root. The autocorrelation coefficients shown in Appendix A indicate that the logarithmic series are I(1).

Given that non-synchronous trading effects depend on the level of liquidity, it is important to assess the relative liquidity of the indices used in this study. One may reasonably expect the

[2] We thank NSE for providing us with trial copies of these data sets.

main index (Nifty) to be more liquid than the smaller capitalisation index (Midcap), since the shares included in the main index are those of larger capitalisation companies which are likely to be more popular amongst traders. Yet, an empirical investigation about the trading frequencies of the shares included in the indices was undertaken as shown below.

Table 3 shows the average trading frequencies for the individual shares in the Nifty and the Midcap index. Given that the individual shares in the indices change periodically, the index composition as at 18 May 2003 was chosen, largely due to the unavailability of index composition historical data. The sampled trading days are included in both the daily and the high frequency data sets, and this restricted the potential sample to 9 trading days (since the latter data set consists of 9 trading days). Alternate trading days from the latter data set were chosen.

Given this, we should note that this preliminary investigation suffers from various limitations. In particular, it is assumed that the index compositions and the respective liquidity levels do not change materially over the sample period.

Table 3 unambiguously shows that the Nifty index is more liquid than the Midcap, in line with our expectations. On average, there is a waiting time of around 6 seconds in between trades for the shares included in the Nifty index, while the waiting time for the shares in the Midcap index is around 20 seconds. This implies that each of our index observations is in fact "outdated" by around 6 seconds and 20 seconds respectively. In estimating the average waiting time, it was assumed that trades occur evenly throughout the day. This is not usually a realistic assumption, and according to Shah and Sivakumar (2000) NSE trading tends to peak at the end of the day. This implies that during some periods throughout the day, the average waiting time will decrease or increase; and in the latter case the non-synchronous trading effects become more pronounced.

Given that non-synchronous trading effects are mainly caused by the less liquid stocks, it also makes sense to look at the waiting time statistics of the less liquid shares in the respective indices. This is not an entirely "pessimistic" approach, given that the outdated trade price information of these shares is incorporated in the index value, resulting in an imprecise yardstick. The waiting time statistics of the ten least frequently traded shares in the indices are shown in Table 4. The table shows that one quintile of the information which is used in estimating Nifty observations is at least five minutes old. Similarly, Midcap observations are based on a data set, a quintile of which is around ten minutes old.

When working with intra-day observations, we should observe considerably higher non-synchronous trading effects, following the notion that intra-day activity is somewhat lower. In case of the daily data set, the index value occurs at the end of the day – which is a trading peak, and therefore one may argue that the non-synchronous trading effects should be less pronounced.

Thus, we may now re-formulate our expectations more specifically. Firstly, we expect that the Midcap index should appear "less efficient" than the Nifty index, in the sense that we should obtain an indication that Midcap returns may be predicted to some degree from Nifty returns. This "inefficiency" is partly the result of non-synchronous trading, and the study aims to infer whether the predictability is more attributable to non-synchronous trading effects or actual delayed adjustments of traders' expectations.

Secondly, we may also expect that non-synchronous trading effects to be more pronounced in the high-frequency data set. Following the observation of Shah and Sivakumar (2000) that NSE trading peaks at the end of the day, end-of-day observations should be based on reasonably recent trade information and therefore these observations should include less non-synchronous trading effects.

The analysis now proceeds with investigating predictability effects in the data using three different methodologies: Pesaran Timmermann Tests, Vector Autoregression and Granger Causality, and Impulse Response Functions.

4. Pesaran-Timmermann Tests

We now turn to the first predictability investigation through Pesaran-Timmermann (1992) tests. The test proposed by Pesaran and Timmermann (1992) measures the dependence between two time series, in terms of whether the series fluctuate in the same direction. Therefore, this non-parametric test considers the direction of the changes and largely ignores the magnitude of the fluctuations. The procedure tests the null hypothesis that the series are independent, and the test statistic is normally distributed in case of large samples. One potential application of such this test is in assessing the predictive power of a model; whereby the values predicted by the particular model are compared to the actual values.

The test statistic for assessing the relationship between variables x_t and y_t is computed as follows:

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

where

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

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

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

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

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

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

and the function *Sign(Z)* takes a value of 1 when the variable is positive and zero otherwise.

Thus, $\hat{P}$ is a measurement of the number of occurrences where both time series changed in the same direction – it takes a value of 1 when all the respective changes are all in the same direction, and a value of 0 when all contemporaneous changes are always in an opposite direction. The terms $\hat{P}_*$, $\hat{V}(\hat{P})$, and $\hat{V}(\hat{P}_*)$ adjust this "crude" measurement by considering the individual proportions of negative and positive changes in both of the series, and scale the original measurement to a normal distribution.

This test may be used in detecting lead-lag effects if we apply it to the relationships between x_t and y_{t-n} or between x_{t-n} and y_t , instead of x_t and y_t as discussed in the general case above.

Given that this procedure is essentially a test on the signs, rather than the magnitude of a time series, it does not matter whether we analyse the simple returns or the log returns. Yet, the test cannot be applied to the original prices given that the latter cannot be negative and the test cannot be applied to series which do not change sign.

Pesaran-Timmermann tests were conducted not only on the contemporaneous relationship between the indices, but also on the relationship between the change of an index at time *t* with the lagged changes of the other index. This would be similar to stating that the changes of a particular index are "Granger-Causing" the changes of the other index; although the terminology "Granger-Causing" requires some specification. Its usage is justified in that if we find relationships between the lagged change directions and the current change directions of the respective indices, this does not imply *actual* causality, for reasons outlined in Section 5. Yet, this methodology is different from Granger-Causality in that the latter considers the magnitude of returns, in addition to the direction of the changes.

Table 5 shows the Pesaran-Timmermann statistic for both the daily and intra-day data series. At both frequencies, the Nifty and Midcap indices tend to contemporaneously move in the same direction, as witnessed by the highly significant statistics of 20.28 for the daily data and 18.83 for the intra-day data. Assuming that the difference between these statistics is not due to the different sample periods, the fact that the test statistic at daily frequency is higher than that at the one minute frequency implies that the synchronicity between the indices is higher at

daily intervals. The less pronounced synchronicity of the intra-day data may indicate that the indices do not adjust as contemporaneously during the day, since otherwise we would have obtained similar statistics for both frequencies. This is in line with the above arguments that adjustments may take a longer series of observations to get priced in when considering high frequency data and that non-synchronous trading effects are more evident in the high frequency data set. Another possible reason might be that high frequency data contains a higher level of "noise".

Inspecting the Pesaran-Timmerman statistics for the daily data set in more detail, both the contemporaneous and the first lag are significant in predicting the change direction of the other index. The significant contemporaneous relationship implies that the indices tend to move in the same direction, as may be reasonably expected. Ignoring the other significant lags for the time being, we may note that the Nifty and Midcap indices are roughly equally relevant for predicting the direction change of the other index. This implies that the difference in liquidity in between the daily data sets is not resulting in differing degrees of predictability. We may thus assume that this predictability effect is not the result of non-synchronous trading or differing degrees of liquidity. One possible explanation might be "runs" in the data, whereby each of the indices changes successively in the same direction over a period of days. Data runs do not necessarily contradict the efficient market hypothesis given that longer runs may be considered as a normal feature when analysing a series that may be classified as a random walk with drift.

The other significant lagged observations in the daily data set are the third Midcap lag, and the fourteenth Nifty lag. As regards the latter, there might be no economic reason why the current direction change of the Midcap is affected by the Nifty direction change of 14 days ago, and thus we may treat it as a rogue observation. The third Midcap lag direction change is significant in predicting the current Nifty direction change. One might again think of this as a rogue observation, yet there might be an economic reason why the lagged change direction of the smaller capitalisation index affects the current change direction of the index of larger companies. Whilst causality is usually assumed to run from larger stocks to smaller ones, there might be a case for arguing that smaller companies may affect larger ones as well. Some events may impact to a higher degree on smaller companies. These may include concessions which aim to reduce the regulatory burden on smaller companies, or a sudden reduction in local consumer demand, when smaller companies rely to a larger extent on domestic trade. In such cases, one might argue that these events are immediately reflected in the prices of smaller company stocks. Yet, one may still expect that the fluctuations eventually spill over to larger company stocks. For instance, a positive news item might make smaller companies more optimistic and therefore they would increase their trading with larger companies, say those who supply them with raw materials. Therefore it might be plausible that Midcap shocks eventually spill over to Nifty prices. Yet, Table 5 indicates that it takes about 3 days for market

participants to adjust for these effects, and it may be debatable whether this is a short enough period to be consistent with such an adjustment, given that in an efficient market responses to news should be priced in instantaneously.

When considering the lagged relationships in the high-frequency data, we note that both indices tend to move in synchronicity with the other index and a number of lags. Given that the first two lagged change directions are significant in both cases, we may again attribute this to data runs. One may also note that at such high frequencies it may not be realistic to expect abrupt price changes given that as new information becomes available it is plausible that "old" limit orders do not get cancelled immediately and are "picked off"; i.e. they trade against an order that was submitted by a trader with more updated information. In this way the stock would still trade at the "old" price, despite the availability of new information.

Yet, the Nifty change directions remain significant for a further five lags. This may indicate that, as may be expected, the Nifty leads the Midcap. This might be due to non-synchronous trading effects where the less liquid stocks appear to take longer to adjust to news, given that they trade less frequently. The remaining significant lags in the high frequency data set, may either be considered as rogue observations, or they may also be consistent with feedback effects running from the Midcap to the Nifty index as discussed above.

Overall, these statistics indicate that the Nifty index leads the Midcap at high frequency data, while the indices tend to move contemporaneously at daily frequency. This may imply that the Midcap index appears to adjust more slowly – but not slowly enough to obtain a clear-cut leading relationship at daily frequency. This seems in line with our prior expectations: any "causality" mainly runs from the Nifty to Midcap, and (at least) part of this predictability is the result of non-synchronous trading effects, which become more pronounced in the high frequency data set.

We now turn to investigate whether alternative methodologies yield the same indications about such lead-lag relationships.

5. Granger-Causality Tests

This section applies Granger-Causality tests through the estimation of Vector Autoregressions (VARs). VAR methodology is based on the principle of Granger-Causality. Granger (1969) argued that if shocks in a particular time series lead to shocks in another time series, then the former series is "Granger-causing" the latter. In this way, VARs model a time series as an AR process, with the added lagged terms of another time series and an error term. If the lags of the second time series are significant, then we may argue that the latter is Granger-Causing the dependent variable. Thus, VARs offer the potential for modelling causal and feedback effects, where two or more time series Granger-cause each other.

The term "Granger-Causality" does not imply actual causality. For instance, it might be the case that the inter-relationships between the time series might in fact be caused by an exogenous variable. Therefore, Granger-Causality modelling should be accompanied by an underlying theoretical relationship since otherwise the model may be incorrectly specified.

A bivariate VAR may be formulated as follows:

$$x_t = \sum_{i=1}^{n} \alpha_{1i} x_{t-i} + \sum_{i=1}^{n} \beta_{1i} y_{t-i} + u_{1t} \tag{8}$$

and

$$y_t = \sum_{i=1}^{n} \alpha_{2i} x_{t-i} + \sum_{i=1}^{n} \beta_{2i} y_{t-i} + u_{2t} \tag{9}$$

where x_t and y_t are the variables that are assumed to Granger-cause each other, whilst u_t is an error term.

The above set of equations may be formulated in the following vector and matrix notation:

$$\begin{bmatrix} x_t \\ y_t \end{bmatrix} = \begin{bmatrix} \alpha_{1i} & \beta_{1i} \\ \alpha_{2i} & \beta_{2i} \end{bmatrix} \begin{bmatrix} x_{t-i} \\ y_{t-i} \end{bmatrix} + \begin{bmatrix} u_{1t} \\ u_{2t} \end{bmatrix} \tag{10}$$

One persisting argument concerning Granger-Causality is whether the latter implies market inefficiencies, in the sense that if an index fluctuation leads to a fluctuation in another index, this would mean that if the first fluctuation was justified on the grounds of new information, the latter fluctuation should have occurred at the same time, ruling out lead-lag effects. Various authors such as Niarchos and Alexakis (1998) argue that Granger-Causality from one price series to the other may be taken as evidence against market efficiency.

Therefore when testing for Granger-Causality using daily data, one should expect contemporaneous (but not lagged) relationships if the markets are efficient and if there are no non-synchronous trading effects. This research also uses high frequency data sampled at one minute intervals. In the latter case, one may reasonably expect some lagged relationships. For instance prices of the most liquid stocks adjust instantly to news, whilst in case of less liquid stocks one may expect that the adjustment occurs later. The latter effect might not necessarily imply that traders are inefficient. Following non-synchronous trading arguments, it might well be the case that the security does not trade immediately after the news; in this way the prices which are used for the purpose of calculating index values might be "outdated ones", giving the impression that the traders did not adjust their view about the value of the security. Thus when considering the high frequency data, we may expect Granger-Causality say, over two lags (corresponding to two minutes).

VAR and Granger-Causality are subject to a number of limitations as outlined underneath. Firstly, Granger-Causality does not necessarily imply actual causality. In the current context,

the indices may be influenced by an exogenous variable, and therefore the actual causality runs from the latter variable to the time series being studied. This argument may be relevant to the Indian stock markets, when considering the empirical evidence presented by Lamba (2003) that Indian equities are influenced by the markets of developed countries such as US, UK and Japan. As empirically shown by Capelle-Blancard and Raymond (2002), cross-country linkages may emanate both from market-wide information and non-fundamental factors such as herding and traders adopting similar trading rules that may lead to over-valuation or under-valuation.

As outlined by Renault, Sekkat and Szafarz (1998) the inferences obtained by the Granger-Causality model may be affected by the sampling process. This emanates from the fact that empirical researchers usually analyse discrete price series, whereas the underlying theoretical models typically assume continuous time. In addition, linear Granger-Causality tests may fail to detect non-linear causal relations as discussed by Baek and Brock (1992) and Hiemstra and Jones (1994). Non-linearity implies that the extent of the dependency between the time series varies during the sample period.

Engle and Granger (1987) argued that VAR estimates obtained when analysing differenced data of cointegrated time series, may be flawed since that the VAR excludes the error correction terms which appear in cointegration models[3].

We now turn to the empirical results. Section 5.1 presents the empirical results for the daily frequency data whilst Section 5.2 shows the results obtained through the high-frequency data set.

5.1 Daily Interval Data

A preliminary 24 order VAR was estimated (using the log returns series) in order to select the optimal order of the VAR. As shown in Appendix B, both the Akaike Information Criterion and the Schwarz Bayesian Criterion selected a VAR(1) model, yet the log-likelihood ratio statistics rejected all orders less than 16. In view of this, two VAR models were estimated: a VAR(1) and a VAR(16). The former model was deemed superior on the basis of higher System Log Likelihood Ratio, higher Log Likelihood Ratios for the individual equations, higher F-statistics, and higher Akaike Information Criterion and Schwarz Bayesian Criterion, as shown in Appendix C. Besides, it is not clear on practical grounds why an index shock taking place on day t, should still affect the price series on day $t+16$.

When inspecting the largest error terms of the VAR(1) regressions, it was not apparent that errors tended to occur on any particular day. If we hypothesise that more volatile returns are realised on Mondays, this would imply that the VAR(1) regressions would perform badly in

[3] Cointegration tests that were performed on the data sets used in this analysis showed that the series were not cointegrated, and therefore this limitation is not highly relevant to this analysis.

forecasting the Monday return, as well as the Tuesday returns (given that the Monday information is then used to forecast the Tuesday return). Therefore, one may expect that the Monday effect would result in larger forecast errors for Mondays and Tuesdays. The largest 40 error terms for each of the VAR regressions were inspected, but the Monday and Tuesday errors did not particularly outnumber those of other trading days.[4] This is in line with the empirical results of Choudhry (2000) who found that the change in volatility on Mondays was not significant on the Indian market. Overall, the introduction of dummy variables to account for day-of-the-week effects does not seem necessary.

The individual regressions of the estimated VAR(1) model are being reproduced in Appendices E and F together with error plots and histograms.

The LM statistic (as shown in Appendices E and F) indicates that error term is heteroskedastic. This may be attributed to exogenous factors which are not being captured by the model. Given that our main interest is the relationship between the indices, this might not be particularly problematic as long as the omitted variables do not lead to spurious results.

Summary statistics of the VAR(1) model are presented in Table 6. In both equations, the coefficients through which one can infer any lead-lag relationship between the indices are insignificant. The most "significant" coefficient in the respective regressions is the lag of the dependent variable. Despite this, the F-tests reject the null hypothesis that the joint coefficients are equal to zero in case of both regressions; yet this may be probably attributed to the lagged dependent variable, rather than the coefficient relating to the other index.

In order to investigate further, Granger non-causality tests were conducted on the Nifty and Midcap Log Return series in the system of equations. This methodology tests the null hypothesis of no causality. The test is χ^2 distributed with one degree of freedom, and test statistics of 0.17453 and 0.13989 did not permit the rejection of the null hypothesis of no-causality for both variables.

A Log Likelihood Ratio (LR) Statistic was computed to test the null hypothesis that the contemporaneous covariance between the fluctuations in the log returns series is equal to zero. The Log-Likelihood statistics for the VAR system and the LRN and LRM Equations (estimated independently through OLS) were 7168.3, 3416.0 and 3124.4. The LR (H_0:H_1) ratio is computed as [2 (7168.3 - 3416.0 - 3124.4)] and is equal to 1255.8. The test is χ^2 distributed, and the test statistic allows us to strongly reject the null hypothesis that the shocks in the log return price series are contemporaneously uncorrelated.

Overall, these results are in line with the inferences obtained through the above Pesaran-Timmermann tests, that there is only a weak lead-lag relationship (if at all) between the Nifty and Midcap indices, as sampled at daily intervals. It might be more accurate to postulate a

[4] A histogram is tabulated in Appendix D.

contemporaneous relationship between the indices at this frequency, as inferred through the LR test and the Pesaran-Timmermann test. Again, this suggests that the differing liquidity levels in between indices do not lead to pronounced non-synchronous trading effects which may be gauged through lead-lag effects in a daily-interval data set. This is probably due to the fact that as trading peaks at the end of the day, the closing index observations are based on reasonably current information.

5.2 High Frequency Data

The above procedure was repeated using the high frequency data set.

A preliminary 24 order VAR was estimated (using the log returns series) in order to select the optimal order of the VAR. As shown in Appendix G, the Akaike Information Criterion selected a VAR(9) model, whilst the Schwarz Bayesian Criterion selected a VAR(3) model. Yet, the log-likelihood ratio statistics rejected all orders which were less than 7, and therefore an order 9 VAR was selected.

The diagnostics for VAR(9) model showed problems in terms of normality and heteroskedasticity. A histogram of the error terms showed that the deviations from normality are not particularly problematic – the histograms were peak-shaped and therefore most of the error terms were very close to zero.

When plotting the error terms, it became apparent that large errors tended to occur at approximately equally-spaced intervals and this partly explains the heteroskedasticity of the error terms. The larger error terms tend to occur on the opening of the trading day – particularly at the first two observations. This is not totally surprising since a higher amount of news is priced during the first observation following the overnight interval. The lagged initial return of the trading day is then used to explain the second return during the trading day in the VAR system, and the first "unusual observation" probably leads to a particularly weak forecast for the return realised during the second minute.

Thus a dummy variable was created that takes the value of 1 during the first two observations of each trading day. A log-likelihood ratio test on the deletion of this dummy variable from the system of equations was conducted. This test is χ^2 distributed with two degrees of freedom, and a statistic of 408.35 indicated that the dummy is highly significant. The "dummy version" of the VAR resulted in regression equations with a higher R^2 and R-bar-squared, yet the diagnostics still indicated that the error term is non-normal and heteroskedastic. As regards the latter, the plot of the error terms still showed regularly-spaced errors. This may be attributed to the fact that in some cases, the large error at the opening occurred at the third, fourth, …sometimes tenth observation for the day. The dummy variable does not account for these observations. However, modifying the dummy variable to include the first ten observations of

each day would "dilute" the dummy with smaller error terms, and this might reduce the effectiveness of the dummy variable.

Granger non-causality tests were conducted on the Nifty and Midcap Log Return series in the system of equations, where the null hypothesis of no causality is tested. The test is χ^2 distributed with 9 degrees of freedom, and test statistics of 328.1 and 51.6 permitted the rejection of the null hypothesis of no-causality for both series.

A Log Likelihood Ratio (LR) Statistic was computed to test the null hypothesis that the contemporaneous covariance between the fluctuations in the log returns series is equal to zero. The Log-Likelihood statistics for the VAR system and the LRN and LRM Equations (estimated independently through OLS) were 36954, 17915 and 18705. The LR (H_0:H_1) ratio is computed as [2 (36954.2 - 17915.8 - 18705.3)] and is equal to 666.2. The test is χ^2 distributed, and the test statistic allows us to strongly reject the null hypothesis that the shocks in the log return price series are contemporaneously uncorrelated.

Thus we may assume the presence of Granger-Causality and contemporaneous effects in between the indices as sampled at one-minute intervals and propose a system of equations as shown in Table 7. [5]

The VARs fitted in this section, for the data sets at daily and one-minute frequency tend to confirm the inferences from the Pesaran-Timmermann tests of the previous section. The strongest relationship is that the Nifty index leads the Midcap at high frequency data, and there is a tendency for a feedback effect from the Midcap to the Nifty. The indices tend to move more or less contemporaneously at daily frequency.

One possible explanation for this might be that market-wide information is first reflected in the Nifty index, which includes the most liquid stocks. Some minutes after this, the information is priced in the Midcap index, and we obtain a lead-lag relationship at high frequency intervals. This may be consistent with the fact that market participants tend to monitor the major stocks (the Nifty stocks) more closely, and therefore they first price in the new information in these stocks. Whether this behaviour is in fact true or otherwise, the "extra waiting time" which is required for the Midcap stocks to trade (the evidence of which was presented in Tables 3 & 4 above), invariably results in non-synchronous trading effects in the high-frequency data set.

Yet, by the end of the day most of the new market-wide information would have been priced in both of the indices, and therefore only a contemporaneous relationship is detected when investigating relationships at daily frequency.

The Granger-Causality from Midcap to Nifty for the high-frequency data set, may be consistent with the spillover effects from smaller to larger stocks as discussed in Section 4 above. The

[5] The actual software output, including diagnostic statistics, plots and histograms of the error terms are shown in Appendices H and I.

feedback effect from Midcap to Nifty is unlikely to be the result of non-synchronous trading, since the companies constituting the main index are more liquid than the smaller capitalisation companies. Thus, the evidence obtained so far is in line with our prior expectations that non-synchronous trading effects result in predictability, yet not *all* of the predictability is the result of non-synchronous trading.

6. Impulse Response Functions

The analysis now proceeds by generating Impulse Response Functions (IRFs) for respective shocks in the Nifty and Midcap indices, using the VAR system obtained for each data frequency. VAR models may be used to generate Impulse Response Functions (IRFs) as shown for instance by Sims (1980). IRFs trace the response of each of the variables in the system, to a shock in a given variable. Usually, a shock to the variable x_t has the largest effect on subsequent realisations of that variable itself. Yet if the VAR model predicts that the variable x_t is affecting other variables in the system, the latter may also respond to the initial shock in x_t – usually the latter responses are lower than the subsequent fluctuation in x_t due to the shock.

The IRF of variable y_t to a shock in variable x_t which occurs at time t, may be viewed as the difference between two time series:

- The realisations of the time series y_t after the shock in x_t has occurred; and
- The realisations of the time series y_t during the same time period, in the absence of the shock in x_t.

The above may be formulated in the following mathematical notation:

$$IRFy(n,\delta,\omega_{t-1}) = E[y_{t+n}|\varepsilon_t = \delta, \varepsilon_{t+1} = ... = \varepsilon_{t+n} = 0, \omega_{t-1}] - E[y_{t+n}|\varepsilon_t = 0, \varepsilon_{t+1} = ... = \varepsilon_{t+n} = 0, \omega_{t-1}] \quad (11)$$

where y is the time series of interest, δ is the shock taking place at time t, ω_{t-1} is the history of the time series, ε denotes an innovation, and the IRF is generated from time t to $t+n$. As shown in the equation, the intermediate innovations are assumed to be equal to zero, and this may be considered as a limitation of the above (orthogonalised version) of the IRF. Another weakness of orthogonalised IRFs is that the actual results may vary, depending on the order in which the variables are specified in the VAR model.

In order to overcome the above limitations, Koop, Pesaran and Potter (1996) introduced the Generalised Impulse Response Function (GIRF), where the expectations of y_{t+n} are only conditional on the shock and the history of the series. In mathematical notation:

$$GIRFy(n,\delta,\omega_{t-1}) = E[y_{t+n}|\varepsilon_t = \delta, \omega_{t-1}] - E[y_{t+n}|\omega_{t-1}] \quad (12)$$

The empirical results are shown in Section 6.1 (Nifty shocks) and Section 6.2 (Midcap shocks).

6.1 Responses to Nifty Shocks

Both the Orthogonalised and Generalised IRF versions showed that a given shock in the Nifty log return series effects the Midcap log return in the same direction. The Orthogonalised and Generalised IRFs yielded the same qualitative results, and therefore only the latter diagrams are being shown in Figure 5. Appendix J Panel A and Appendix K Panel A show the four IRF statistics (orthogonalised and generalised versions, for daily and one-minute frequencies respectively).

In case of the daily data sets, the shocks practically die out within one day. This is in line with the above notions that no high degree of causality may be detected using this sampling frequency. In case of the high-frequency data sets, the shocks die out approximately within 10 minutes, since they are based on a VAR system of order 9.

For both data sets, when considering the latter part of the effect (i.e. just before the shocks die out) the effect on the Midcap is larger than that on the Nifty itself. In case of the high-frequency data set this may be attributed to the "lead-lag" relationship which partly emanates from non-synchronous trading effects discussed above. In case of the daily sets, the larger effect on the Midcap seems visible right from the initial part of the IRF.

The observation that a shock in variable x has a greater impact on variable y rather than on itself is rather unusual in empirical exercises involving IRFs. This may suggest that any actual causality effects which run from Nifty to Midcap are also amplified by non-synchronous trading effects. Yet, there might also be an additional explanation as to why a Nifty shock might lead to a higher Midcap shock, and this relates to the risks of the stocks. If we assume that the smaller companies are more risky, they should have higher *betas* and therefore should fluctuate more widely as compared to the larger company stocks. Thus, a given news item may in fact have a higher impact on the Midcap rather than on the Nifty index. Given that this research does not account for news releases the results present an "illusion" that a Nifty shock results in a larger shock in the Midcap; yet the latter movement might in fact be a response to news rather than the Nifty shock itself.

6.2 Responses to Midcap Shocks

We now turn to the responses following a shock in the Midcap index. This time, the Orthogonalised and Generalised versions of the IRFs yielded differing results. Figure 6 shows the Orthogonalised IRFs whilst Figure 7 shows the Generalised IRFs. Actual statistics are respectively shown in Appendix J Panel B (daily data set) and Appendix K Panel B (high frequency data set).

The Orthogonalised IRFs show that the Nifty index is practically unresponsive to Midcap shocks, especially when considering the daily data set.

The Generalised IRFs seem to indicate that the Nifty index responds to Midcap shocks. A shock in the Midcap leads to a Nifty fluctuation in the same direction which dies out after one day in case of the daily data set, and before 10 minutes in case of the high-frequency data set.

Given that Midcap stocks trade less frequently than Nifty stocks as discussed above, we cannot attribute any predictability that runs from Midcap to Nifty to non-synchronous trading effects. Overall, the IRFs seem to confirm that there may be a feedback effect from the Midcap to the Nifty, although this does not seem particularly pronounced. Possible reasons why a small capitalisation index may affect the main index were discussed in Section 4 above.

Overall, the IRFs (together with the previous tests) suggest that the lead-lag effects which are present in the data sets are a combination of actual "causality" or delayed adjustments to news, and non-synchronous trading effects.

7. Inefficiency or Non-Synchronous Trading Effects?

Considering the above three predictability investigations between the Nifty and Midcap indices, one observation which was consistently confirmed is that at high frequency intervals the Nifty index unambiguously leads the Midcap index. The "causality" from Nifty to the Midcap may be explained both by non-synchronous trading arguments and by the possibility that market participants do not adjust their expectations immediately since they do not follow lower-capitalisation companies as closely. A further elaboration on this argument may be inferred from the research of Niarchos and Alexakis (1998) in relation to the Greek Stock market. The authors argued that foreign investors tend to restrict their holdings in a particular category of shares. Since foreign investors are typically more sophisticated the latter share category is more efficient.[6] Therefore, if in the case of NSE overseas investors restrict their holdings to the shares in the main index (which is quite plausible), one would expect the Nifty to be more efficient as compared to the Midcap. Yet, from the trading frequency statistics presented in Tables 3 & 4, we may also deduce that part of this "predictability" is in fact related to non-synchronous trading.

The analysis of trading break and post-trading break returns is relevant for inferring whether delayed price adjustments in the data set, mainly emanate from traders' delays in adjusting their expectations or whether they are more attributable to non-synchronous trading. We may assume that during a trading break, market participants have enough time to adjust their judgements regarding the fundamental value of the firms, and that any outdated limit orders are cancelled. In this investigation, the term "trading break" refers to the ceasing of trading activity at the end of the day till the subsequent morning (and at times till after the weekend).

[6] Similar conclusions that foreign investors contribute more actively to market efficiency were empirically discovered by Tian and Wan (2004) in an investigation of the Chinese and Hong Kong share market. Potentially contrasting evidence was presented by Panagiotidis (forthcoming), who found that upon the adoption of the Euro, the main indices of the Athens Stock Exchange still did not follow a random walk.

In this way, the former assumption that the trading break provides enough time for participants to adjust their expectations is reasonable.

Given this, one may assume that any trades which occur immediately after a trading break will reflect the underlying market value of the particular firms and we may rule out any delayed price adjustments on part of the traders. This implies that if we note that lead-lag effects in between the indices persist in the post-trading break data, they are mainly due to non-synchronous trading effects rather than mispriced trades. Non-synchronous trading effects can still coexist with trading breaks given that an infrequently traded stock, might still take longer to trade following the trading break. This results in a delayed adjustment of market price data – yet it is reasonable to assume that traders' expectations would have already been adjusted by the time that trading was resumed.

Thus we now turn to test which of the above effects is the main cause of the predictability by applying the above methodology. The VAR model shown in Section 5.2 indicates that the first three and the sixth (one-minute) Nifty lags are significant in determining the value of the Midcap, in case of the high frequency data set. Therefore we look at the Nifty overnight returns, the Midcap overnight returns and the Midcap returns during the first six minutes of the trading day.

Now, if the Nifty overnight returns can explain the Midcap overnight returns, we may argue that the lagged Midcap response observed above consists of inherent inefficiency. This follows the argument that in the course of the trading day, participants adjust their judgement regarding the fundamental value of the Midcap stocks with a six-minute delay, for possible reasons which were already discussed. The situation changes somewhat in case of overnight returns. We may assume that if the traders adjust the Midcap fundamental value with a six-minute delay, most of the adjustment should still appear in the Midcap overnight fluctuation, given that the non-trading period is much longer than six minutes. In this way, the Nifty overnight return should correspond with the Midcap overnight return.

Conversely, if we note that the Nifty overnight returns can explain the Midcap return for the first six-minutes of the following trading day, we may attribute this to non-synchronous trading. This follows the notion that whilst traders adjust their Midcap expectations overnight (prior to the start of the subsequent trading session), it still takes six minutes for these adjustments to get reflected in the prices – since it takes about this time for sufficient transactions to take place in the less liquid stocks to adjust the actual trading prices.

Therefore the following OLS regression is estimated:

$$OR(N)_{t\to t+1} = \alpha + OR(M)_{t\to t+1} + IR(M)_{t+1} + \varepsilon \qquad (13)$$

where $OR(N)_{t\to t+1}$ is the Nifty overnight (log) return between day t and $t+1$, $OR(M)$ is the Midcap overnight (log) return, $IR(M)$ is the Midcap initial six-minute (log) return, α is a

constant, and ε is an error term. Usually the dependent variable of the regression is shown on the left hand side, yet given that in this case there are two (possibly) dependent variables, they are shown on the right hand side. This does not affect the estimated coefficients given that OLS regression does not imply any inferences about causality, and therefore the methodology is equally valid.

The sample period consisted of 112 observations ranging from 11th June 1999 to 16th November 1999. These dates were deliberately chosen to obtain a sample period where no initial call auctions were held. Yet, the data set featured one missing observation in the Midcap return for the first six minutes given that the intra-day file for the 22nd September 1999 was unavailable. Thus the sample period was split up, and two regressions were estimated. The results of the two regressions were qualitatively the same, and they are being reproduced in Table 8. The regressions have a comprehensive explanatory power in terms of R^2 and R-bar-squared. Both regressions indicate that the Nifty Overnight Return is more correlated with the Initial Midcap Return of the subsequent trading day rather than the contemporaneous Overnight Midcap Return. The former variable is highly significant, whilst the latter is insignificant. This lead-lag relationship is unlikely to be attributable to traders delaying the adjustment of their judgment regarding the value of the securities, since during the overnight period one may assume that participants have ample time to adjust expectations!

Thus, the lead-lag relationship from Nifty to Midcap at high-frequency data is more attributable to non-synchronous trading effects. Yet, the lead-lag effect which runs from Midcap to Nifty (which cannot be attributed to non-synchronous trading) implies that part of this predictability effect constitutes of an actual lead-lag relationship. Possible economic explanations for such a relationship might be spillover effects amongst stocks, though one cannot rule out the possibility that such an effect is a mere coincidence.

The above results are in line with the discussion by Atchison, Butler and Simonds (1987) that stock indices tend to exhibit higher autocorrelation than that which may be expected from non-synchronous trading effects. Similar conclusions were presented by Lo and MacKinlay (1990). The authors argued that portfolios of smaller stocks usually feature high autocorrelation levels which may not be explained by non-synchronous trading alone, and therefore one cannot rule out the presence of actual lead-lag effects running from larger to smaller stocks, in addition to non-synchronous trading effects.

The latter investigation thus yields an important contribution as regards the interpretation of predictability. Predictability and Granger-Causality effects do not necessarily imply actual causality. This does not simply reflect the possibility that the analysed time series may be responding to an exogenous variable. Absence of actual causality may also be due to non-synchronous trading which results in less liquid stocks (apparently) taking longer to adjust to new information.

The above discussion also relates to the issue of market efficiency. In particular, predictability effects in the data do not necessarily imply market inefficiency. In the above context, whilst one may attempt to partly predict the Midcap value from the lagged Nifty values particularly when using high frequency data, this does not necessarily translate into profitable opportunities. Although the Midcap value might be temporarily mispriced being calculated through stock prices prevailing in past trades, it does not mean that traders are still prepared to transact at such "outdated" prices. In addition, there is also a possibility of observing some transitory transactions at "outdated" prices in an efficient market. This rests on the scenario that following the release of news, some limit orders which were submitted prior to the news release are not cancelled immediately and efficient traders "pick off" these orders. This would result in mispriced transactions – yet it does not mean that the market is inefficient.

8. Conclusion

One of the main branches of research in market microstructure relates to the area of market efficiency. As discussed above, non-synchronous trading effects can give the impression that market participants are adjusting their expected value of the traded assets following a delay. In this way non-synchronous trading effects may lead to flawed inferences as regards market efficiency. This investigation shows that predictability in stock prices does not necessarily contradict market efficiency. Factors such as non-synchronous trading and market participants "picking off" mispriced orders following the arrival of new information, may give the impression that the traders are not adjusting their expectations immediately. This confirms the notion that the main criterion for an inefficient market is the existence of profitable trading opportunities and not predictability.

This research has investigated the lead-lag effects in NSE stock price data using different methodologies. The main empirical observation is that the Nifty index leads the Midcap index – particularly when considering a high frequency data set. When analysing the trading break returns it was noted that such lead-lag effects persist. In the latter case, such predictability cannot be attributed to traders' delayed expectation adjustments, since during an overnight period market participants have sufficient time to adjust their expected values of stocks. Thus we may conclude that lead-lag effects are mainly caused by non-synchronous trading, and that this predictability is not likely to result in abnormal profit opportunities. Yet, in line with previous studies, we may also note that non-synchronous trading is not the exclusive cause of such predictability. In particular, the feedback effect from the Midcap to the Nifty index may not be attributed to non-synchronous trading given that the former index is composed of less liquid stocks as compared to the latter.

The main contributions of this investigation to the literature are three-fold. Firstly, it was shown how non-synchronous trading effects may be detected through lead-lag effects in stock prices, whereas previous studies tended to focus on the serial correlation structure. Secondly,

the investigation proposed a simple methodology constituting of the analysis of trading break and post-trading break returns in order to infer whether predictability is more attributable to non-synchronous trading as opposed to actual delayed price adjustments.

Thirdly, it was formally shown that non-synchronous trading effects tend to become more pronounced in a high-frequency data set. This partly rests on the tendency for the level of trading activity to vary throughout the trading day. This is also in line with the empirical evidence of Papachristou (1999), who compared daily frequency serial correlation to weekly frequency serial correlation and found that non-synchronous trading effects are more evident in higher frequency data. This implies that researchers aiming to study non-synchronous trading effects stand a better chance of obtaining significant empirical evidence if they use a high-frequency set. Conversely, researchers who use high-frequency data in investigating unrelated issues should consider the possibility that their empirical research may be biased due to such effects. For instance, non-synchronous trading effects should be taken into consideration when evaluating the adequacy of a given market setup in terms of its efficiency and volatility.

Another appealing feature of this analysis is the application of non-synchronous trading concepts in the context of an emerging market. Non-synchronous trading might be even more relevant to emerging markets, given that such markets are often found to be less liquid.

This investigation also suggests further research issues. Firstly, one potential avenue might lie in the re-interpretation of previous studies. Most of the latter studies were based on daily data. In this way, the likely increases in trading activity at the end of the day probably diminished non-synchronous trading effects, and this might amount to an under-estimation of the effects of non-synchronous trading on stock price data. Another potential research issue lies in the specific investigation of individual stock price high-frequency data. Finally, following the notion that trading activity varies throughout the trading day, one may also inquire how non-synchronous trading effects become more pronounced during the middle of the day when trading activity tends to abate.

References

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

Baek, E., and W. Brock, 1992, A General Test for Nonlinear Granger Causality: Bivariate Model, *Working Paper*, Iowa State University and University of Wisconsin-Madison.

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

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

Capelle-Blancard, G., and H. Raymond, 2002, What Is Behind International Stock Index Linkages?, *Working Paper: Université Paris - Université Lille*.

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

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

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

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

Fisher, L., 1966, Some New Stock Market Indexes, *Journal of Business*, 39, 191-225.

Granger, C.W.J., 1969, Investigating Causal Relations By Econometric Methods And Cross Spectral Methods, *Econometrica*, 37(3), 424-438.

Hiemstra, C., and J.D. Jones, 1994, Testing For Linear And Nonlinear Granger Causality In The Stock Price-Volume Relation, *The Journal of Finance*, 49(5), 1639-1664.

Kadlec, G.B., and D.M. Patterson, 1999, A Transactions Data Analysis of Nonsynchronous Trading, *The Review of Financial Studies*, 12(3), 609-630, Fall.

Koop, G., M.H. Pesaran and S.M. Potter, 1996, Impulse Response Analysis In Nonlinear Multivariate Models, *Journal of Econometrics*, 74(1), 119-147.

Lamba, A.S., 2003, An Analysis Of The Dynamic Relationships Between South Asian And Developed Equity Markets, *Unpublished Working Paper: NSE Research Initiative.*

Lo, A.W., and A.C. MacKinlay, 1990, When Are Contrarian Profits Due To Stock Market Overreaction, *The Review of Financial Studies*, 3(2), 175-205.

Mills, T.C., and J.V. Jordanov, 2000, Lead-Lag Patterns Between Small and Large Size Portfolios in the London Stock Exchange, *Applied Financial Economics*, 11(5), 489-495.

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

Panagiotidis, T., forthcoming, Market capitalization and efficiency. Does it matter? Evidence from the Athens Stock Exchange, *Applied Financial Economics*.

Papachristou, G., 1999, Stochastic behaviour of the Athens Stock Exchange: a case of institutional nonsynchronous trading, *Applied Financial Economics*, 9, 239-250.

Pesaran, M.H., and A. Timmermann, 1992, A Simple Non-Parametric Test Of Predictive Performance, *Journal of Business and Economic Statistics*, 10(4), 461-465.

Renault, E., K. Sekkat, and A. Szafarz, 1998, Testing For Spurious Causality In Exchange Rates, *Journal of Empirical Finance*, 5, 47-66.

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

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

Sims, C., 1980, Macroeconomics and Reality, *Econometrica*, 48, 1-48.

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

Wood, R.A., T.H. McInish, and J.K. Ord, 1985, An Investigation of Transactions Data for NYSE Stocks, *Journal of Finance*, 40(3), 723-739.

Table 1: Daily Interval Data Set

Properties of Nifty and Midcap Logarithmic Series (LN and LM) and Nifty and Midcap Log Returns (LRN and LRM)

Variable	LN	LM	LRN	LRM
Maximum Value	7.539	8.534	0.075	0.074
Minimum Value	6.750	6.953	-0.077	-0.090
Mean	7.083	7.565	0.001	0.001
Std. Deviation	0.167	0.345	0.016	0.020
Skewness	0.438	0.584	-0.174	-0.547
Excess Kurtosis	-0.743	-0.321	2.614	1.972
Jarque-Bera Test	69.2	76.9	363.8	266.1
ADF Test Statistics (Excluding Trend):				
ADF (1)	-1.11	-0.71	-25.12	-23.59
ADF (5)	-1.18	-0.76	-14.15	-13.68
ADF (10)	-1.30	-1.07	-10.15	-8.75
ADF Test Statistics (Including Trend):				
ADF (1)	-1.00	-0.40	-25.13	-23.60
ADF (5)	-1.08	-0.47	-14.16	-13.71
ADF (10)	-1.20	-0.86	-10.17	-8.79

The table shows the distributional properties and ADF test statistics for the logs and log returns of Nifty and Midcap daily series. The ADF tests do not reject the null hypothesis of a unit root for the logarithmic series. The null hypothesis of a unit root is rejected when considering the logarithmic returns.

Table 2: One-Minute Interval Data Set

Properties of Nifty and Midcap Logarithmic Series (LN and LM) and Nifty and Midcap Log Returns (LRN and LRM)

Variable	LN	LM	LRN	LRM
Maximum Value	7.113	7.592	0.010	0.004
Minimum Value	7.012	7.522	-0.006	-0.003
Mean	7.070	7.560	0.0000	0.0000
Std. Deviation	0.024	0.015	0.0006	0.0005
Skewness	-0.804	-0.167	2.624	0.579
Excess Kurtosis	-0.319	-0.344	39.803	7.278
Jarque-Bera Test	332.5	28.4	199,395	6,719
ADF Test Statistics (Excluding Trend):				
ADF (1)	-1.63	-1.06	-31.93	-32.29
ADF (5)	-1.65	-1.26	-20.80	-18.29
ADF (10)	-1.63	-1.28	-15.72	-15.13
ADF Test Statistics (Including Trend):				
ADF (1)	-0.91	-1.63	-31.96	-32.29
ADF (5)	-1.01	-1.91	-20.84	-18.28
ADF (10)	-0.93	-1.94	-15.78	-15.13

The table shows the distributional properties and ADF test statistics for the logs and log returns of Nifty and Midcap series sampled at one minute frequencies. The ADF tests do not reject the null hypothesis of a unit root for the logarithmic series. The null hypothesis of a unit root is rejected when considering the logarithmic returns.

Table 3: Average Trading Frequencies for Nifty and Midcap Shares

	Nifty Shares		Midcap Shares	
	Av. # Trades per Share	Av. Waiting Time (seconds)	Av. # Trades per Share	Av. Waiting Time (seconds)
15-Jun-99	3181	6.2	956	20.7
17-Jun-99	3798	5.2	858	23.1
21-Jun-99	3081	6.4	1034	19.1
23-Jun-99	3568	5.5	989	20.0
25-Jun-99	3180	6.2	965	20.5

The table shows the average number of transactions for the shares included in the Nifty and Midcap indices for five different trading days. Assuming that these transactions occur evenly throughout a trading day of five and a half hours, we can estimate the average waiting time – i.e. the average interval between trades.

Table 4: Average Trading Frequencies for Nifty and Midcap Least Frequently Traded Shares

	Nifty (Least Traded) Shares		Midcap (Least Traded) Shares	
	Av. # Trades per Share	Av. Waiting Time (minutes)	Av. # Trades per Share	Av. Waiting Time (minutes)
15-Jun-99	73	5	42	8
17-Jun-99	66	5	26	13
21-Jun-99	73	5	31	11
23-Jun-99	86	4	51	7
25-Jun-99	71	5	32	10

The table shows the average number of transactions for the ten least frequently traded shares included in the Nifty and Midcap indices. Statistics for five different trading days are shown. Assuming that these transactions occur evenly throughout a trading day of five and a half hours, we can estimate the average waiting time – i.e. the average interval between trades.

Table 5: Pesaran-Timmermann Statistics

	Daily Data				Intra-Day Data			
Lag (a)	$S(M_t,N_{t-a})$		$S(N_t,M_{t-a})$		$S(M_t,N_{t-a})$		$S(N_t,M_{t-a})$	
0	20.28	***	20.28	***	18.83	***	18.83	***
1	3.99	***	3.65	***	13.41	***	11.61	***
2	-0.31		0.60		8.18	***	3.37	***
3	1.60		2.97	***	5.96	***	0.34	
4	1.58		1.69		4.47	***	0.87	
5	0.19		0.30		5.19	***	0.52	
6	-1.72		-1.49		3.92	***	-1.08	
7	-0.55		-0.44		1.95	*	0.48	
8	-0.35		1.07		0.35		-1.74	
9	1.16		1.21		-0.36		-2.42	
10	1.36		-0.01		0.90		-1.56	
11	0.99		0.13		0.67		-2.49	
12	0.91		0.73		-0.31		-0.89	
13	0.31		0.82		0.23		-1.10	
14	2.06	**	-0.75		-0.71		-0.79	
15	0.37		-0.56		-0.95		-0.51	
16	0.58		0.45		-1.82		-1.63	
17	0.15		0.82		-0.11		-0.99	
18	-0.17		0.85		-1.52		-1.56	
19	-1.06		-0.04		-0.14		-1.25	
20	-0.91		0.22		-0.71		-1.75	
21	-1.28		-0.15		0.82		0.11	
22	-0.56		-0.47		1.06		-0.68	
23	0.50		0.37		0.27		-0.66	
24	-0.50		1.43		1.32		0.72	
25	0.50		0.60		1.00		1.59	
26	0.02		-1.15		1.61		2.01	**
27	1.03		0.49		1.12		3.51	***
28	-0.44		0.52		1.87	*	0.42	
29	-0.01		-0.32		0.27		-1.33	
30	-0.44		-0.06		0.84		-0.87	

The table shows Pesaran-Timmermann Test Statistics (S) for the relationship between the Nifty Index (N) and the Midcap Index (M). Both contemporaneous and lagged relationships are investigated. Significance at the 99%, 95% and 90% level of confidence is denoted by ***, ** and * respectively.

Table 6: Nifty and Midcap Regression Coefficients for Daily Data VAR

	Nifty Regression			Midcap Regression		
Regressor	Coefficient	Standard Error	T-Ratio	Coefficient	Standard Error	T-Ratio
LRN(-1)	0.058	0.046	1.242	-0.024	0.058	0.417
LRM(-1)	0.014	0.036	0.374	0.168	0.046	3.669
CONST	0.001	0.000	1.214	0.001	0.001	0.941
F-Statistic	F(2, 1252): 3.2571			F(2, 1252): 15.1028		

The first column shows the regressor, where LRN and LRM stand for Nifty and Midcap Log Return, whilst Const is the intercept of the regression. The first lag of a variable is denoted as (-1). For both the Nifty and Midcap equations, the table shows the regression coefficients, standard errors and T-ratios. The F-statistics shown in Appendices E and F reject the null hypothesis that the estimated coefficients are equal to zero.

Table 7: Nifty and Midcap Regression Coefficients for VAR estimated through high-frequency data.

	Nifty Regression			Midcap Regression		
Regressor	Coefficient	Standard Error	T-Ratio	Coefficient	Standard Error	T-Ratio
LRN(-1)	0.203	0.019	10.507	0.241	0.015	16.260
LRN(-2)	0.012	0.020	0.606	0.077	0.015	4.962
LRN(-3)	0.010	0.020	0.500	0.075	0.016	4.842
LRN(-4)	-0.010	0.020	-0.487	0.019	0.016	1.210
LRN(-5)	0.013	0.020	0.636	0.023	0.016	1.465
LRN(-6)	-0.030	0.020	-1.496	0.034	0.016	2.179
LRN(-7)	0.048	0.020	2.368	0.011	0.016	0.720
LRN(-8)	0.000	0.020	-0.005	-0.045	0.015	-2.898
LRN(-9)	0.015	0.020	0.771	0.039	0.015	2.585
LRM(-1)	0.162	0.027	6.061	0.027	0.020	1.306
LRM(-2)	-0.036	0.027	-1.349	-0.080	0.020	-3.918
LRM(-3)	0.041	0.027	1.529	0.031	0.020	1.522
LRM(-4)	-0.028	0.027	-1.055	-0.011	0.020	-0.530
LRM(-5)	0.033	0.027	1.256	-0.010	0.020	-0.502
LRM(-6)	-0.028	0.027	-1.064	0.031	0.020	1.514
LRM(-7)	-0.078	0.026	-2.973	-0.042	0.020	-2.065
LRM(-8)	-0.002	0.026	-0.064	0.036	0.020	1.783
LRM(-9)	-0.032	0.025	-1.282	0.011	0.019	0.554
CONST	0.000	0.000	-0.451	0.000	0.000	-0.692
O	0.003	0.000	19.179	0.000	0.000	1.293
F-Statistic	F(19, 2940): 36.3434			F(19, 2940): 31.0200		

The first column shows the regressor, where LRN and LRM stand for Nifty and Midcap Log Return, Const is the intercept of the regression, whilst O is a dummy variable which takes the value of 1 for the first 2 observations of the trading day and zero otherwise. Lags are denoted as (-1), (-2), etc. For both the Nifty and Midcap equations, the table shows the regression coefficients, standard errors, and T-ratios. The F-statistics shown in Appendices H and I strongly reject the null hypothesis that the estimated coefficients are equal to zero.

Table 8: Nifty Overnight Return Regressions [$OR(N)_{t\rightarrow t+1}$]

Sample Period:	11-Jun-1999 to 21-Sep-1999			23-Sep-1999 to 16-Nov-1999		
# Observations	72			39		
	Coeff.	Std. Error	T-Ratio	Coeff.	Std. Error	T-Ratio
α	0.00004	0.00006	0.61879	0.00022	0.00011	1.98830
$OR(M)_{t\rightarrow t+1}$	0.05813	0.05581	1.04160	-0.08314	0.13644	-0.60939
$IR(M)_{t+1}$	0.03733	0.00846	4.41180	0.02847	0.00918	3.10090
R-squared		0.32683			0.24793	
R-bar-squared		0.30732			0.20615	

The above table shows the results for the Nifty Overnight Return [$OR(N)_{t\rightarrow t+1}$] Regressions. The regressors were the Midcap Overnight Return between day t and t+1 [$OR(M)_{t\rightarrow t+1}$], the Midcap Initial Return during the first six minutes of the subsequent trading day [$IR(M)_{t+1}$], and an intercept α. The sample period was split into two due to a missing observation. In both regressions, $OR(N)_{t\rightarrow t+1}$ is more correlated with $IR(M)_{t+1}$. This is an indication that the lead-lag relationship between the Nifty and the Midcap is more related to non-synchronous trading, rather than a delayed adjustment of traders' expectations.

Figure 1: Logarithmic Series for Nifty (LN) and Midcap (LM) (Daily Frequency)

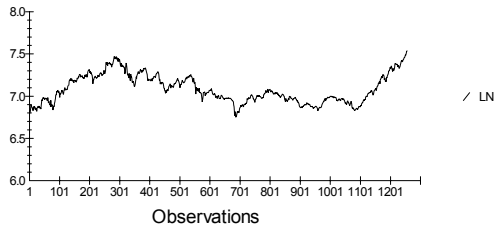

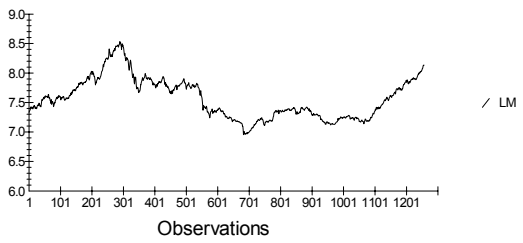

Figure 2: Log Return Series for Nifty (LRN) and Midcap (LRM) (Daily Frequency)

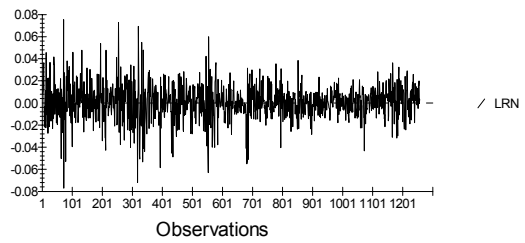

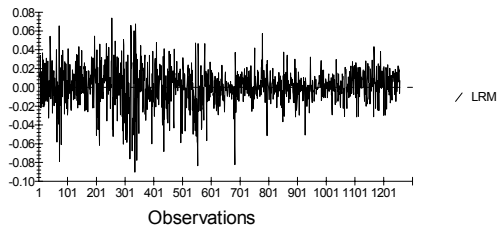

Figure 3: Logarithmic Series for Nifty (LN) and Midcap (LM) (One Minute Frequency)

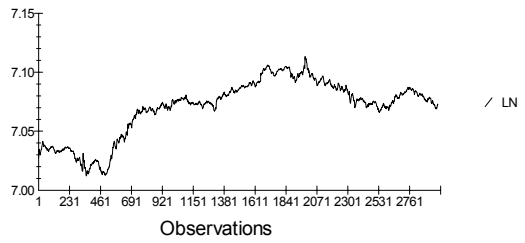

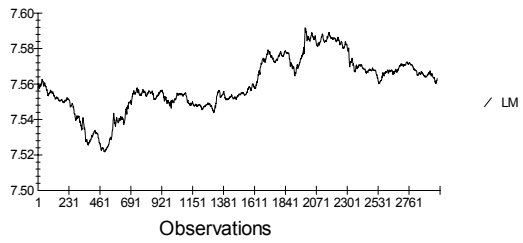

Figure 4: Log Return Series for Nifty (LRN) and Midcap (LRM) (One Minute Frequency)

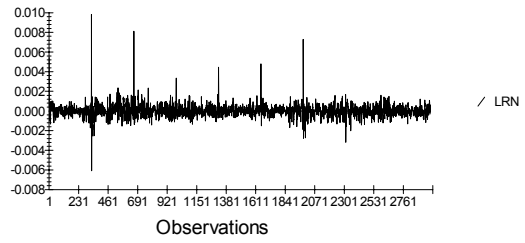

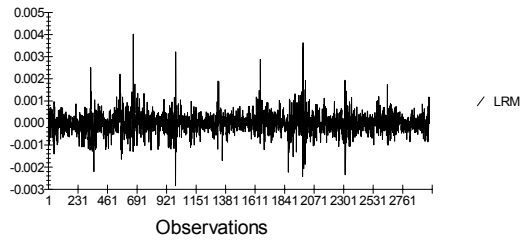

Figure 5: Generalised IRFs following a Nifty Shock

Daily Data Set

Generalised Impulse Responses to one SE shock in the equation for LRN

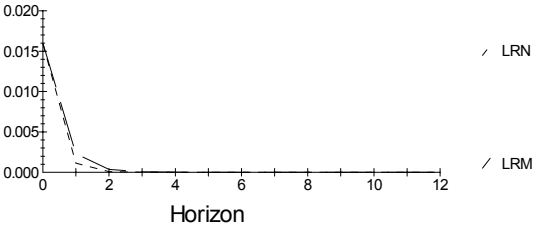

High-Frequency Data Set

Generalised Impulse Responses to one SE shock in the equation for LRN

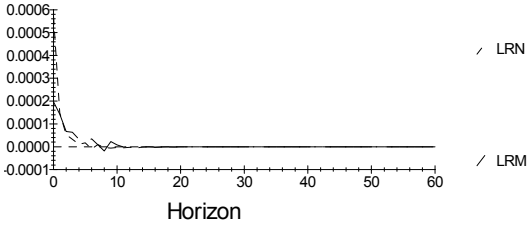

Figure 6: Orthogonalised IRFs following a Midcap Shock

Daily Data Set

Orthogonalised Impulse Responses to one SE shock in the equation for LRM

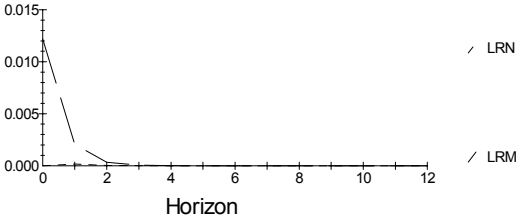

High-Frequency Data Set

Orthogonalised Impulse Responses to one SE shock in the equation for LRM

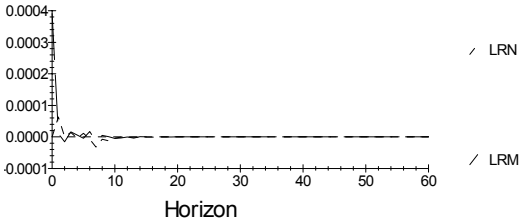

Figure 7: Generalised IRFs following a Midcap Shock

Daily Data Set

Generalised Impulse Responses to one SE shock in the equation for LRM

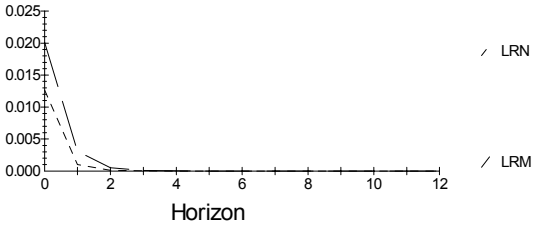

High-Frequency Data Set

Generalised Impulse Responses to one SE shock in the equation for LRM

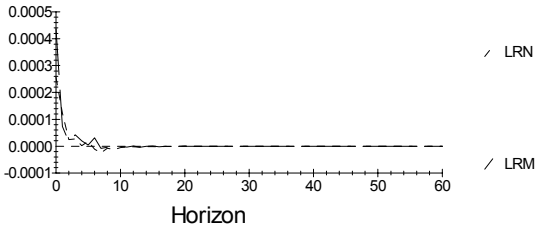

Appendix A: Serial Correlation Statistics For The First Five Lags of Log Nifty, Log Midcap, Log Return Nifty and Log Return Midcap.

Order	Autocorrelation Coefficient	Std. Error	Box-Pierce Statistic	Ljung-Box Statistic
Log Nifty (Daily Frequency)				
1	0.991	0.028	1235.2	1238.2
2	0.982	0.049	2447.4	2454.3
3	0.973	0.062	3637.8	3649.4
4	0.965	0.073	4807.6	4824.7
5	0.956	0.083	5956.9	5980.5
Log Midcap (Daily Frequency)				
1	0.997	0.028	1249.5	1252.5
2	0.994	0.049	2490.2	2497.2
3	0.990	0.063	3722.3	3734.2
4	0.987	0.074	4945.7	4963.4
5	0.983	0.084	6160.4	6184.9
Log Return Nifty (Daily Frequency)				
1	0.071	0.028	6.361	6.376
2	-0.033	0.028	7.727	7.747
3	0.003	0.028	7.741	7.761
4	0.045	0.028	10.258	10.289
5	0.018	0.028	10.674	10.708
Log Return Midcap (Daily Frequency)				
1	0.153	0.028	29.404	29.475
2	-0.003	0.029	29.418	29.488
3	0.035	0.029	31.000	31.077
4	0.015	0.029	31.266	31.344
5	0.016	0.029	31.570	31.649

Order	Autocorrelation Coefficient	Std. Error	Box-Pierce Statistic	Ljung-Box Statistic
Log Nifty (One-Minute Frequency)				
1	0.999	0.018	2964.8	2967.8
2	0.998	0.032	5924.1	5931.1
3	0.997	0.041	8877.6	8889.6
4	0.996	0.048	11825.1	11843.1
5	0.995	0.055	14766.7	14791.6
Log Midcap (One-Minute Frequency)				
1	0.999	0.018	2967	2970
2	0.999	0.032	5929.7	5936.7
3	0.998	0.041	8887.7	8899.7
4	0.997	0.048	11840.2	11858.2
5	0.996	0.055	14786.8	14811.7
Log Return Nifty (One-Minute Frequency)				
1	0.272	0.018	219.470	219.692
2	0.097	0.020	247.344	247.603
3	0.052	0.020	255.419	255.692
4	0.016	0.020	256.179	256.453
5	0.015	0.020	256.876	257.152
Log Return Midcap (One-Minute Frequency)				
1	0.222	0.018	145.853	146.000
2	0.093	0.019	171.496	171.678
3	0.126	0.019	218.941	219.203
4	0.065	0.020	231.639	231.926
5	0.039	0.020	236.157	236.455

Appendix B: Selecting the Order of the VAR for Daily Data

The table below shows AIC, SBC and LR test statistics for a preliminary 24 order VAR which was estimated on the log returns series of Nifty and Midcap. Both the Akaike Information Criterion and the Schwarz Bayesian Criterion select a VAR(1) model, yet the log-likelihood ratio statistics reject all orders less than 16. Therefore, both a VAR(1) and a VAR(16) model were estimated, and their diagnostics were compared as shown in Appendix N.

Test Statistics and Choice Criteria for Selecting the Order of the VAR Model
Based on 1232 observations. Order of VAR = 24
Variables included in the unrestricted VAR: LRN, LRM
Deterministic and/or exogenous variables: CONST

Order	LL	AIC	SBC	LR test	Adjusted LR test
24	7109.9	7011.9	6761.2	------	------
23	7109.1	7015.1	6774.7	CHSQ(4)= 1.5566[.817]	1.4947[.828]
22	7104.0	7014.0	6783.8	CHSQ(8)= 11.8435[.158]	11.3724[.181]
21	7102.4	7016.4	6796.4	CHSQ(12)= 15.1076[.236]	14.5067[.270]
20	7102.1	7020.1	6810.4	CHSQ(16)= 15.5466[.485]	14.9282[.530]
19	7100.1	7022.1	6822.5	CHSQ(20)= 19.6788[.478]	18.8961[.529]
18	7099.1	7025.1	6835.8	CHSQ(24)= 21.6473[.600]	20.7863[.651]
17	7091.1	7021.1	6842.1	CHSQ(28)= 37.5427[.107]	36.0495[.141]
16	7087.9	7021.9	6853.0	CHSQ(32)= 44.0813[.076]	42.3281[.105]
15	7083.4	7021.4	6862.8	CHSQ(36)= 53.0531[.033]	50.9430[.051]
14	7081.6	7023.6	6875.2	CHSQ(40)= 56.7236[.042]	54.4676[.063]
13	7075.2	7021.2	6883.0	CHSQ(44)= 69.5188[.008]	66.7539[.015]
12	7068.3	7018.3	6890.4	CHSQ(48)= 83.3006[.001]	79.9875[.003]
11	7066.9	7020.9	6903.2	CHSQ(52)= 86.1284[.002]	82.7029[.004]
10	7065.5	7023.5	6916.1	CHSQ(56)= 88.8453[.003]	85.3117[.007]
9	7058.9	7020.9	6923.7	CHSQ(60)= 102.0775[.001]	98.0176[.001]
8	7055.6	7021.6	6934.6	CHSQ(64)= 108.6157[.000]	104.2958[.001]
7	7053.6	7023.6	6946.9	CHSQ(68)= 112.5511[.001]	108.0746[.001]
6	7052.7	7026.7	6960.2	CHSQ(72)= 114.4210[.001]	109.8701[.003]
5	7051.0	7029.0	6972.7	CHSQ(76)= 117.7967[.002]	113.1116[.004]
4	7050.1	7032.1	6986.0	CHSQ(80)= 119.6825[.003]	114.9224[.006]
3	7046.6	7032.6	6996.8	CHSQ(84)= 126.5716[.002]	121.5375[.005]
2	7041.7	7031.7	7006.1	CHSQ(88)= 136.3747[.001]	130.9507[.002]
1	7039.0	7033.0	7017.7	CHSQ(92)= 141.7437[.001]	136.1062[.002]
0	7014.5	7012.5	7007.4	CHSQ(96)= 190.8167[.000]	183.2274[.000]

AIC=Akaike Information Criterion SBC=Schwarz Bayesian Criterion

Appendix C: Selecting the Order of the VAR for Daily Data

The table shows various explanatory power statistics for the VAR(1) and the VAR(16) models. The VAR(1) model was selected on the basis of higher System Log Likelihood Ratio, higher Log Likelihood Ratios for the individual equations, higher F-statistics, and higher Akaike Information Criterion and Schwarz Bayesian Criterion. Besides, it is not clear on practical grounds why an index shock taking place on day *t*, should still affect the price series on day *t+16*.

VAR (1)	**Nifty Regression**	**Midcap Regression**
System Log-likelihood	7168.3	7168.3
Equation Log-likelihood	3416	3124.4
F-Statistic F(2,1252)	3.2571	15.1028
R-Bar-Squared	0.0036	0.0220
Akaike Info. Criterion	3413	3121.4
Schwarz Bayesian Criterion	3405.3	3113.7
VAR (16)	**Nifty Regression**	**Midcap Regression**
System Log-likelihood	7135.8	7135.8
Equation Log-likelihood	3399.8	3111.7
F-Statistic F(32,1207)	1.3277	2.5443
R-Bar-Squared	0.0084	0.0384
Akaike Info. Criterion	3366.8	3078.7
Schwarz Bayesian Criterion	3282.3	2994.2

Appendix D: An Analysis of the Occurrences of Forecast Error Terms for the VAR(1) regressions.

The largest 40 error terms of the VAR(1) regressions were tabulated to inquire whether such errors may be attributed to a particular day of the week. The table shows that the error terms are distributed across the week. (A lower number of Saturday occurrences may be explained by the fact that the exchange is not usually open for trading on this day). If we hypothesise that more volatile returns are realised on Mondays, this would imply that the VAR(1) regressions would perform badly in forecasting the Monday return, as well as the Tuesday returns (given that the Monday information is then used to forecast the Tuesday return). A higher number of Monday and Tuesday return forecast errors is not clearly evident from the table. This is in line with the empirical results of Choudhry (2000) who found that the change in volatility on Mondays was not significant on the Indian market.

	Log Return Nifty Regression	Log Return Midcap Regression
Monday	11	9
Tuesday	6	6
Wednesday	5	8
Thursday	7	5
Friday	9	11
Saturday	2	1
TOTAL	40	40

Appendix E: Regression Coefficients, Diagnostic Statistics and Error Plots for Nifty Equation in the VAR (daily data).

```
          OLS estimation of a single equation in the Unrestricted VAR
_____________________________________________________________________________
Dependent variable is LRN
1255 observations used for estimation
_____________________________________________________________________________
Regressor              Coefficient       Standard Error         T-Ratio[Prob]
LRN(-1)                   .057459           .046267             1.2419[.215]
LRM(-1)                   .013572           .036329             .37358[.709]
CONST                    .5462E-3          .4499E-3             1.2140[.225]
_____________________________________________________________________________
R-Squared                    .0051762   R-Bar-Squared                 .0035870
S.E. of Regression            .015928   F-stat.    F(  2,1252)    3.2571[.039]
Mean of Dependent Variable   .5888E-3   S.D. of Dependent Variable     .015957
Residual Sum of Squares       .31763    Equation Log-likelihood         3416.0
Akaike Info. Criterion        3413.0    Schwarz Bayesian Criterion      3405.3
DW-statistic                  1.9939    System Log-likelihood           7168.3
_____________________________________________________________________________
                              Diagnostic Tests
_____________________________________________________________________________
    Test Statistics  *        LM Version        *         F Version
_____________________________________________________________________________
                     *                          *
 A:Serial Correlation*CHSQ(   1)=   2.0324[.154]*F(   1,1251)=   2.0292[.155]
                     *                          *
 B:Functional Form   *CHSQ(   1)=   3.0133[.083]*F(   1,1251)=   3.0109[.083]
                     *                          *
 C:Normality         *CHSQ(   2)= 391.8065[.000]*       Not applicable
                     *                          *
 D:Heteroscedasticity*CHSQ(   1)=  56.8301[.000]*F(   1,1253)=  59.4308[.000]
_____________________________________________________________________________
  A:Lagrange multiplier test of residual serial correlation
  B:Ramsey's RESET test using the square of the fitted values
  C:Based on a test of skewness and kurtosis of residuals
  D:Based on the regression of squared residuals on squared fitted values
```

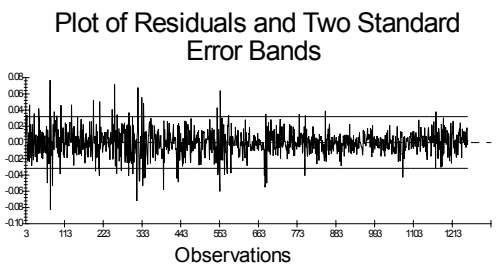

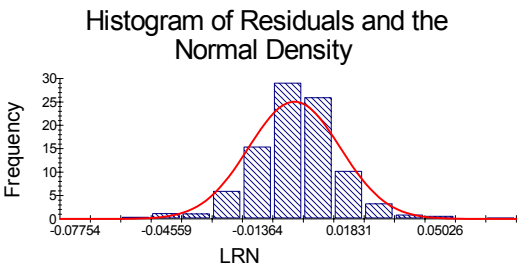

Appendix F: Regression Coefficients, Diagnostic Statistics and Error Plots for Midcap Equation in the VAR (daily data).

```
          OLS estimation of a single equation in the Unrestricted VAR
____________________________________________________________________________
Dependent variable is LRM
1255 observations used for estimation
____________________________________________________________________________
Regressor              Coefficient       Standard Error         T-Ratio[Prob]
LRN(-1)                  -.024358            .058373            -.41728[.677]
LRM(-1)                    .16819            .045834             3.6694[.000]
CONST                    .5341E-3           .5676E-3             .94096[.347]
____________________________________________________________________________
R-Squared                      .023558   R-Bar-Squared                  .021998
S.E. of Regression             .020095   F-stat.    F(  2,1252)   15.1028[.000]
Mean of Dependent Variable   .6260E-3   S.D. of Dependent Variable     .020320
Residual Sum of Squares         .50558   Equation Log-likelihood         3124.4
Akaike Info. Criterion          3121.4   Schwarz Bayesian Criterion      3113.7
DW-statistic                    1.9918   System Log-likelihood           7168.3
____________________________________________________________________________
                              Diagnostic Tests
____________________________________________________________________________
    Test Statistics  *        LM Version        *         F Version
____________________________________________________________________________
                     *                          *
 A:Serial Correlation*CHSQ(   1)=   .83678[.360]*F(   1,1251)=   .83467[.361]
                     *                          *
 B:Functional Form   *CHSQ(   1)=   .38251[.536]*F(   1,1251)=   .38140[.537]
                     *                          *
 C:Normality         *CHSQ(   2)= 218.3047[.000]*       Not applicable
                     *                          *
 D:Heteroscedasticity*CHSQ(   1)= 145.7278[.000]*F(   1,1253)= 164.6097[.000]
____________________________________________________________________________
  A:Lagrange multiplier test of residual serial correlation
  B:Ramsey's RESET test using the square of the fitted values
  C:Based on a test of skewness and kurtosis of residuals
  D:Based on the regression of squared residuals on squared fitted values
```

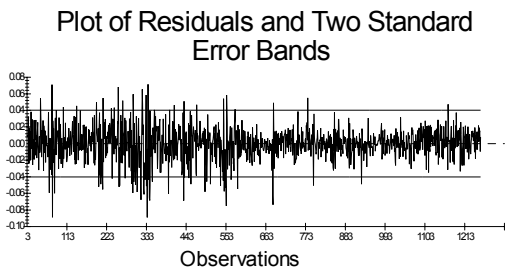

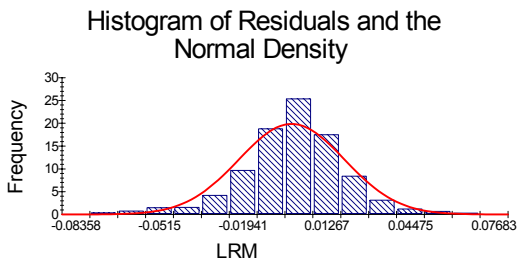

Appendix G: Selecting the Order of the VAR for High Frequency Data

The table below shows AIC, SBC and LR test statistics for a preliminary 24 order VAR which was estimated on the log returns series of Nifty and Midcap. The Akaike Information Criterion selects an order 9 VAR, whilst the Schwarz Bayesian Criterion selects an order 3 VAR. The log-likelihood ratio statistics rejects all orders less than 7, and therefore an order 9 VAR was selected.

Test Statistics and Choice Criteria for Selecting the Order of the VAR Model

Based on 2945 observations. Order of VAR = 24
Variables included in the unrestricted VAR: LRN, LRM
Deterministic and/or exogenous variables: CONST

Order	LL	AIC	SBC	LR test	Adjusted LR test
24	36597.6	36499.6	36206.2	------	------
23	36596.7	36502.7	36221.3	CHSQ(4)= 1.8854[.757]	1.8540[.763]
22	36595.2	36505.2	36235.8	CHSQ(8)= 4.8352[.775]	4.7547[.783]
21	36594.3	36508.3	36250.8	CHSQ(12)= 6.6465[.880]	6.5359[.887]
20	36589.2	36507.2	36261.7	CHSQ(16)= 16.7996[.399]	16.5200[.417]
19	36588.6	36510.6	36277.0	CHSQ(20)= 18.1479[.578]	17.8460[.598]
18	36588.3	36514.3	36292.8	CHSQ(24)= 18.6322[.771]	18.3222[.787]
17	36586.6	36516.6	36307.1	CHSQ(28)= 22.0035[.781]	21.6374[.798]
16	36584.9	36518.9	36321.3	CHSQ(32)= 25.5161[.785]	25.0916[.802]
15	36583.5	36521.5	36335.9	CHSQ(36)= 28.2575[.818]	27.7873[.835]
14	36582.9	36524.9	36351.3	CHSQ(40)= 29.4061[.891]	28.9168[.903]
13	36578.7	36524.7	36363.0	CHSQ(44)= 37.9461[.728]	37.3147[.752]
12	36577.0	36527.0	36377.3	CHSQ(48)= 41.2341[.744]	40.5481[.769]
11	36574.0	36528.0	36390.3	CHSQ(52)= 47.3077[.659]	46.5206[.688]
10	36569.9	36527.9	36402.2	CHSQ(56)= 55.4008[.497]	54.4790[.533]
9	36566.2	36528.2	36414.5	CHSQ(60)= 62.8028[.377]	61.7578[.413]
8	36559.6	36525.6	36423.8	CHSQ(64)= 76.1749[.142]	74.9074[.165]
7	36553.8	36523.8	36434.0	CHSQ(68)= 87.6988[.054]	86.2396[.067]
6	36547.6	36521.6	36443.8	CHSQ(72)= 100.0129[.016]	98.3488[.021]
5	36533.2	36511.2	36445.3	CHSQ(76)= 128.8739[.000]	126.7297[.000]
4	36530.1	36512.1	36458.2	CHSQ(80)= 135.1758[.000]	132.9267[.000]
3	36526.9	36512.9	36471.0	CHSQ(84)= 141.4691[.000]	139.1153[.000]
2	36499.0	36489.0	36459.1	CHSQ(88)= 197.1981[.000]	193.9170[.000]
1	36480.0	36474.0	36456.1	CHSQ(92)= 235.2070[.000]	231.2936[.000]
0	36244.7	36242.7	36236.7	CHSQ(96)= 705.9187[.000]	694.1733[.000]

AIC=Akaike Information Criterion SBC=Schwarz Bayesian Criterion

Appendix H: Regression Coefficients, Diagnostic Statistics and Error Plots for Nifty Equation in the VAR (one-minute data).

OLS estimation of a single equation in the Unrestricted VAR

Dependent variable is LRN
2960 observations used for estimation

Regressor	Coefficient	Standard Error	T-Ratio[Prob]
LRN(-1)	.20302	.019322	10.5071[.000]
LRN(-2)	.012243	.020205	.60591[.545]
LRN(-3)	.010131	.020268	.49985[.617]
LRN(-4)	-.0098900	.020305	-.48706[.626]
LRN(-5)	.012902	.020301	.63553[.525]
LRN(-6)	-.030319	.020266	-1.4960[.135]
LRN(-7)	.048013	.020275	2.3681[.018]
LRN(-8)	-.9586E-4	.020231	-.0047380[.996]
LRN(-9)	.015361	.019935	.77057[.441]
LRM(-1)	.16175	.026685	6.0613[.000]
LRM(-2)	-.036006	.026699	-1.3486[.178]
LRM(-3)	.040859	.026730	1.5286[.126]
LRM(-4)	-.028102	.026648	-1.0546[.292]
LRM(-5)	.033457	.026631	1.2563[.209]
LRM(-6)	-.028304	.026596	-1.0642[.287]
LRM(-7)	-.078469	.026392	-2.9732[.003]
LRM(-8)	-.0016913	.026223	-.064497[.949]
LRM(-9)	-.032466	.025335	-1.2815[.200]
CONST	-.4754E-5	.1055E-4	-.45056[.652]
O (Dummy Variable)	.0027613	.1440E-3	19.1789[.000]

R-Squared	.19020	R-Bar-Squared	.18497
S.E. of Regression	.5710E-3	F-stat. F(19,2940)	36.3434[.000]
Mean of Dependent Variable	.1383E-4	S.D. of Dependent Variable	.6325E-3
Residual Sum of Squares	.9585E-3	Equation Log-likelihood	17915.8
Akaike Info. Criterion	17895.8	Schwarz Bayesian Criterion	17835.8
DW-statistic	2.0194	System Log-likelihood	36954.2

Diagnostic Tests

Test Statistics	LM Version	F Version
A:Serial Correlation	CHSQ(1)= 3.5580[.059]	F(1,2939)= 3.5371[.060]
B:Functional Form	CHSQ(1)= 57.4915[.000]	F(1,2939)= 58.2143[.000]
C:Normality	CHSQ(2)= 104511.6[.000]	Not applicable
D:Heteroscedasticity	CHSQ(1)= 975.0851[.000]	F(1,2958)= 1453.1[.000]

A:Lagrange multiplier test of residual serial correlation
B:Ramsey's RESET test using the square of the fitted values
C:Based on a test of skewness and kurtosis of residuals
D:Based on the regression of squared residuals on squared fitted values

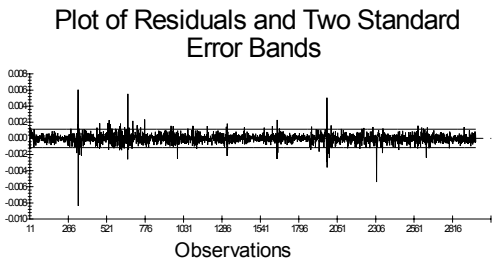
Plot of Residuals and Two Standard Error Bands
Observations

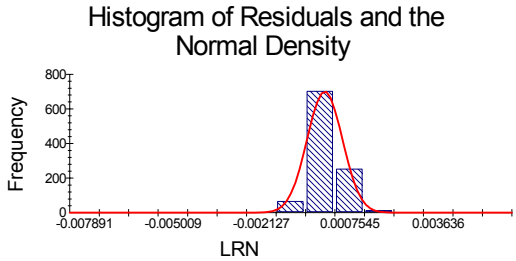
Histogram of Residuals and the Normal Density
Frequency
800
600
400
200
0
-0.007891
-0.005009
-0.002127
0.0007545
0.003636
LRN

Appendix I: Regression Coefficients, Diagnostic Statistics and Error Plots for Midcap Equation in the VAR (one-minute data).

OLS estimation of a single equation in the Unrestricted VAR

Dependent variable is LRM
2960 observations used for estimation

Regressor	Coefficient	Standard Error	T-Ratio[Prob]
LRN(-1)	.24062	.014798	16.2598[.000]
LRN(-2)	.076787	.015475	4.9621[.000]
LRN(-3)	.075153	.015523	4.8415[.000]
LRN(-4)	.018818	.015551	1.2100[.226]
LRN(-5)	.022783	.015548	1.4654[.143]
LRN(-6)	.033814	.015521	2.1785[.029]
LRN(-7)	.011175	.015528	.71969[.472]
LRN(-8)	-.044903	.015495	-2.8980[.004]
LRN(-9)	.039470	.015267	2.5853[.010]
LRM(-1)	.026681	.020437	1.3055[.192]
LRM(-2)	-.080116	.020448	-3.9180[.000]
LRM(-3)	.031154	.020472	1.5218[.128]
LRM(-4)	-.010816	.020409	-.52996[.596]
LRM(-5)	-.010245	.020396	-.50232[.615]
LRM(-6)	.030844	.020369	1.5143[.130]
LRM(-7)	-.041733	.020213	-2.0647[.039]
LRM(-8)	.035813	.020083	1.7832[.075]
LRM(-9)	.010745	.019403	.55379[.580]
CONST	-.5593E-5	.8081E-5	-.69208[.489]
O (Dummy Variable)	.1425E-3	.1103E-3	1.2926[.196]

R-Squared	.16699	R-Bar-Squared	.16161
S.E. of Regression	.4373E-3	F-stat. F(19,2940)	31.0200[.000]
Mean of Dependent Variable	.1591E-5	S.D. of Dependent Variable	.4776E-3
Residual Sum of Squares	.5622E-3	Equation Log-likelihood	18705.3
Akaike Info. Criterion	18685.3	Schwarz Bayesian Criterion	18625.4
DW-statistic	2.0031	System Log-likelihood	36954.2

Diagnostic Tests

Test Statistics	LM Version	F Version
A:Serial Correlation	CHSQ(1)= 2.8695[.090]	F(1,2939)= 2.8519[.091]
B:Functional Form	CHSQ(1)= 29.4679[.000]	F(1,2939)= 29.5531[.000]
C:Normality	CHSQ(2)= 8917.9[.000]	Not applicable
D:Heteroscedasticity	CHSQ(1)= 65.6122[.000]	F(1,2958)= 67.0542[.000]

A:Lagrange multiplier test of residual serial correlation
B:Ramsey's RESET test using the square of the fitted values
C:Based on a test of skewness and kurtosis of residuals
D:Based on the regression of squared residuals on squared fitted values

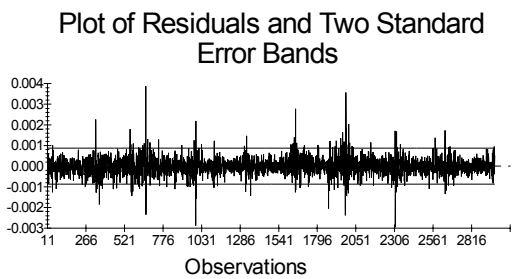
Plot of Residuals and Two Standard Error Bands
0.004
0.003
0.002
0.001
0.000
-0.001
-0.002
-0.003
11
266
521
776
1031
1286
1541
1796
2051
2306
2561
2816
Observations

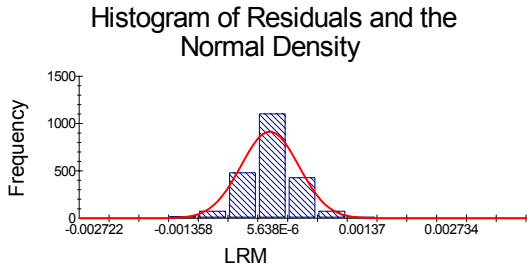
Histogram of Residuals and the Normal Density
Frequency
1500
1000
500
0
-0.002722
-0.001358
5.638E-6
0.00137
0.002734
LRM

Appendix J: Impulse Response Functions for Daily Data.

Orthogonalised and Generalised Impulse Response Functions for a shock in the Nifty Log Return Series are shown in Panel A. Panel B shows Orthogonalised and Generalised Impulse Response Functions for a shock in the Midcap Log Return Series.

Panel A: LRN Shock				
	Orthogonalised		Generalised	
Horizon	LRN	LRM	LRN	LRM
0	0.0159280000	0.0159800000	0.0159280000	0.0159800000
1	0.0011321000	0.0022997000	0.0011321000	0.0022997000
2	0.0000962600	0.0003592000	0.0000962600	0.0003592000
3	0.0000104100	0.0000580700	0.0000104100	0.0000580700
4	0.0000013860	0.0000095130	0.0000013860	0.0000095130
5	0.0000002087	0.0000015660	0.0000002087	0.0000015660
6	0.0000000333	0.0000002583	0.0000000333	0.0000002583
7	0.0000000054	0.0000000426	0.0000000054	0.0000000426
8	0.0000000009	0.0000000070	0.0000000009	0.0000000070
9	0.0000000001	0.0000000012	0.0000000001	0.0000000012
10	0.0000000000	0.0000000002	0.0000000000	0.0000000002
11	0.0000000000	0.0000000000	0.0000000000	0.0000000000
12	0.0000000000	0.0000000000	0.0000000000	0.0000000000
Panel B: LRM Shock				
	Orthogonalised		Generalised	
Horizon	LRN	LRM	LRN	LRM
0	0.0000000000	0.0121840000	0.0126660000	0.0200950000
1	0.0001654000	0.0020492000	0.0010005000	0.0030712000
2	0.0000373100	0.0003406000	0.0000991700	0.0004922000
3	0.0000067670	0.0000563800	0.0000123800	0.0000803600
4	0.0000011540	0.0000093170	0.0000018020	0.0000132100
5	0.0000001928	0.0000015390	0.0000002829	0.0000021790
6	0.0000000320	0.0000002541	0.0000000458	0.0000003595
7	0.0000000053	0.0000000420	0.0000000075	0.0000000594
8	0.0000000009	0.0000000069	0.0000000012	0.0000000098
9	0.0000000001	0.0000000011	0.0000000002	0.0000000016
10	0.0000000000	0.0000000002	0.0000000000	0.0000000003
11	0.0000000000	0.0000000000	0.0000000000	0.0000000000
12	0.0000000000	0.0000000000	0.0000000000	0.0000000000

The Table shows the Impulse Response Function for Shocks in the Nifty Log Return (Panel A) and the Midcap Log Return (Panel B), on the variables included in the VAR which was estimated using daily data. Both the Orthogonalised and Generalised versions of the Functions are shown, for a 12-day horizon.

Appendix K: Impulse Response Functions for High Frequency Data.

Orthogonalised and Generalised Impulse Response Functions for a shock in the Nifty Log Return Series are shown in Panel A. Panel B (overleaf) shows Orthogonalised and Generalised Impulse Response Functions for a shock in the Midcap Log Return Series.

Panel A: LRN Shock

	Orthogonalised		Generalised	
Horizon	LRN	LRM	LRN	LRM
0	0.000571000	0.000196300	0.000571000	0.000196300
1	0.000147700	0.000142600	0.000147700	0.000142600
2	0.000052970	0.000067450	0.000052970	0.000067450
3	0.000032140	0.000063490	0.000032140	0.000063490
4	0.000011170	0.000032250	0.000011170	0.000032250
5	0.000017350	0.000019250	0.000017350	0.000019250
6	-0.000010080	0.000034890	-0.000010080	0.000034890
7	0.000008881	0.000008155	0.000008881	0.000008155
8	-0.000003777	-0.000018410	-0.000003777	-0.000018410
9	-0.000005474	0.000022950	-0.000005474	0.000022950
10	-0.000003544	0.000007513	-0.000003544	0.000007513
11	-0.000005020	-0.000000671	-0.000005020	-0.000000671
12	-0.000002600	0.000001181	-0.000002600	0.000001181
13	-0.000005990	-0.000001526	-0.000005990	-0.000001526
14	-0.000001089	-0.000001025	-0.000001089	-0.000001025
15	-0.000000632	0.000000088	-0.000000632	0.000000088
16	-0.000002573	-0.000001776	-0.000002573	-0.000001776
17	-0.000000853	-0.000000826	-0.000000853	-0.000000826
18	-0.000001350	-0.000000078	-0.000001350	-0.000000078
19	-0.000000700	-0.000000788	-0.000000700	-0.000000788
20	-0.000000418	-0.000000539	-0.000000418	-0.000000539

Table ___ Panel A shows the Impulse Response Function for Shocks in the Nifty Log Return, on the variables included in the VAR which was estimated using data at one-minute intervals. Both the Orthogonalised and Generalised versions of the Functions are shown, for a 20-minute horizon.

… continued overleaf

Appendix K (continued)

Panel B: LRM Shock				
	Orthogonalised		Generalised	
Horizon	LRN	LRM	LRN	LRM
0	0.000000000	0.000390700	0.000256300	0.000437300
1	0.000063200	0.000010430	0.000122800	0.000073350
2	0.000000448	-0.000015820	0.000024180	0.000016150
3	0.000013900	0.000015880	0.000026850	0.000042690
4	-0.000003950	0.000005917	0.000001487	0.000019770
5	0.000011270	-0.000004384	0.000017860	0.000004724
6	-0.000007349	0.000016920	-0.000011090	0.000030780
7	-0.000032240	-0.000013800	-0.000024820	-0.000008669
8	-0.000006968	0.000004473	-0.000007922	-0.000004266
9	-0.000012130	0.000000515	-0.000013300	0.000010760
10	-0.000003888	-0.000005109	-0.000005065	-0.000001192
11	-0.000000872	-0.000003480	-0.000003033	-0.000003411
12	-0.000001323	-0.000000169	-0.000002349	0.000000379
13	-0.000000823	-0.000004060	-0.000003425	-0.000004313
14	-0.000001035	0.000000584	-0.000001414	0.000000061
15	-0.000001083	0.000000095	-0.000001251	0.000000125
16	-0.000000937	-0.000001974	-0.000001992	-0.000002561
17	-0.000000326	-0.000000059	-0.000000674	-0.000000424
18	-0.000000094	-0.000000471	-0.000000690	-0.000000456
19	0.000000029	-0.000000693	-0.000000289	-0.000000973
20	0.000000352	0.000000095	0.000000126	-0.000000157

Table ___ Panel B shows the Impulse Response Function for Shocks in the Midcap Log Return, on the variables included in the VAR which was estimated using data at one-minute intervals. Both the Orthogonalised and Generalised versions of the Functions are shown, for a 20-minute horizon.