

Portfolio Return Autocorrelation and Non-Synchronous Trading in UK Equities

G.S. Morgan¹, P.N. Smith² and S.H. Thomas^{1*}

August 30, 2001

Abstract

Although infrequent trading in equity stocks is more prevalent in the United Kingdom (and other non-United States countries), in contrast to much of the existing literature we find that it is proportionally more important in explaining the degree of serial correlation in stock returns in the US than in the UK. We show that infrequent trading cannot explain more than a small proportion of the serial correlation observed in monthly UK stock returns and hence, other explanations for return predictability must be sought. Many studies have shown that stock market returns in the UK and other international markets are substantially and significantly serially correlated. The success of an investment strategy that is based on the apparent predictability of returns depends on whether the serial correlation is truly random and period specific or due to time varying risk premia or to

^{*1}University of Southampton,²Corresponding author; University of York, Department of Economics and Related Studies, Heslington, York YO1 5DD, United Kingdom, (pns2@york.ac.uk), who is grateful to the Economics Program, RSSS, Australian National University for providing a very congenial environment in which to complete this paper. The authors are also grateful to Andrew Clare, Keith Cuthbertson, Karen Mumford, Don Poskitt, Chris Skeels, Jim Steeley and Mike Wickens for helpful detailed and constructive comments which have greatly improved the paper and to the ESRC and INQUIRE for financial support.

market microstructure effects. A frequently noted explanation for this serial correlation is market thinness or, more precisely, the infrequency with which a substantial number of UK stocks are traded. Non-synchronous trading results in a measurement error in the observed data for returns on individual stocks, portfolios and market indices. This measurement error generates serial correlation in the observed returns. Here, we assess the extent to which the observed serial correlation in returns can be explained by equity non-trading behaviour. This will reveal whether there is any residual serial correlation left to be explained by alternative sources. We find that, whilst a proportion of the serial correlation in the returns of portfolios of low value stocks can be explained by non-trading, much of it still remains unexplained.

1. Introduction

There is now extensive evidence that security returns are predictable, see (Keim and Stambaugh (1986) and Campbell and Hamao (1992)). An important question from the point of view of an investor seeking to exploit such predictability is whether it is due to time varying risk premia or market microstructure effects, including infrequent trading. It is widely considered that non-US equity markets are characterized by thin trading which induces serial correlation in returns and distorts measures of risk (Dimson (1979), Clare, Morgan, and Thomas (2002), Lange (1999)). Here we challenge this conventional view by suggesting that infrequent trading is proportionally *more important* in the US security market than in a major alternative market such as that in the UK. This is important in assessing the sources of predictability in these markets and in testing asset pricing theories.

There are two forms of infrequent trading: *non-synchronous* trading which implies that trading takes place in every consecutive time interval but not necessarily at the close of the interval of measurement of the returns data (see Fisher (1966), Dimson (1979), Lo and MacKinlay (1990), Stoll and Whaley (1990) and Muthuswamy (1990)), while, *non-trading* occurs when securities do not trade in every consecutive time interval (see Cohen, Maier, Schwartz, and Whitcomb

(1978) and Stoll and Whaley (1990)). As Miller, Muthuswamy, and Whaley (1994) emphasize, the key to this distinction lies in the length of the interval of measurement of the returns data. If returns are calculated on a monthly basis then nearly all NYSE (New York Stock Exchange), and to a slightly lesser extent LSE (London Stock Exchange), stocks will have traded at least once in the month. However, by no means all stocks will have traded precisely at the close of trading on the last day of the month when the data is collected. If we foreshorten the interval of measurement to ten minutes, then many stocks will not have traded in the period of measurement. Hence, non-synchronous trading becomes non-trading as the interval of measurement shrinks.

In this paper we focus on non-synchronous trading within the one month period of measurement for UK stock returns available on the LSPD (London Share Price Database). This data source is the most important stock return database for risk measurement and the testing of asset pricing theories in the UK, (eg Dimson (1979)). We have total return information on all UK stocks from 1975 - 1995 together with a non-trading marker which indicates the number of days from the end of the month when the stock last traded. This aspect of the LSPD has not been exploited, to our knowledge, since Dimson (1979) and there have been substantial changes in it's behaviour since then.¹ Data for returns calculated on an end of month basis will reflect the price of the last recorded trade which, for an infrequently traded stock could be many days old as we show below. Fama (1965) and Fisher (1966) show that this non-synchronous trading leads to measurement error in the recorded returns on individual stocks, portfolios and indices. In particular it induces positive serial correlation in portfolio and index returns which may be mistaken for predictability.

In Section 2 we describe the relation between trading frequency and firm size for the UK. Although many empirical studies acknowledge the importance of thin trading, very few give detailed information on its nature and extent. Here we compare our data with that for the US described by Foerster and Keim (1993).

¹Although see Clare, Morgan, and Thomas (2002) for a recent analysis. An alternative approach emphasising liquidity is in Datar, Naik, and Radcliffe (1998)

This leads naturally to a discussion of the time-series behaviour of portfolio returns in Section 3, where we report finding substantial, significant first-order serial correlation in the more infrequently traded, medium to low market capitalisation stocks. In Section 4 we discuss two models of infrequent trading which will allow us to calibrate an estimate of the degree of autocorrelation induced by such trading patterns and hence the extent to which autocorrelation may be present for other reasons. The first model is due to Lo and MacKinlay (1990) and is based on the simple assumption that the non-trading indicator for a stock, for a fixed time interval, follows a simple Markov process and is time independent. This model turns out not to be a good description of the pattern of non-trading in the UK as the implied degree of non-trading at longer durations is substantially below that observed in the data. Consequently, we consider a second model which introduces time dependence in the non-trading indicator using a two-state Markov model (see Campbell, Lo, and MacKinlay (1997)). Although this provides a much improved fit of the empirical distribution, we still underestimate the persistence in infrequent trading in the data. In Section 5 we examine the implications of the two Markov models and the empirical distribution of infrequent trading for the calculation of serial correlation in portfolio returns. We calculate how much of the observed serial correlation in returns is due to infrequent trading. We also assess the importance of assuming within-decile homogeneity in non-synchronous trading behaviour. Unlike Boudoukh, Richardson, and Whitelaw (1994), we find that allowing for heterogeneity generates little change in our results. Finally, in Section 6 we discuss our results and a trading rule based upon our findings before presenting our conclusions.

2. Infrequent Trading and Firm Size

2.1. The London Share Price Database

Our sample consists of all stocks listed on the LSPD for any period of time between 1975 and 1995. In particular, our sample includes stocks traded on the Alternative

(one-time Unlisted) Securities Market as well as the London Stock Exchange.² It thus has a complete representation of stocks of all sizes, including the smallest. The data has a monthly frequency.³ In order to be admitted to the sample for a particular month, a stock is required to have a valid return for that month⁴, and a market capitalisation value for January of the current year⁵. Stocks enter and leave the sample each month according to whether or not these admission criteria are satisfied, so the sample is free of survival bias problems. The number of stocks in the sample varies from one month to another, but is consistently greater than 2,000.

In this paper we exploit an aspect of the LSPD which has attracted little attention in the literature, namely the non-trading marker. Each month-end return has a corresponding non-trading marker which gives the number of days before the end of the month that the last trade occurred, making it possible to examine patterns in the extent of infrequent trading over time. A value of zero indicates that a stock was traded on the last trading day of the month, while values of 1-31 mean that the last trade occurred between one and 31 days respectively. The number 32 indicates that a stock was not traded during the current month⁶, and so the corresponding return will generally be zero, unless a dividend payment occurred during the month⁷. There is, therefore, an element of non-trading in our data. All stocks which pass the selection criteria are assigned to

²For a detailed description of the structure of the London Stock Exchange see Buckle and Thomson (1998).

³Some of the work done in this area has been carried out using weekly or daily data. Boudoukh, Richardson, and Whitelaw (1994) focuses on weekly data. The LSPD is the most comprehensive data on stock returns in the UK and the only publicly available one, to our knowledge, with information on trading frequency.

⁴The LSPD identifies a missing monthly (log) return by recording a value of -10. Any other value for the return was taken as evidence that the stock was alive during that month.

⁵The LSPD only provides annual market capitalisation data.

⁶In Table 1A, discussed later, it can be seen that the proportion of stocks traded outside of the month is negligible.

⁷Occasionally the non-trading indicators have values between 33 and 63, which relate to special cases in which the previous month's price was unavailable. Further details are given in LBS (1996).

one of 10 portfolios, according to their market capitalisation, and portfolio returns were calculated for each decile. Returns are equally weighted within portfolios. Monthly rebalancing resulted in a maximum time series of 251 observations for each of the size-ranked portfolios. Over the period 1975-1995, Clare, Morgan, and Thomas (2002) find that almost 44% of all stocks listed on the LSPD fail to trade on the last day of a given month, a figure which is significantly higher than for stocks in the US (Foerster and Keim (1993)). However, closer examination of the data reveals that stocks are much more likely to be recorded as not trading on the last day of the month in the period prior to April 1981 than is the case after this date. This is due to changes in the recording requirements relating to trades which were introduced by the London Stock Exchange (LSE) in March 1981. Prior to this date, the requirements relating to the marking of trades were less stringent, with the result that a significant proportion of transactions were not recorded on the LSPD,(see LSE (1981)). As Clare, Morgan, and Thomas (2002) show, this discontinuity has important implications for risk measurement. On average, over two thirds of stocks failed to trade on the last day of the month in the period before April 1981, but this figure falls to around one third of the sample in the second sub-period. As the method employed in the latter period is considered to be much more accurate (see LSE (1981) and Clare, Morgan, and Thomas (2002)), we restrict our analysis to the period April 1981 - December 1995 or 176 observations.⁸

2.2. Non-Synchronous Trading in Equities and Firm Size

We now consider the nature of non-synchronous trading in our data. In Table 1 we present some summary statistics for non-synchronous trading in each of the value-ranked portfolios. The results demonstrate a clear relationship between firm size or market capitalisation and the likelihood of it's stock being traded. Part A of the table gives the proportions of the stocks in each portfolio that have traded on the day(s) concerned. For example, in the fourth highest value portfolio 73.6%

⁸Note that Foerster and Keim (1993) noticed a similar data inconsistency on the CRSP tapes which could have led researchers to come to incorrect inferences about returns.

of the stocks traded on the last day, a further 12.0% traded last on the day before, and so forth. In the final trading week of the month, 95.0% of stocks have traded. This increases to 98% when we consider the last two weeks of the month. The corresponding trading figures for the smaller value portfolios are both substantially smaller for the last day and more spread out over time. In the smallest value portfolio only 36.9% traded on the last day, 13.3% on the penultimate day and 67.4% in the last week. This extends to 78% in the last two weeks and to only 90.3% over the month as a whole. The relationship between market capitalisation and non-trading is monotonic over the ten portfolios examined here. This is consistent with the US results presented by Foerster and Keim (1993) and provides support for our choice of firm value as the basis for portfolio construction. An alternative would be that of portfolios ranked by dividend yield. However, further analysis of our data by Clare, Morgan, and Thomas (2002) demonstrates that there is no systematic relationship between dividend yield and the incidence of non-trading. Moreover, Clare, Smith, and Thomas (1997) exploit this fact in testing the conditional CAPM. There is evidence that zero-dividend stocks are more prone to non-trading. However, this is a group of low value stocks which confirms that size is highly correlated with trading frequency.

3. Serial Correlation in Portfolio Returns

The autocorrelation patterns in the portfolio return data are summarised in Table 2. The portfolios are arranged by decreasing market capitalisation. The values of the autocorrelation function (ACF) and partial autocorrelation function (PACF) show significant first order serial correlation for all but the three largest size deciles. The order of the serial correlation in the returns for portfolios of all sizes clearly does not exceed two. Comparison of the ACF and PACF for each portfolio return shows significant first order coefficients and significant higher order effects in the ACF but not in the PACF suggesting an AR(1) model structure. We consider the ARMA(1,1) model to be an adequate model of the serial correlation having found no significant residual serial correlation. In the last column we

present likelihood ratio tests of the restriction of an ARMA(1,1) model to either AR(1) or MA(1) models. These statistics are distributed $\chi^2(1)$ under the null hypothesis and show that the MA(1) is in general rejected whereas the AR(1) is not. We conclude that an AR(1) model is most appropriate, although there is some evidence that the return on the smallest value portfolio may be best modelled by an ARMA(1,1).

The size and sign of the first order autocorrelation coefficients estimated are also of interest. From Table 2 they can be seen to range from 0.45 for the smallest value portfolio to 0.004 for the largest. Comparative estimates for the US are much lower; those provided by Campbell, Lo, and MacKinlay (1997) are 0.21 and 0.01, respectively, for a comparable period of monthly data on returns of NYSE and AMEX stocks taken from the CRSP files for the period 1962 - 1990. Returns appear always to be positively autocorrelated. The limited negative autocorrelation in the highest value portfolio returns is not significant, even at the 90% level. If the period 1975 -1981 is examined, significant positive first-order and negative second order serial correlation are found for the two largest-value deciles. However, we regard this as a reflection of the inadequacy of the method of data collection rather than as a reflection of the reliability of the estimates.⁹ The estimates of the first elements of the ACF provide us with a benchmark against which to assess the importance of non-synchronous trading.

A second issue concerns cross-serial correlation between portfolio returns. Table 3 presents a summary of the first-order cross-serial correlation coefficients for the largest, smallest and mid-value portfolios. These clearly show that higher value portfolio returns appear to lead the lowest value portfolio returns. None of the lower value portfolio returns appear to predict the highest value portfolio return. Lo and MacKinlay (1990) refer to this lead-lag relationship as a stylised fact which supports, among other things, the potential profitability of contrarian trading strategies. They go on to show that it could be explained by non-synchronous

⁹A complete set of estimates for the full period and two sub-periods is presented in Clare, Morgan, and Thomas (2002). These results show that the estimates for the lower value portfolio returns are affected much less by the sample split.

trading. However, others (eg Boudoukh, Richardson, and Whitelaw (1994)) have pointed out that this may only be an artefact of two observations: small value stocks are significantly autocorrelated; and, high and low value portfolio returns are highly contemporaneously correlated. Taken together, these two effects would generate the significant lead-lag relationship shown in Table 3. An additional piece of evidence supporting this view for our UK data is provided in the second part of Table 3. Using a one-lag VAR in all ten of the portfolio returns as the basis of the analysis, we present Granger causality tests for each of the lagged returns. None of the statistics presented are significant at any reasonable level of significance suggesting that none of the portfolio returns Granger-causes any of the others. In other words, once we condition on own- and other-value portfolio returns, there is no significant additional predictability of the small (or higher) value portfolio returns from the highest value return.

Given that the lead-lag evidence is unconvincing, we examine the observed data and models of non-synchronous trading to provide more compelling evidence on the relationship between infrequency of trading and portfolio return autocorrelation.

4. How Important is Non-Synchronous Trading?

Thus far we have shown that there is substantial, significant first-order serial correlation in medium to low value stock returns and that there is also a high degree of regular non-synchronous trading behaviour in these stocks. In this section we examine two models of non-synchronous trading and assess their empirical relevance. In the following section we can then estimate the degree of autocorrelation induced by non-synchronous trading behaviour.

The best known model is due to Lo and MacKinlay (1990) and is based on the most straightforward assumption to make concerning whether an individual stock trades over a fixed period of time which is that the non-trading indicator variable for any stock i , δ_{it} , follows a Markov process and is independent over time. More formally, if the stock i does not trade in period t , then $\delta_{it} = 1$ with

probability $(1-\pi_i)$ and if it does trade then $\delta_{it} = 0$ with probability π_i . It is further assumed that π_i is constant over time and that the stocks in each portfolio are homogeneous with common non-trading probability π . The assumption of time-independence makes much more feasible the closed-form solution of single equity and portfolio return autocorrelations. These solutions are presented in Lo and MacKinlay (1990) and Boudoukh, Richardson, and Whitelaw (1994). However, as we will show, this time-independence does not seem to describe the UK data very well and we relax this assumption in our second model.

For each portfolio, if we take the proportion of stocks that trade on the last day of the month (day 0) as our measure of the time-invariant, and independent, trading probability (ie, π), we can calculate the implied distribution of trading on day k over the average month of the sample period as $(1 - \pi)^{k-1}\pi$. Part B of Table 1 gives these predicted proportions for the size decile portfolios. It is clear that the implied degree of non-trading at longer durations is substantially underestimated using this assumption. For all portfolios, the longer durations are much more important in the data than time independence suggests. The point is made graphically in Figure 1 which shows the predicted trading proportions for the smallest value decile. The integral under each of the curves shown is equal to one, so the fact that the empirical curve is substantially below that generated by the time-invariant model demonstrates that it must lie above the time-invariant model curve for longer durations. In short, the empirical distribution of trades is much more persistent, a feature we now build into the theoretical model.

An alternative assumption concerning the non-trading indicator variable is that it displays some formal time dependence. A two-state Markov model can deliver such a result.¹⁰ The model for δ_{it} now has transition probabilities as follows:

¹⁰This structure for the formal time-dependence of non-trading is foreshadowed in comments in Section 5 of Lo and MacKinlay (1990) and in Campbell, Lo, and MacKinlay (1997). They do not present any calculated trading proportions or predicted serial correlations.

			δ_{it}	
			trade	no trade
			0	1
δ_{it-1}	trade	0	π_i	$1 - \pi_i$
	no trade	1	$1 - \phi_i$	ϕ_i

so $P(\delta_{it} = 1 | \delta_{it-1} = 1) = \phi_i$ etc.

Therefore, the probability of the stock being traded in period t depends, in part, on whether it was traded in period $t-1$. This specification has the attractive feature that there is persistence in the probability of non-trading. It also allows for the presence of clustering of information flows across days. To calculate the implied distribution of trades over time we need the steady-state unconditional probabilities of trading and non-trading in period t . Define Ψ_{it} and Θ_{it} , respectively, as the unconditional probabilities of $\delta_{it} = 0$ and $\delta_{it} = 1$ and P and Q as their steady-state values. The transition probabilities above imply that the unconditional probabilities evolve through time as:

$$\begin{aligned}\Psi_{it} &= \pi_i \Psi_{it-1} + (1 - \phi_i) \Theta_{it-1} \\ \Theta_{it} &= (1 - \pi_i) \Psi_{it-1} + \phi_i \Theta_{it-1}\end{aligned}\tag{4.1}$$

The steady-state version of this system can be derived using eigenvalue methods or otherwise (see Hamilton (1995)) as:

$$\begin{aligned}P_i &= \frac{1 - \phi_i}{2 - (\pi_i + \phi_i)} \\ Q_i &= \frac{1 - \pi_i}{2 - (\pi_i + \phi_i)}\end{aligned}\tag{4.2}$$

The proportion of trades which take place in period k is therefore:

$$Z_{ik} = (\pi_i P_i + (1 - \phi_i) Q_i \pi_i) (1 - \pi_i) \phi_i^{k-1}\tag{4.3}$$

We concentrate on the smallest value decile portfolio as this allows comparison with US studies such as Boudoukh, Richardson, and Whitelaw (1994), consists of the most infrequently traded stocks and has the largest serial correlation coefficient. The series Z_k for this portfolio in the UK data is calibrated as follows. Taking the observed empirical trading proportions for periods $k=0$ and 1 , values for π_i and ϕ_i are calculated. These values are then used to compute the remaining trading proportions. For the smallest value portfolio $Z_0 = 0.37$ and $Z_1 = 0.13$. This implies from simultaneous solution of (4.3) for $k=0,1$ that $\pi = 0.650$ and $\phi = 0.794$. From (4.1) it can be seen that this implies a degree of persistence in both the probabilities of trading and non-trading, as we would expect.

The Markov time-dependent trading proportion for the smallest-value-stock decile is shown in Figure 1 along with the other two alternative series. It is clear that whilst the persistence in the Markov time dependent series is also not as great as in the empirical series, it is much closer and any comparison based on the correlation with the empirical proportions would clearly favour the Markov time-dependent model.¹¹

The last rows of Tables 1A and 1B give some comparative results for small US stocks from Lo and MacKinlay (1990) and Boudoukh, Richardson, and Whitelaw (1994). The last row of Table 1A is constructed assuming time independence as above, the final row of Table 1B gives the empirical distribution of trading frequency presented in Foerster and Keim (1993). Comparison of these rows with those above shows two things. First, the degree of non-trading in the US is much smaller. The behaviour of the smallest stocks resembles that of the fourth highest decile in the UK. Second, non-trading is much more persistent in the UK than in the US. Behaviour of the fourth highest UK decile over two weeks resembles that of the smallest group of US stocks over a week. This suggests that the time-dependent model is of even more value in modelling the UK than the US trading

¹¹Of course, if we increase the number of states in the Markov model we would generate an increasing degree of serial correlation in the trading proportion which could come ever closer to the empirical distribution. Our result shows that a large degree of improvement in matching the empirical proportions is achieved by adding only one extra state.

distribution.

5. Implications for Serial Correlation in Portfolio Returns

The results above show that in modelling the relationship between non-trading and portfolio return autocorrelation, we can expect the assumption of temporal independence to under-estimate the degree of persistence in non-trading. We can also anticipate that it would also lead to an under-estimate of the predicted extent of autocorrelation in returns. In order to estimate the degree of serial correlation induced by the non-synchronous trading, we analyse a model built on the foundations of the work of Scholes and Williams (1977). This model requires that all stocks trade within a fixed time interval but any distribution of non-trading behaviour can be accommodated and the period can be extended appropriately. We begin by requiring homogeneity of stocks within each portfolio in terms of their non-trading probability, an assumption we relax in the next section. Let s_{it} be the time between the last trade of stock i and the end of the trading period, expressed as a proportion of the length of the trading period. In our data this is a month and we allow the representative month to be made up of 31 days.¹² If stock i is not traded for one day it will have an s_{it} of $1/31$. The s_{it} are assumed to be independently distributed over time. Atchison, Butler, and Simonds (1987) show that the first order autocorrelation coefficient of the return of a well-diversified portfolio of n stocks is:

$$\text{corr}(R_{pt}, R_{pt-1}) \approx \frac{\sum_{i=1}^n \sum_{j=1, \neq i}^n \text{cov}(R_{it}, R_{jt-1})}{\sum_{i=1}^n \sum_{j=1, \neq i}^n \text{cov}(R_{it}, R_{jt})} \quad (5.1)$$

If we further assume that the s_{it} and s_{jt} are independent or, at least that the returns are variable enough, then Boudoukh, Richardson, and Whitelaw (1994) show that for an equally-weighted portfolio of n stocks,

$$\text{corr}(R_{pt}, R_{pt-1}) \approx \frac{\sum_{i=1}^n \sum_{j=1, \neq i}^n E[\max(s_{it} - s_{jt}, 0)]}{\sum_{i=1}^n \sum_{j=1, \neq i}^n 1 - E[\max(s_{it}, s_{jt})] + E[\min(s_{it}, s_{jt})]} \quad (5.2)$$

¹²The choice of 31 or 30 days makes very little difference to our results.

Assuming, further, that the s_{it} are homogeneous, and therefore all have the same distribution, allows (5.2) to be written in terms of the expectation for any pair of stocks, ie:

$$\text{corr}(R_{pt}, R_{pt-1}) \approx \frac{E[\max(s_{it} - s_{jt}, 0)]}{1 - E[\max(s_{it}, s_{jt})] + E[\min(s_{it}, s_{jt})]} \quad (5.3)$$

This solution can be simulated for any set of s_{it} . We consider two types of distribution of s_{it} , that from a time-independent distribution and that from the empirical, time dependent distribution. We use the distributions for the market-value based decile portfolios from Parts A and B of Table 1, respectively. After some experimentation we find we can get a non-degenerate solution from a 10,000 replication simulation. The results are presented in Tables 4 and 5. In Table 4 we concentrate in detail on the results for the smallest value decile portfolio under a number of alternative models, since we have comparable figures for the US from Boudoukh, Richardson, and Whitelaw (1994) and this decile has the largest degree of non-synchronous trading and the largest serial correlation coefficient. In Table 5 we present results for all deciles under a single model. From Table 4 it can be seen that the time-independent theoretical distribution gives a calculated first-order autocorrelation coefficient of 0.036, whereas the Markov time-dependent distribution gives a value of 0.082 and the empirically based distribution gives a value of 0.116. It is clear that the assumption of time-independence of non-trading probabilities makes a substantial difference to the implied autocorrelation in returns. The time dependent Markov model gives a much closer estimate to that given by the empirical distribution emphasising the importance of time dependence.

These values are, however, much smaller than the estimated autocorrelation coefficient for the return on the smallest value portfolio reported in Table 2 which is 0.45. Put differently, our results show that, *non-trading explains, at most, about one quarter of the autocorrelation in the returns to the smallest value portfolio in the UK*. The various percentages for the other models are shown in Table 4.

Comparable results for monthly stock returns in the US can be constructed from the figures and calculations in Lo and MacKinlay (1990), Campbell, Lo, and

MacKinlay (1997) and Boudoukh, Richardson, and Whitelaw (1994).¹³ These are also laid out in Table 4. They show, for the time independent Markov case, that non-synchronous trading generates an AR(1) coefficient of 0.013 or 6.2% of the estimated first-order autocorrelation coefficient (0.21) in stock returns for the smallest value decile portfolio of the NYSE-AMEX stocks reported on the CRSP files for the period 1962 - 1990. The empirical distribution of non-trading delivers a value of the first-order autocorrelation coefficient of 0.061 which is about 29% of the estimated first-order autocorrelation coefficient. The evidence taken together says that *non-trading explains more of the degree of serial correlation in returns in the US than it does in the UK*, whilst both are smaller in the US than in the UK.

We can also consider the predicted degree of autocorrelation in higher market value portfolio returns. These are shown in Table 5 for the empirical distribution of non-trading. Comparison of the three columns in Table 5 shows that non-trading and the relationship between non-trading and the observed degree of first-order autocorrelation is monotonically increasing (although not linearly) as decile value declines. To highlight one entry, the fourth highest value decile portfolio resembles that of the lowest value in the US data most closely in terms of its non-trading behaviour. In terms of autocorrelation in the returns data we have an estimate of 0.253 from Table 2, above. The first-order autocorrelation predicted by simulation of equation (5.3) with empirical non-trading probabilities, is 0.0267 or some 11% of the observed value. This is a substantially smaller proportion than the 29% Boudoukh, Richardson, and Whitelaw (1994) report for the US.¹⁴

¹³An alternative approach based on simulation of the impact of non-synchronous trading on individual firms stock returns is presented by Kadlec and Patterson (1999). Their estimates of the autocorrelation generated by non-synchronous trading are a little lower than that of Boudoukh, Richardson, and Whitelaw (1994).

¹⁴The high figure of 31.5% for the highest value decile reflects the imprecision of the estimation of two tiny coefficients and is spuriously large.

5.1. Heterogeneity within Portfolios

Thus far we have treated all stocks within a portfolio as homogenous with respect to non-synchronous trading behaviour, making possible the application of equation (5.3) for the various models of non-trading. The evidence from Boudoukh, Richardson, and Whitelaw (1994) is that within-portfolio heterogeneity may also lead to a higher implied degree of autocorrelation. A straightforward way of assessing the importance of heterogeneity is to apply the approach of Lo and MacKinlay (1990) for the time-independent Markov model. We examine the component percentiles of the lowest value decile portfolio and then compute the implied autocorrelation coefficient of the aggregated decile return. Lo and MacKinlay consider an underlying model of the unobservable return on an individual or group of stocks which is determined by one factor of the form

$$R_{it}^* = \mu_i + \beta_i \Lambda_t + \epsilon_{it} \quad (5.4)$$

where Λ_t is the single factor (such as the market return) and which has a constant variance σ_λ^2 and the errors ϵ_{it} are iid. As before, the stock trades in the current period with a probability π . Here we assume that this probability is time independent. As before, the observed return of a stock trading in the current period depends upon information accumulated over the prior periods over which the stock did not trade. So if we consider the aggregation of observed stock returns over q periods indexed τ , we get $R_{i\tau}(q) = \sum_{i=(\tau-1)q+1}^{q\tau} R_{it}$. The intertemporal n th order covariance, given this set of assumptions was shown by Boudoukh, Richardson, and Whitelaw (1994) to be:

$$\begin{aligned} cov[R_{i\tau}(q), R_{j\tau+n}(q)] &= \left[q - \frac{(1 - \pi_i)(1 - (1 - \pi_i)^q)\pi_j^2 + (1 - \pi_j)(1 - (1 - \pi_j)^q)\pi_i^2}{1 - (1 - \pi_i)(1 - \pi_j)\pi_i\pi_j} \right] \beta_i \beta_j \sigma_\lambda^2 \\ &\text{for } n = 0 \\ &= \frac{\pi_i \pi_j}{1 - (1 - \pi_i)(1 - \pi_j)} \left(\frac{1 - (1 - \pi_j)^q}{\pi_j} \right)^2 (1 - \pi_j)^{(n-1)q+1} \beta_i \beta_j \sigma_\lambda^2 \\ &\text{for } n > 0 \end{aligned} \quad (5.5)$$

This expression can be combined with equation (5.1) to calculate the return autocorrelation coefficient.

We examine the impact of within-portfolio heterogeneity on the autocorrelation of the returns for the smallest value decile. The non-trading probabilities $(1 - \pi_i)$ for the ten percentiles that make up the smallest value decile are presented in Table 6. The percentile portfolios are constructed in an identical way to the deciles discussed earlier. Using these figures in equations (5.5) and (5.1) generates the autocorrelations shown in the lower panel of Tables 6. First we give the AR(1) coefficient assuming that all of the ten percentile portfolios have the same non-trading probability equal to the average for the decile. This figure of 0.0389 is equivalent to that given in Table 4 using the simulation method above. The small difference in the figures is due to sampling error in the simulations. Two figures are given based on the heterogeneity in non-trading probability across the percentiles shown in the top part of Table 6. The first assumes that the β on the single factor in equation (5.4) is equal to one for all percentiles. The second employs the estimated β from the market model as an example of a single factor model.¹⁵ The declining β shown, whilst looking counter-intuitive, is consistent with the results for the UK market in various studies (eg Clare, Priestley, and Thomas (1998)). The autocorrelation coefficients generated allowing for heterogeneity are 0.0403 and 0.0394, respectively. These are very similar to that assuming within-decile homogeneity and suggest that within-decile heterogeneity is not an important issue in assessing the importance of non-synchronous trading on serial correlation in the UK. Boudoukh, Richardson, and Whitelaw (1994) finds a more substantial impact from heterogeneity for the US.¹⁶ The main reason for this is that there is no evidence in our data for the UK of the much greater within-decile heterogeneity in the non-trading probability found for the US by Foerster and Keim (1993) which is used by Boudoukh, Richardson, and Whitelaw (1994) to generate the increased autocorrelation coefficient.

¹⁵The β 's are estimated using the adjustment for non-synchronous trading proposed by Scholes and Williams (1977), see Clare, Morgan, and Thomas (2002) for discussion of the impact of non-synchronous trading on the estimation of β .

¹⁶Kadlec and Patterson (1999) discuss the assumptions of Boudoukh, Richardson, and Whitelaw (1994) and their impact on the calculation of implied autocorrelation.

6. Discussion and Conclusions

In this paper we examine the extent of serial correlation in the returns to portfolios of UK equities and the degree to which this can be explained by the extent of observed non-synchronous trading in those stocks. Before reflecting on the importance of infrequent trading, we discuss whether the autocorrelation itself is a potential source of excess profits and evidence of inefficiency in this market.

Return predictability suggests the construction of trading rules for these portfolios which might be profitable. To provide some measure of the likely profitability of such rules we examine two alternative investment strategies for the smallest value portfolio. First, we consider a risk neutral, return maximisation rule which is a simple strategy of buying (selling) the portfolio in month t if it produced a positive (negative) return in $t-1$, a momentum strategy. The alternative investment is assumed to be one which receives the risk-free interest rate. To calculate profitability, the strategy can be compared with a buy-and-hold rule of buying the portfolio in the first month of our data period and holding it until the end, 176 months later. If we compare the compound nominal return for the two alternatives, we find that both are profitable, producing 431.15 for a 100 unit investment in the trading momentum strategy and 273.05 for the buy-and-hold strategy. The trading strategy is 1.2% more profitable per one-way trade, which may be profitable for risk-neutral institutional investors but most certainly not for retail investors. This result supports market efficiency for the market as a whole as all of the higher value portfolios are less substantially autocorrelated.

Whilst apparently not the basis of excess profits, autocorrelation in returns presents an anomaly to be explained. In this paper we have considered one explanation which has had wide coverage but, thus far, had received very little direct analysis or calibration. We have established that non trading is an important source of autocorrelation and that the modelling of the incidence of non-trading had to be extended from the most widely used model. The two-state, time dependent Markov model analysed in this paper generates a pattern of non trading behaviour which is much closer to the observed empirical data than does the

classic time independent model (see Clare, Morgan, and Thomas (2002)).

It is the case, though, that non trading is not the whole of the story. It explains about one quarter of the first order autocorrelation in monthly returns for small stocks. Alternative explanations for the observed autocorrelation include time-varying risk premia,¹⁷ the presence of agents employing feedback trading strategies¹⁸ and the presence of adjustment costs for traders.¹⁹ A complete explanation should probably include components of each of these alternatives. However, the parameterisation proposed in this paper offers the basis for a more precise calibration of CAPM β s and APT sensitivities for markets where stocks are more infrequently traded than in the US. Further, the success, or otherwise, of momentum or contrarian investment strategies may well be influenced by the interaction between the measurement period of returns and the extent of infrequent trading in the measurement interval.

¹⁷See Clare, Smith, and Thomas (1997) for an example with this data set.

¹⁸See DeLong, Schleifer, Summers, and Waldman (1990)

¹⁹See Mech (1993).

References

- ATCHISON, M., K. BUTLER, AND R. SIMONDS (1987): “Nonsynchronous Security Trading and Market Index Autocorrelation,” *Journal of Finance*, 42, 533–553.
- BOUDOUKH, J., M. RICHARDSON, AND R. WHITELAW (1994): “A Tale of Three Schools: Insights on Autocorrelations in Short-Horizon Stock Returns,” *Review of Financial Studies*, 7, 539–573.
- BUCKLE, M., AND J. THOMSON (1998): *The U.K. Financial System*. Manchester University Press, 3rd edn.
- CAMPBELL, J. Y., AND Y. HAMAOKA (1992): “Predictable Stock Returns in the United States and Japan: A Study of Long-term Capital Market Integration,” *Journal of Finance*, pp. 43 – 69.
- CAMPBELL, J. Y., A. W. LO, AND A. C. MACKINLAY (1997): *The Econometrics of Financial Markets*. Princeton University Press.
- CLARE, A., G. MORGAN, AND S. THOMAS (2002): “Direct Evidence of Non-Synchronous Trading on the London Stock Exchange,” *Journal of Business Finance and Accounting*.
- CLARE, A., R. PRIESTLEY, AND S. THOMAS (1998): “Reports of Beta’s Death are Premature: Evidence from the UK,” *Journal of Banking and Finance*, pp. 1207 – 1229.
- CLARE, A., P. SMITH, AND S. THOMAS (1997): “UK Stock Returns and Robust Tests of Mean Variance Efficiency,” *Journal of Banking and Finance*, 21, 641 – 660.
- COHEN, K., S. MAIER, R. SCHWARTZ, AND D. WHITCOMB (1978): “The Returns Generating Process, Returns Variance and the Effect of Thinness in Securities Markets,” *Journal of Finance*, pp. 149–167.

- DATAR, V. T., N. Y. NAIK, AND R. RADCLIFFE (1998): “Liquidity and Stock Returns: An Alternative Test,” *Journal of Financial Markets*, 1, 203 – 219.
- DELONG, J., A. SCHLEIFER, L. SUMMERS, AND R. WALDMAN (1990): “Positive Feedback Investment Strategies and Destabilising Rational Speculation,” *Journal of Finance*, 45, 379 – 395.
- DIMSON, E. (1979): “Risk Measurement When Shares are Subject to Infrequent Trading,” *Journal of Financial Economics*, 7, 197–226.
- FAMA, E. (1965): “The Behaviour of Stock Market Prices,” *Journal of Business*, 38, 34 – 105.
- FISHER, L. (1966): “Some New Stock Market Indices,” *Journal of Business*, 39, 191 – 225.
- FOERSTER, S., AND D. KEIM (1993): “Direct Evidence of Non-Trading of NYSE and AMEX Stocks,” *University of Pennsylvania*, (19-93).
- HAMILTON, J. (1995): *Time Series Econometrics*. Oxford University Press.
- KADLEC, G. B., AND D. M. PATTERSON (1999): “A Transactions Data Analysis of Nonsynchronous Trading,” *Review of Financial Studies*, 12(3), 609–630.
- KEIM, D., AND R. STAMBAUGH (1986): “Predicting Returns in the Stock and Bond Markets,” *Journal of Financial Economics*, pp. 357 – 390.
- LANGE, S. (1999): “Modeling Asset Market Volatility in a Small Market: Accounting for Non-Synchronous Trading Effects,” *Journal of International Financial Markets, Institutions and Money*, 9, 1 – 18.
- LBS (1996): “London Business School, London Share Price Database Reference Manual,” .
- LO, A., AND A. MACKINLAY (1990): “An Econometric Analysis of Nonsynchronous Trading,” *Journal of Econometrics*, 45, 181–211.

- LSE (1981): “The Stock Exchange Official List - Reporting and Publication of Bargains: Amendments to Rules,” *Stock Exchange Weekly Official Intelligence*, 311, 108 – 109.
- MECH, T. (1993): “Portfolio Return Autocorrelation,” *Journal of Financial Economics*, 34, 307–344.
- MILLER, M., J. MUTHUSWAMY, AND R. WHALEY (1994): “Mean Reversion of Standard and Poor’s 500 Index Basis Changes: Arbitrage-Induced or Statistical Illusion?,” *Journal of Finance*, pp. 479–517.
- MUTHUSWAMY, J. (1990): “Nonsynchronous Trading and the Index Autocorrelation Problem,” Ph.D. thesis, Graduate School of Business, University of Chicago.
- SCHOLES, M., AND J. WILLIAMS (1977): “Estimating Betas from Nonsynchronous Data,” *Journal of Financial Economics*, 5, 309–327.
- STOLL, H., AND R. WHALEY (1990): “The Dynamics of Stock Index and Stock Index Futures Returns,” *Journal of Financial and Quantitative Analysis*, pp. 441–468.

Table 1 : Proportions Traded Per Day

Days (t) since last trade

	$0 \leq t < 1$	$1 \leq t < 2$	$2 \leq t < 3$	$3 \leq t < 4$	$4 \leq t < 5$	$5 \leq t < 10$	$10 \leq t < 15$	$15 \leq t < 20$	$20 \leq t < 25$	$25 \leq t < 30$
A: Empirical Proportions										
Highest 1	98.74	0.57	0.09	0.14	0.07	0.13	0.04	0.02	0.01	0.01
2	91.81	4.76	0.83	0.81	0.50	0.72	0.16	0.07	0.05	0.03
3	81.36	10.19	2.23	1.93	1.19	2.16	0.29	0.13	0.10	0.04
4	73.62	12.05	3.63	2.83	1.92	3.96	0.78	0.29	0.27	0.09
5	66.74	14.17	4.30	3.65	2.42	5.64	1.35	0.49	0.34	0.13
6	61.65	14.66	4.80	3.84	2.69	7.29	1.85	0.78	0.58	0.25
7	57.21	14.40	5.35	4.11	3.08	8.67	2.45	1.06	0.91	0.44
8	52.14	14.97	5.68	4.60	3.44	10.14	3.20	1.37	1.09	0.44
9	46.30	14.53	6.02	4.74	3.72	11.85	4.08	1.93	1.69	0.66
Lowest 10	36.94	13.28	6.15	4.71	4.00	13.82	5.09	2.74	2.50	1.05
Lowest US(*)	73.00	12.00	6.00	4.50	4.50					
B: Time Independent Proportions										
Highest 1	98.74	1.24	0.02	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2	91.81	7.52	0.62	0.05	0.00	0.00	0.00	0.00	0.00	0.00
3	81.36	15.17	2.83	0.53	0.10	0.02	0.00	0.00	0.00	0.00
4	73.62	19.42	5.12	1.35	0.36	0.09	0.00	0.00	0.00	0.00
5	66.74	22.20	7.38	2.45	0.82	0.27	0.00	0.00	0.00	0.00
6	61.65	23.64	9.07	3.48	1.33	0.51	0.00	0.00	0.00	0.00
7	57.21	24.48	10.48	4.48	1.92	0.82	0.01	0.00	0.00	0.00
8	52.14	24.95	11.94	5.72	2.74	1.31	0.03	0.00	0.00	0.00
9	46.30	24.86	13.35	7.17	3.85	2.07	0.09	0.00	0.00	0.00
Lowest 10	36.94	23.20	14.69	9.26	5.84	3.68	0.37	0.04	0.00	0.00
Lowest US(*)	73.00	19.71	5.32	1.44	0.39					

(*) These figures come originally from Foerster and Keim (1993) and are presented in Lo and MacKinlay (1990) and Boudoukh, Richardson and Whitelaw (1994).

Table 2:AR and MA Coefficients						
Portfolio		1	2	3	LR Test	
Largest	ACF	0.004	-0.144	-0.099	AR(1)	na
	PACF	0.004	-0.144	-0.100	MA(1)	na
2	ACF	0.081	-0.099	-0.111	AR(1)	0.312
	PACF	0.081	-0.106	-0.095	MA(1)	0.042
3	ACF	0.194*	-0.067	-0.090	AR(1)	0.598
	PACF	0.194*	-0.109	-0.057	MA(1)	0.046
4	ACF	0.253**	-0.010	-0.064	AR(1)	0.326
	PACF	0.253**	-0.079	-0.045	MA(1)	0.488
5	ACF	0.318**	0.024	-0.029	AR(1)	0.528
	PACF	0.318**	-0.086	-0.011	MA(1)	1.21
6	ACF	0.320**	0.075	-0.023	AR(1)	0.002
	PACF	0.320**	-0.030	-0.043	MA(1)	3.18
7	ACF	0.363**	0.108	0.029	AR(1)	0.03
	PACF	0.363**	-0.027	-0.002	MA(1)	4.83*
8	ACF	0.335**	0.138 [†]	0.028	AR(1)	1.08
	PACF	0.335**	0.029	-0.030	MA(1)	7.92**
9	ACF	0.408**	0.172 [†]	0.054	AR(1)	0.61
	PACF	0.408**	0.007	-0.023	MA(1)	11.48**
Smallest	ACF	0.450**	0.250**	0.152	AR(1)	5.22*
	PACF	0.450**	0.060	0.025	MA(1)	29.74**
Marginal Significance <1%** , 5%* ,10% [†]						

Table 3 Cross-Serial Correlations and Granger Causality Tests							
Cross-Serial Correlations					Granger Causality Tests		
					t-1		
Portfolio Number	1	5	10		1	5	10
	1	-0.00451	-0.0661	-0.121	-	0.628	1.083
t	5	0.296	0.313	0.178	0.426	-	0.00137
	10	0.418	0.562	0.450	0.181	1.809	-
					Distributed F(1,66)		

Table 4 Computed and Estimated Autocorrelation Coefficients		
Smallest Decile		
UK	AR(1)	% of estimated
Markov Time Independent	0.036	8.0
Markov Time Dependent	0.082	18.2
Empirical	0.112	24.9
Estimated	0.450	
US		
Markov Time Independent	0.013	6.2
Empirical	0.061	28.9
Estimated	0.211	
<p>The table presents the first-order autocorrelation coefficient for the return to the smallest value decile portfolio. In addition to those computed for the three methods of treating non-trading using equation (5.3), the estimated AR(1) coefficient for this decile return is reproduced from Table 2. The computed figures are also expressed as a percentage of the estimated value. The US figures are from Boudoukh, Richardson, and Whitelaw (1994)</p>		

Table 5 Computed and Actual Autocorrelation Coefficients : All Deciles			
Decile	Empirical Non –Trading Proportions calibrated autocorrelation coefficient	Estimated autocorrelation coefficient	%
	(a)	(b)	$100 \times (a)/(b)$
Largest	0.00126	0.004	31.5
2	0.00697	0.081	8.6
3	0.0155	0.194	7.99
4	0.0267	0.253	10.6
5	0.0365	0.318	11.5
6	0.0474	0.320	14.8
7	0.0599	0.363	16.5
8	0.0702	0.335	21.0
9	0.0880	0.408	21.6
Smallest	0.112	0.450	24.9

The figures in column (a) are computed from equation (5.3) using the data on empirical non-trading proportions in Table 1. The estimated AR(1) coefficients in column (b) are reproduced from Table 2.

Table 6 Heterogeneity within the Smallest Value Decile		
Percentile	Proportion not traded on last day	Market Model β
91	0.565	0.898
92	0.586	0.844
93	0.592	0.891
94	0.597	0.959
95	0.612	0.902
96	0.615	0.907
97	0.639	0.741
98	0.680	0.789
99	0.668	0.847
Smallest	0.723	0.571
Decile Average		0.631
Computed Autocorrelation Coefficients from Empirical Non-Trading Proportions		
Homogenous Decile		0.0389
Heterogenous Decile		
	$\beta_i = 1$	0.0403
	Estimated β_i	0.0394
<p>These AR(1) coefficients are computed from equations (5.1) and (5.5) using the figures in the upper part of the table. The β's are estimated using the adjustment for non-synchronous trading proposed by Scholes and Williams (1977), see Clare, Morgan, and Thomas (2002) for discussion of the impact of non -synchronous trading on the estimation of β in this dataset.</p>		

Figure 6.1: Proportions Traded Day-by-Day : Smallest Decile

