Stock Portfolio Returns, Autocorrelations, and Nonsynchronous Trading

Rakesh Bharati, Susan J. Crain, and Manu Gupta

Abstract

Short-horizon portfolio return autocorrelation is often attributed to nonsynchronous trading. Nonsynchronous trading arises as last trades of most stocks occur randomly before the market close. However, prominent studies cast doubt on whether it can entirely explain the observed magnitude of index autocorrelation. We study if there is a market wide component to nonsynchronous trading, as it is often viewed as idiosyncratic to a stock. The results indicate that there is a significant market-wide factor and all stocks are sensitive to this factor. Further, higher size decile stocks are far more sensitive to the market-wide factor than lower size deciles. Theoretical models also imply that portfolio return autocorrelation should increase in lagged dispersion in last trade time; but we find no such sensitivity for equally and value weighted portfolios. Similar to the known relationship between stock price changes and volume, we find reliable evidence that longer trading days garner positive returns while shorter trading days yield negative or zero returns.

I. Introduction

Significant, positive portfolio autocorrelation in the short horizon has been a persistent anomaly since Cowles and Jones (1937). Since then, many studies have documented positive portfolio autocorrelations for stocks as well as other asset classes [e.g., see Khandani and Lo (2011) for evidence on emerging market stocks, corporate bonds, and mortgage-backed securities]. Fisher (1966) proposed that nonsynchronous trading (or nontrading) was the cause of this observed autocorrelation. The phenomenon occurs because all stocks do not last-trade exactly at 4:00 pm EST (the official close of the US markets). While there is some empirical support for Fisher's nonsynchronous trading based explanation [e.g., Boudoukh, Richardson, and Whitelaw (1994) and Ahn, Boudoukh, Richardson, and Whitelaw (2002)], several studies show that nonsynchronous trading only explains a small fraction of the observed magnitudes of autocorrelations [e.g., Perry (1985), Atchison, Butler and Simonds (1987), Lo and MacKinlay (1990), and Kadlec and Patterson (1999)]. Nevertheless, a variety of quantitative methods have been developed to minimize the impact of these autocorrelations in portfolio/index return parameter estimates [e.g., Hayashi and Yoshida (2005), Christensen, Kinnebrock and Podolskij (2010), and Griffin and Oomen (2011)]. These estimators risk discarding valuable information in autocorrelations if the autocorrelations have economically meaningful causes [e.g., Hou (2007)].

Alternative explanations of the autocorrelation are considered marginal. For example, time variation in expected returns due to a changing investment opportunity set can only be low-frequency and thus is not a credible explanation [see Boudoukh, Richardson, and Whitelaw (1994)]. Using index futures data, Ahn, Boudoukh, Richardson, and Whitelaw (2002) reject the bid-ask bounce as a cause since autocorrelation-implied bid-ask spreads are nearly ten times the actual spreads. Stock price discreteness, along with non-normality of returns, is ruled out as the cause by Kadlec and Patterson (1999). Therefore a deeper understanding of the anomaly is important as the observed autocorrelations appear to contradict the efficient markets hypothesis.

Rakesh Bharati is a finance professor at Southern Illinois University, Edwardsville, Illinois 62026, Susan J. Crain is a finance professor at Missouri State University, Springfield, Missouri 65897, and Manu Gupta is a finance professor at Virginia Commonwealth University, Richmond Virginia 23059. They can be contacted at rbharati@siue.edu, susaincrain@missouristate.edu, and mgupta@vcu.edu .

Availability of trade-by-trade data permits us to study the pattern of last trade time across stocks and see if it affects the serial correlations in two widely studied portfolios (equally-weighted and value-weighted market indices). First, using actual time stamps, we test if nonsynchronous trading is indeed idiosyncratic to a stock as presumed by most studies. Second, we test if the dispersion in last trade time predicts the level of serial correlation in the next-day return, as required by theoretical models. Finally, we investigate if mean last trade time is correlated with index returns. As Campbell, Grossman, and Wang (1993) note, "…volume tends to be higher when stock prices are increasing than when prices are falling." One should expect similar results with the mean trade time as, on high volume days, most stocks trade closer to the end of the day.

Except Kadlec and Patterson (1999), this is the only study to our knowledge to use the actual last trade time data to explore nonsynchronous trading. Unlike Kadlec and Patterson (1999), we use more recent data from a period that is twice as long. Further, while Kadlec and Patterson merely permit last trade time to be correlated across a group of securities in their simulation, we directly test if there is a market wide factor to nonsynchronous trading. To accomplish this goal, we characterize the distribution of the last time of trade across stocks with two market-wide measures where the first can be understood as the mean last trade time and the second is the dispersion in last trade time.

Our key result is that last trade time is positively correlated across stocks. This supports the presence of a systematic factor which causes stocks to last trade later or sooner. Thus, models of nonsynchronous trading need to assume a market wide factor affecting nonsynchronous trading. Further, a theoretical prediction is that thinner trading should lead to higher serial correlations [see, for example, Campbell, Lo, and MacKinlay (1997)]. We test this proposition by using the last trade time dispersion across the market as a proxy for thin trading and find that the thinness of trading does not seem to affect the strength of the serial correlation in equally and value weighted market portfolios. Finally, analogous to price-volume studies, we find that index returns are positively associated with the length of the trading day and those days with smaller mean last trade time have low or even negative index returns.

The paper is organized as follows. Section II presents a brief review of the literature. Section III discusses our data, methodology, and hypotheses. Section IV presents the results and section V concludes.

II. Literature Review

The fact that the close of the market is often not the time of the last trade for stocks was first analyzed by Fisher (1966). Subsequently, Perry (1985) showed that a large stock portfolio has an increase in serial correlation upon inclusion of additional securities, thus indicating that serial correlation is not merely an artifact of nonsynchronous trading. Atchison, Butler, and Simonds (1987) use the Scholes-Williams (1977) model to derive a theoretical value for portfolio autocorrelation and conclude that the empirically obtained values are well beyond what is attributable to nonsynchronous trading. Lo and MacKinlay (1990) also find that nonsynchronous trading does not explain a large fraction of the observed autocorrelations.

Boudoukh, Richardson, and Whitelaw (1994) allow for the heterogeneity of betas and nonsynchronous trading probabilities, as well as time dependence of nonsynchronous trading – issues not addressed by Lo and MacKinlay (1990). They show that, under such conditions, nonsynchronous trading can lead to comparatively larger autocorrelations than what was earlier predicted by Lo and MacKinlay (1990). In response, Kadlec and Patterson (1999) use five years of transactional data to obtain the last trade time for the securities to endogenize across-security correlation in nontrading intervals, time dependence and heterogeneity of trading, and difference in betas. They find autocorrelations in daily and weekly returns of large, small, and randomly selected stock portfolios to be larger than the simulated values.

Ahn, Boudoukh, Richardson, and Whitelaw (2002) find that index futures return autocorrelations are close to zero while the spot index return autocorrelations are expectedly positive. Further, the difference in the autocorrelations between the spot index returns and the index futures returns does not depend on transaction costs, according to their model. Thus, they claim support for the nonsynchronous trading based explanation. Clearly further explorations are necessary as the weight of the evidence suggests that a large part of return autocorrelation remains unexplained.

III. Data, Methodology, and Key Hypotheses

In selecting the period for the data, our goal is to select a long time series that is relatively free of regime shifts. The recent decade has seen momentous changes in the way securities are traded in the USA. Electronic Communications Networks (ECNs), offering faster trades, have lured institutional investors away from exchanges giving rise to a loss of liquidity. Market fragmentation due to ECNs, prominent users of ECNs (program and high-frequency traders), and dark pools (unregulated trading networks) are often quoted as new challenges for the markets.[1] Easley, Lopez de Prado, and O'Hara (2011) find pull-back of individual investors along with a reduction and concentration in liquidity providers, coupled with high-frequency trading, to be a major cause for the flash crash of May 6, 2010. The financial crisis of 2008 similarly created another externality. Other crucial changes include decimalization, Regulation SHO relaxing the long-existing regulations on short selling, and Regulation FD prohibiting selective disclosure of material information. Since these changes occurred at various time-points and their impact has been felt increasingly over time, it is not possible to choose a clear cut-off date without significantly reducing the sample. Still, it would be best to minimize the adverse impact of these changes on our estimates. In order to select a relatively uniform set of regimes, we cut off our sample at the end of the year 2002. Consequently, we sampled the Trades and Quotes Data (TAQ), from January 1993 through December 2002, to obtain the last trade time for listed securities on the New York Stock Exchange, American Stock Exchange and NASDAQ. In addition, we obtain the value- and equally-weighted index returns, volume, and turnover related information from the Center of Research in Security Prices (CRSP) database.

During 1993-2002, the market close was 4:00 PM EST and the open was 9:30 AM. The last trade and its time stamp are obtained from the TAQ data base. If no trade was found

[1] As Philips (2012) reports in a Businessweek on-line article, NYSE share of trade volume has gradually dropped from nearly 80 percent in the late 1990s to 22 percent in the early part of 2012. He attributes this decline to the emergence of competing ECNs and dark pools.

between 9:30 AM and 4 PM, no last trade time was recorded.[2] We converted the last trade time into a fraction of the trading day as shown below. A logit transformation of the fraction provides the last trade time measure (LTTM) for each security.

$$LTTM_{it} = \ln\left(\frac{x_{it}}{1 - x_{it}}\right) \tag{1}$$

where $x_{it} = \dfrac{Time\ in\ seconds\ from\ 9{:}30\ AM\ to\ the\ last\ trade\ for\ stock\ i\ on\ day\ t}{Total\ Number\ of\ Seconds\ between\ 9{:}30\ AM\ and\ 4{:}00\ PM}$

The logit transformation has desirable distributional properties, unlike the proportion which runs from zero to one. Finally, for each daily cross-section, we compute two statistics on LTTM across all available stocks – mean of the last trade time measure and the dispersion of the last trade time measure.[3] These two statistics essentially provide us with measures of market-wide nonsynchronous trading.

$$MLTTM_t = \frac{1}{N}\sum_{i=1}^{N} LTTM_{it} \tag{2}$$

$$DLTTM_t = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(LTTM_{it} - MLTTM_t)^2} \tag{3}$$

These two time series are used to test three hypotheses. The first hypothesis is a direct test of the common assumption in nonsynchronous trading models that nontrading probabilities are expected to be independent across stocks [for example, Lo and MacKinlay (1990)]. While it is a mathematically convenient assumption, there is no direct evidence on this crucial issue to date.

H1: The last trade time of individual stocks is independent of mean last trade time (MLTTM).

Next, we know that return autocorrelation is higher for the small-stock dominated equally-weighted market index compared to the value-weighted market index [see, for example, Atchison, Butler, and Simonds (1987)]. Small stocks tend to close much sooner than large stocks which we propose could be measured by the dispersion in the last trade time. Based on this reasoning, we should observe two patterns in data. First as noted above, we should observe lower serial correlation in the NYSE/AMEX/NASDAQ value-weighted portfolio while the NYSE/AMEX/NASDAQ equally-weighted portfolio should show strong positive serial correlation. Also, days with high dispersion in last trade time should demonstrate higher forward serial correlation. Days with low dispersion in last trade time, on the other hand, should have lower forward serial correlation as all stocks have nearly captured the same information, thereby reducing the predictability of the next day return using the previous day index return. This reasoning provides us with our second hypothesis.

[2] We dropped all dates when the market closed before 4 PM (26 days).

[3] McInish and Wood (1986) proxy thin trading by a similar variable based on the mean last trade time to market close in their comparison of beta adjustment techniques. Dimson and Marsh (1983) similarly employ a measure of trading infrequency with monthly returns to develop an adjustment to the systematic risk.

H2: Forward index return serial correlation should increase in dispersion in last trade time (DLTTM).

The third hypothesis is based on the intuitive similarities between our mean last trade time measure and the trading volume. Thus, days with later mean last trade time should demonstrate higher positive returns, compared to earlier mean last trade time days.

H3: Index return will increase in mean last trade time (MLTTM).

IV. Results

Figure 1 plots the values of the two measures, MLTTM and DLTTM.

Figure 1: Nonsynchronous Trading Measures

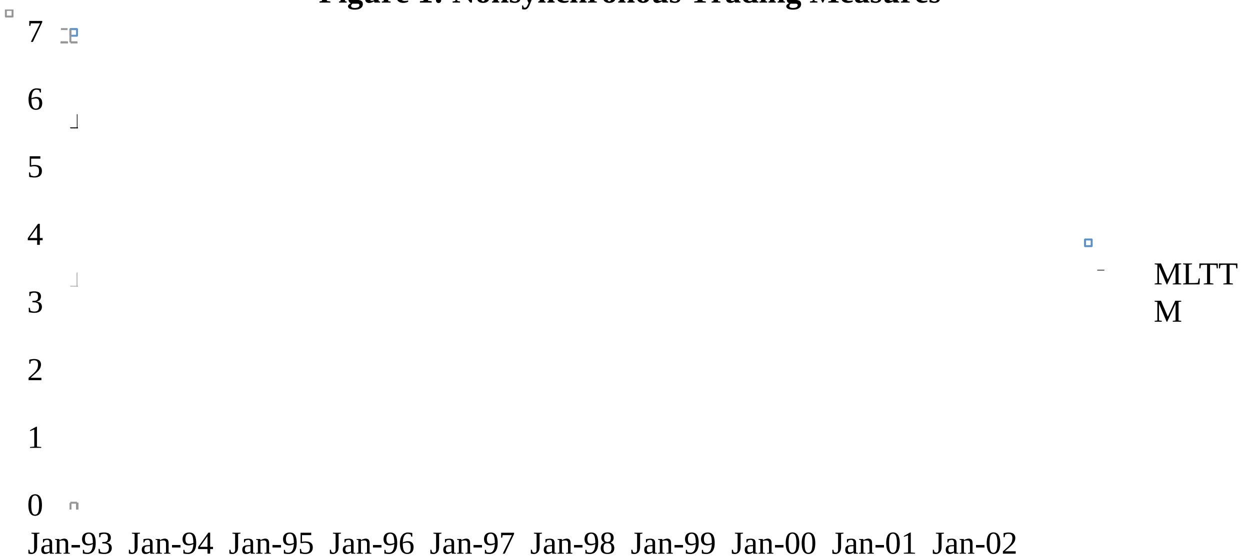

The mean last trade time and its dispersion rise over time indicating that stocks continue to trade later into the trading day, but last trade time dispersion also rises over time. Also, both time series show considerable variation.[4] As MLTTM appears intuitively similar to the volume related variables, we computed four market-wide volume proxies – average volume, average adjusted volume, average dollar volume and average turnover. We found them to only explain 15 percent of the total variation in our MLTTM measure.[5] Therefore, MLTTM appears to capture another dimension of market liquidity, of which the above four are also proxies. We now investigate if the mean last trade time measure is systematically associated with the last trade times across individual stocks.

The first hypothesis (H1) of idiosyncratic nonsynchronous trading implies that the Last Time to Trade Measure for a security (LTTM) and Mean Last Time to Trade Measure across securities (MLTTM) must be independent. We difference the two time series to obtain the

[4] We checked for changes in the two measures MLTTM and DLTTM immediately around the periods of change in the minimum tick size and found no significant difference. Therefore, tick size changes from 1/8th to 1/16th of a dollar and decimalization do not appear to have an impact on our measures, at least in the short term.

[5] Detailed descriptive statistics are available from the authors upon request.

change in the Last Trade Time Measure over its lag for the security and for the market. The resulting stationarity in the two time series helps minimize the probability of finding spurious correlations. We then regress the differenced LTTM (ΔLTTM) for each security on the differenced MLTTM (ΔMLTTM) and expect to find a coefficient of zero, consistent with H1. Table 1 reports average coefficients and proportions of positive and significant coefficients for Equation 4 by size deciles (as size is often a proxy for liquidity).

Table I
Regressions of Differenced Last Time to Trade Measure for each security (ΔLTTM) on the Differenced Mean Last Time to Trade Measure across securities (ΔMLTTM) by Size Deciles

$$\Delta LTTM = \beta_0 + \beta_1 \Delta MLTTM + \varepsilon$$

Size Decile	Proportion of stocks with positive β_1			Average β_1
		Significant at 95%	Significant at 90%	
Lowest	0.74	0.27	0.33	0.77
2	0.81	0.31	0.39	0.79
3	0.84	0.38	0.44	0.84
4	0.87	0.40	0.48	0.93
5	0.89	0.49	0.57	1.03
6	0.90	0.53	0.61	1.09
7	0.92	0.61	0.67	1.20
8	0.92	0.65	0.70	1.27
9	0.93	0.63	0.68	1.12
Highest	0.92	0.52	0.59	0.76

We fail to find support for the hypothesis of idiosyncratic nonsynchronous trading (H1) as there is a positive relationship between ΔLTTM and ΔMLTTM across deciles, indicating that there is a market-wide component to nonsynchronous trading. For example, 73.81% of the regressions at the lowest decile yielded positive coefficients and 33.46% (27.02%) were significant at the 90% (95%) confidence level. The average slope coefficient (β_1) on ΔMLTTM is 0.77 for the first decile. Moving up the size deciles, the proportion of positive slope coefficients, the proportion of significant coefficients, and the average value of the coefficients all monotonically increase. The peak occurs at the 8th decile where 65.34% of all regression coefficients are positive and significant at the 95% confidence level and the average slope coefficient is 1.27. Over 92% of all regression coefficients are positive.

Large stocks (5th decile and above) tend to be far more sensitive to this factor than small stocks. Interestingly, there is a decline in significances and average coefficient values as we move to the 9th and 10th deciles, with a large drop in the average β_1 for the highest decile. This may be due to the market-on-close orders related to S&P 500 index trading. Overall we find pervasive evidence of correlated last trade times across stocks, suggesting that nonsynchronous trading, at least in part, is driven by a market-wide phenomenon (perhaps arrival of economy wide information). This contradicts the typical assumption of idiosyncratic nonsynchronous trading, and lends support to the critique of Boudoukh, Richardson and Whitelaw (1994).

Dispersion in the Last Trade Time and Forward Index Return Correlation

As discussed earlier, models of nonsynchronous trading indicate that previous day trade time dispersion should lead to greater serial correlation in the next day index return (H2). This explains the high observed index return autocorrelation in the equally weighted index, while the correlation is nearly zero for the value weighted index (as larger stocks trade more frequently) [e.g., Perry (1985) and Atchison, Butler, and Simonds (1987)]. To test the second hypothesis, we assign each trading day into one of five groupings based on the dispersion in the last trade time (DLTTM) with the lowest dispersion days in the first quintile and the highest dispersion days in the fifth quintile. Thus as we move from the first to the fifth quintile, the sample has progressively higher dispersion in the last trade times. Next, we study the forward serial correlation of index returns. In other words, if the quintile assignment is made with the dispersion on day 0, the correlation between day 0 and day 1 is studied.

Table II presents serial correlation in value and equally weighted portfolio returns based on the sort by lagged dispersion in last trade time (DLTTM). The dispersion in the last trade time has risen over time. So, without an adjustment, we will simply end up with a sort where most recent data will be concentrated in the 5th quintile while the older data will fall in smaller quintiles. To counter this time trend, we adopt three different strategies. First, we perform a quintile sort within each calendar year and then merge the like quintiles across years. Second, we subtract the previous twenty-day average of DLTTM from the current value to eliminate the time trend. Third, we simply subtract the lagged value of DLTTM from the current value to eliminate the time trend.

Table II
Forward Serial Correlation in Value and Equally Weighted Index Returns for Quintiles formed on the Lagged Dispersion in the Nonsynchronous Trading Factor (DLTTM)

Quintile	$DLTTM_{t-1}$ sort by year			$DLTTM_{t-1} – SDLTTM_{t-i}/20$ where i = 2, 21			$DLTTM_{t-1}$ - $DLTTM_{t-2}$		
	$r_{t,t-1}$	p-value	R^2	$r_{t,t-1}$	p-value	R^2	$r_{t,t-1}$	p-value	R^2
Value-weighted Index Returns (Unconditional r = 0.04)									
1	-0.03	0.41	0.00	-0.09	0.05	0.01	-0.03	0.50	0.0
2	-0.04	0.38	0.00	0.09	0.05	0.01	0.05	0.25	0.0
3	0.11	0.04	0.01	0.05	0.28	0.00	0.04	0.37	0.0
4	0.18	<.00	0.03	0.09	0.04	0.01	0.17	0.00	0.03
5	0.09	0.06	0.01	0.10	0.03	0.01	0.05	0.37	0.00
Equally-weighted Index Returns (Unconditional r = 0.23)									
1	0.21	<.00	0.06	0.17	<.00	0.04	0.24	<.00	0.08
2	0.23	<.00	0.05	0.36	<.00	0.11	0.26	<.00	0.07
3	0.27	<.00	0.06	0.25	<.00	0.05	0.21	<.00	0.04
4	0.36	<.00	0.11	0.22	<.00	0.05	0.36	<.00	0.12
5	0.21	<.00	0.04	0.28	<.00	0.07	0.15	0.01	0.01

Across all three methodologies, the equally weighted portfolio return correlations are generally positive and significant. However, value weighted quintiles generally present insignificant correlations. This is to be expected for the two indices and has been documented in extant literature. However, going from the highest dispersion to lowest dispersion quintile for the equally weighted portfolio, there does not appear to be a systematic decline in correlation for all three methodologies, as one would expect under the second hypothesis. For the value weighted index, the correlation values for the bottom quintile are negative, while the higher quintiles generally have positive return correlations. However, most of these correlations are not significantly different from zero at the 95% confidence level. Thus our analysis fails to find support for the second hypothesis (H2).

Nonsynchronous Trading and Index Returns

We now test our third hypothesis (H3) that index returns are increasing in mean last trade time (MLTTM), similar to volume. In order to do so, we first sort trading days by mean last trade time. The upward trend of the MLTTM is handled in the same manner as was done for DLTTM above. Across the five quintiles, in Table III, we present the median return, the mean return and the p-value of the null hypothesis of the mean being equal to zero for the two indices.

Table III
Value and Equally Weighted Index Returns for Quintiles formed on Mean Last Trade Time Measure (MLTTM)

Quintile	$MLTTM_{t-1}$ sort by year			$MLTTM_{t-1}$ - $SMLTTM_{t-i}/20$ where i = 2, 21)			$MLTTM_{t-1}$ - $MLTTM_{t-2}$		
	Median	Mean	p-val	Median	Mean	p-val	Median	Mean	p-val
Value-weighted Index Returns									
1	-0.07%	-0.16%	0.00	-0.09%	-0.14%	0.00	-0.08%	-0.13%	0.00
2	0.01%	-0.04%	0.37	0.07%	-0.02%	0.57	-0.03%	0.01%	0.91
3	0.10%	0.09%	0.07	0.10%	0.11%	0.03	0.09%	0.06%	0.23
4	0.21%	0.13%	0.01	0.11%	0.11%	0.03	0.21%	0.13%	0.01
5	0.18%	0.18%	0.00	0.24%	0.15%	0.02	0.21%	0.15%	0.01
Equally-weighted Index Returns									
1	0.04%	-0.04%	0.11	0.02%	-0.04%	0.15	0.04%	-0.02%	0.47
2	0.11%	0.04%	0.24	0.15%	0.05%	0.09	0.11%	0.09%	0.01
3	0.19%	0.10%	0.00	0.20%	0.12%	0.00	0.20%	0.10%	0.01
4	0.28%	0.17%	0.00	0.28%	0.20%	0.00	0.26%	0.19%	0.00
5	0.34%	0.24%	0.00	0.29%	0.16%	0.00	0.33%	0.16%	0.00

The mean and median returns on the value-weighted index produce a remarkable pattern where the lowest quintile has consistently negative median, and mean returns with extremely strong significances. The mean return of -0.155% for the smallest quintile is contrasted with a 0.178% return for the highest quintile using the by-year sort.

The values produced by the other two sorting approaches are equally strong and are consistent with the by-year approach. There is a clear, monotonic pattern of increasing returns from the smallest to the largest quintile for all three approaches. The results for the equally-weighted index returns are similar, although the negative return for the lowest quintile is no longer significant at the conventional 95% level. Thus, the pattern of monotonically increasing returns in quintiles is strong and the average returns for the top three quintiles are large and strongly significant at the 90% confidence level or better for both indexes.

We interpret the results in Table III as strong support for H3. Longer trading days, on average, earn larger positive index returns while shorter days actually earn negative or zero returns. Therefore our mean last trade time measure (MLTTM) intuitively behaves in a way similar to market volume.

V. Conclusion

Nonsynchronous trading is commonly assumed to cause positive serial correlation in portfolio returns. However, a majority of studies have argued that the autocorrelations due to nonsynchronous trading are significantly less than the observed autocorrelations, which raises doubts about the efficient markets hypothesis. Most studies also assume that nonsynchronous trading induced autocorrelations are idiosyncratic, though there is no direct evidence on the issue. This study addresses the gap and adds to our understanding of market behavior by testing if nonsynchronous trading is indeed idiosyncratic. By using trade-by-trade data and collecting the last trade time for individual stocks, we are able to create statistical measures of market-wide nonsynchronous trading that can be used for further analysis, namely a mean last trade time measure (MLLTM) and dispersion of last trade time measure (DLLTM).

Our results indicate that the time of a stock's last trade is positively related to mean last trade time for the market – contrary to the assumption of idiosyncrasy. Thus, there appears to be a systematic factor which causes stocks to last trade later or sooner. Furthermore, higher size deciles are far more sensitive to this factor than lower size deciles. Next, nonsynchronous trading based models imply that portfolio return autocorrelation should increase with thin trading. However, we find no such increase for the equally and value-weighted portfolios using our proxy of the lagged dispersion in last trade time. This may indicate that return correlation may indeed be due to economic phenomena beyond nonsynchronous trading, furthering the argument for a systematic factor. Finally, volume is known to be positively related to contemporaneous price changes so longer trading days, as measured by mean last trade time, should have similar results. We find reliable evidence that days with later mean last trade time garner positive returns while those with earlier mean trade time yield negative or zero returns. Thus our proxy for volume appears to capture another dimension of market wide liquidity by finding evidence of a systematic factor that governs trade times.

For future research, it would be interesting to attempt to identify the systematic factor(s) that we have evidenced, perhaps by examining news releases of a broad economic nature.

References

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

Ahn, D-H, Boudoukh, J., Richardson, M. and Whitelaw, R. 2002, Partial Adjustment or Stale Prices? Implications from Stock Index and Futures Return Autocorrelations, Review of Financial Studies 15, 655-689.

Boudoukh, J., Richardson, M., and Whitelaw, R. 1994, A tale of three schools: insights on autocorrelations of short-horizon stock returns, Review of Financial Studies 7, 539-573.

Campbell, J., Grossman, S. and Wang, J. 1993, Trading Volume and Serial Correlation in Stock Returns, The Quarterly Journal of Economics 108, 905-939.

Campbell, J., Lo, A., and MacKinlay, A. 1997, The Econometrics of Financial Markets. Second printing, Princeton Press.

Christensen, K., Kinnebrock, S., & Podolskij, M. 2010, Pre-averaging estimators of the ex-post covariance matrix in noisy diffusion models with non-synchronous data, Journal of Econometrics, 159, 116-133.

Cowles, A., and Jones, H. 1937. Some *a posteriori* probabilities in stock return action. Econometrica 5, 280–94.

Dimson, E., and Marsh, P. 1983, The stability of U.K. risk measures and the problem of thin trading, Journal of Finance 38, 753-783.

Easley, D. Lopez de Prado, M, and O'Hara, M. 2011, The Volume Clock: Insights into the High Frequency Paradigm, Journal of Portfolio Management, 39, 19-29.

Fisher, L. 1966, Some new stock-market indexes, Journal of Business 16, 191-225.

Fowler, D. and Rorke, C. 1983, Risk Measurement When Shares are Subject to Infrequent Trading, Journal of Financial Economics 12, 279-283

Griffin, J. and Oomen, R. 2011, Covariance measurement in the presence of non-synchronous trading and market microstructure noise, Journal of Econometrics, 160, 58-68.

Hayashi, T. and Yoshida, N. 2005, On covariance estimation of non-synchronously observed diffusion processes, Bernoulli, 11, 359-379.

Hou, K. 2007, Industry Information Diffusion and the Lead-lag Effect in Stock Returns, The Review of Financial Studies, 20, 1113-1138.

Kadlec G. and Patterson D. 1999, A transactions data analysis of nonsynchronous trading, Review of Financial Studies 12, 609-630.

Khandani, A. and Lo, A. 2011, Illiquidity premia in asset returns: an empirical analysis of hedge funds, mutual funds, and US equity portfolios, The Quarterly Journal of Finance, 1, 205-264.

Lo, A. and MacKinlay, A. 1990, An econometric analysis of nonsynchronous trading, Journal of Econometrics 45, 181-211.

McInish, T. and Wood, R. 1986, Adjusting for beta bias: An assessment of alternate techniques: A note. Journal of Finance 41, 277-286.

Perry, P. 1985, Portfolio Serial Correlation and Nonsynchronous Trading, Journal of Financial and. Quantitative Analysis, 20, 517–523.

Philips, M. 2012, NYSE Dims the Lights to Attract Trading, Business Week (Markets and Finance Section, On-line edition), July 11.

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