

UC Berkeley

Other Recent Work

Title

Stock Return Autocorrelation is Not Spurious

Permalink

<https://escholarship.org/uc/item/9s35b82c>

Authors

Anderson, Robert M.
Eom, Kyong Shik
Hahn, Sang Buhm
et al.

Publication Date

2005-04-02

UNIVERSITY OF CALIFORNIA AT BERKELEY

Department of Economics

Berkeley, California 94720-3880

Working Paper No. E05-342

Stock Return Autocorrelation is Not Spurious

Robert M. Anderson

Departments of Economics and Mathematics
University of California, Berkeley

Kyong Shik Eom

Department of Economics
University of California, Berkeley

Sang Buhm Hahn

Korea Securities Research Institute

Jong-Ho Park

Department of Business Administration
Sunchon National University

April 2005

Keywords: Stock return autocorrelation, nonsynchronous trading, partial price adjustment, market microstructure, intraday return, SPDRs

JEL Classification: G12, G14, D40, D82

Abstract: We find compelling evidence that stock return autocorrelation is not spurious. Specifically, we find that partial price adjustment is an important source, and in some cases the main source, of the autocorrelation. In contrast to previous tests, our tests of partial price adjustment are direct, using disjoint time intervals, separated by a trade, to eliminate the nonsynchronous trading effect. We find compelling evidence of partial price adjustment in several settings, involving both individual stocks and portfolios. We also find evidence for partial price adjustment in an unlikely setting: the incorporation of very public, non-firm-specific information into the price of individual stocks. Several of our tests allow us to estimate lower bounds on the fraction of the autocorrelation that comes from partial price adjustment; in each case, we find the fraction is very substantial.

Stock Return Autocorrelation is Not Spurious

ROBERT M. ANDERSON

University of California at Berkeley, Departments of Economics and Mathematics

KYONG SHIK EOM*

University of California at Berkeley, Department of Economics

SANG BUHM HAHN

Korea Securities Research Institute

JONG-HO PARK

Sunchon National University, Department of Business Administration

ABSTRACT

We find compelling evidence that stock return autocorrelation is not spurious. Specifically, we find that partial price adjustment is an important source, and in some cases the main source, of the autocorrelation. In contrast to previous tests, our tests of partial price adjustment are direct, using disjoint time intervals, separated by a trade, to eliminate the nonsynchronous trading effect. We find compelling evidence of partial price adjustment in several settings, involving both individual stocks and portfolios. We also find evidence for partial price adjustment in an unlikely setting: the incorporation of very public, non-firm-specific information into the price of individual stocks. Several of our tests allow us to estimate lower bounds on the fraction of the autocorrelation that comes from partial price adjustment; in each case, we find the fraction is very substantial.

First Version: February, 2003

This Version: April 2, 2005

Key words: Stock return autocorrelation, nonsynchronous trading, partial price adjustment, market microstructure, intraday return, SPDRs

JEL Classification: G12, G14, D40, D82

* Corresponding Author: University of California at Berkeley, Department of Economics, 549 Evans Hall #3880, Berkeley, CA 94720-3880, Tel. #: +1-510-642-5248, Fax #: +1-415-621-7996, Email: kseom@econ.berkeley.edu. Anderson's research was supported by Grant SES-0214164 from the U.S. National Science Foundation. We are grateful to Jonathan Berk, Greg Duffee, Bronwyn Hall, Mark Rubinstein, Jacob Sagi, and Adam Szeidl for helpful comments.

There is an extensive literature concerning the autocorrelation of stock returns. The autocorrelation patterns of daily returns have been attributed to two main sources: spurious autocorrelation arising from market microstructure biases, including the nonsynchronous trading effect (in which correlations are calculated using stale prices) and bid-ask bounce, and genuine autocorrelation arising from partial price adjustment (i.e. trade takes place at prices that do not fully reflect the information possessed by traders). There has been considerable controversy over whether the autocorrelation is entirely spurious, in part because all of the tests to date have been indirect. In this paper, we propose and carry out several direct tests. The hypothesis that the autocorrelation of individual stock returns and portfolio returns is spurious is strongly rejected in all of our tests involving small and medium firms. In addition, we find compelling evidence of partial price adjustment by establishing that returns on an Exchange-Traded Fund (ETF) based on a broad index are positively correlated with subsequent returns on small, medium and large firms, even when the returns are calculated in a way that eliminates the use of stale prices. We conclude that partial price adjustment must be a significant source, and in some cases the main source, of autocorrelation in stock returns.

The principle underlying all of our tests is simple. By calculating returns on disjoint intervals separated by a trade, we can ensure that our correlations never involve stale prices, and hence we can eliminate the nonsynchronous trading effect. Once we eliminate the nonsynchronous trading effect, we are left only with partial price adjustment and bid-ask bounce. Since bid-ask bounce is negative, the absence of partial price adjustment implies that the covariance of returns across disjoint time intervals separated by a trade must be less than or equal to zero. Consequently, if we find that the covariance of returns across disjoint time intervals is positive, there must be a positive partial price adjustment effect.

Previous empirical work has found that daily individual stock returns exhibit either positive or negative autocorrelation; the average autocorrelation over all stocks is not statistically significant. By contrast, daily (and longer-term) returns on *portfolios* exhibit positive autocorrelation; this is large and very significant for portfolios of small and medium firms, while the findings on statistical significance are somewhat mixed for portfolios of large firms.

Three hypotheses have been advanced to explain the stylized facts just described (see Boudoukh, Richardson, and Whitelaw (1994))¹:

- market microstructure biases including the nonsynchronous trading effect (leading to negative daily autocorrelation for individual stock returns, positive daily autocorrelation for portfolio returns) and bid-ask bounce (leading to negative daily autocorrelation for both individual stock and portfolio returns²);

- time-varying risk premia;
- partial price adjustment: that some stock trades occur at prices that do *not* fully reflect the information available at the time of the trade. In principle, partial price adjustment could lead to either positive or negative daily autocorrelation. If informed traders strategically conceal their private information by carrying out a series of small trades, spread out in time, positive autocorrelation should result. However, uninformed traders may play positive-feedback strategies, which could lead to overshooting and negative autocorrelation (see Eom, Hahn, and Joo (2004)).

Time-varying risk premia have not been advanced as a source of daily return autocorrelation over periods of two years, the length of the subperiods considered in this paper. However, time-varying risk premia could induce a small bias in daily return autocorrelation. We analyze and estimate this bias, and demonstrate that it is too small to affect our results, in Appendix B. For the remainder of the body of the paper, we shall assume that the nonsynchronous trading effect, bid-ask bounce, and partial price adjustment are the only plausible candidates to explain autocorrelation in daily returns. All three candidates can be sources of negative autocorrelation.

For individual stock returns, partial price adjustment is the only plausible source of positive daily return autocorrelation. Consequently, in the absence of partial price adjustment, the return autocorrelation of every stock must be less than or equal to zero. Previous studies of individual stock return autocorrelation have focused on the average, over all stocks, of the autocorrelation coefficients and have obtained statistically weak results. In this paper, we test the autocorrelation, averaged over groups of stocks stratified by firm size, as well as the autocorrelation of individual firms. We strongly reject the hypothesis that the autocorrelation of every stock is less than or equal to zero. Therefore, partial price adjustment must be an important source of individual firm return autocorrelation.

For portfolios, the nonsynchronous trading effect and partial price adjustment are both plausible sources of positive daily return autocorrelation. However, by focusing on disjoint intervals separated by a trade, our tests avoid the use of stale prices, eliminating the nonsynchronous trading effect, and hence allow us to distinguish between the nonsynchronous trading and partial price adjustment effects as sources of positive portfolio autocorrelation. We find compelling evidence that partial price adjustment is an important source of portfolio return autocorrelation for small and medium firms. In addition, we find that returns on Standard and Poor's Depository Receipts (SPDRs), an ETF based on the Standard and Poor's 500 Index (S&P 500) are positively correlated with subsequent returns on small, medium, and large firms; since our methodology eliminates the nonsynchronous trading effect, this positive correlation

can only be attributed to partial price adjustment.

This paper makes two main methodological innovations. First, as described above, we choose the time intervals over which stock and portfolio returns are calculated in a way that avoids the use of stale prices and thus eliminates the nonsynchronous trading effect. Second, we develop a nonparametric test of return autocorrelation that does not depend on the assumption that returns are independent across stocks. Suppose that individual stock returns exhibit zero autocorrelation. Given 100 firms, let X be the number of firms whose sample autocorrelations are significantly positive at the 2.5% level. If stock returns were independent across firms, X would have the binomial distribution $B(100, 0.025)$ which has mean 2.5 and standard deviation approximately 1.56. Unfortunately, stock returns are not independent across firms, and this lack of independence means that the standard deviation of X is not readily ascertainable, and is almost certainly higher than that of the binomial. We *can* say, however, that X is a nonnegative integer-valued random variable with mean 2.5. If we test in n disjoint time periods, we obtain n independent observations of X , with order statistics denoted X_1, \dots, X_n . Suppose the realizations of X_1 and X_2 are x_1 and x_2 respectively. We can readily calculate $p_1 = P(X_1 > x_1)$ and $p_2 = P(X_2 > x_2)$; we use a test combining p_1 and p_2 to test whether stock returns exhibit nonpositive autocorrelation. In this paper, we break 10 years of intraday data from the New York Stock Exchange (NYSE) into five two-year periods ($n = 5$). The number of independent observations of X is thus quite small. Remarkably, despite the small number of independent observations, we strongly reject the hypotheses that stock and portfolio return autocorrelation is spurious. We also conducted tests with 10 one-year periods ($n = 10$) and found qualitatively similar results.³

Several of our tests, for both individual stocks and portfolios, allow us to establish lower bounds on the fraction of the autocorrelation coming from partial price adjustment.

All of our tests make use of tests for correlation: the Pearson correlation test; a modified version of the Pearson test, due to Andrews (1991), which takes into account heteroskedasticity and autocorrelation of errors; and the nonparametric Kendall tau test.

We also find partial price adjustment in an unexpected place. We test the cross-autocorrelation of individual stock returns to the returns of SPDRs, in a way that eliminates the nonsynchronous trading effect. Because the information contained in the SPDRs price is very public, rather than private, and is not firm-specific, these tests capture only a small fraction of the total partial price adjustment. Nonetheless, we find that this small piece of the total partial price adjustment is an important source of the autocorrelation. The fact that we find partial price adjustment in this unlikely setting indicates how pervasive it must be.

The following subsections outline our tests and findings.

A. Individual Stock Returns

The nonsynchronous trading effect and bid-ask bounce both predict negative autocorrelation in daily individual stock returns. Previous studies have tested the average daily return autocorrelation over all stocks in a given market, and have generally found that average to be statistically insignificant. Chan (1993) reports average daily return autocorrelation in deciles of NYSE and AMEX stocks, finding that average autocorrelations are negative and significant for small firms; insignificant for medium firms; and positive and significant for large firms, while noting that the grand average among firms of all size is not significant. We shall explain why Chan's data provide strong evidence for partial price adjustment for large firms in the 1980s, even though the grand average autocorrelation is not significant. We carry out a similar analysis, and find strong evidence of partial price adjustment for small firms in the period 1993-2002.

While previous work has focused on the average daily return autocorrelation, we look at individual stocks, and test the hypothesis that each stock's daily return autocorrelation is nonpositive. We find that this hypothesis is rejected among small and medium firms; highly significant numbers of small firms, and significant numbers of medium firms, exhibit positive daily return autocorrelation. For these firms, there must be something other than nonsynchronous trading and bid-ask bounce at work, and partial price adjustment is the only plausible candidate. We conclude that partial price adjustment is present in these firms, and that its effect is larger than the combined effect of nonsynchronous trading and bid-ask bounce.

We also find that a significant fraction of medium firms exhibit negative autocorrelation in daily returns.⁴ This finding could be explained by the nonsynchronous trading effect and bid-ask bounce alone. In the light of the finding just described, however, it is more plausible that the negative autocorrelation results from negative partial price adjustment, or from positive partial price adjustment that only partially offsets the nonsynchronous trading effect and bid-ask bounce.

Earlier research that focused on the average autocorrelation over all stocks missed strong evidence of partial price adjustment contained in the average daily return autocorrelation of size-related groups of firms, and in the daily return autocorrelation of individual firms.

In addition to studying the autocorrelation of conventional daily stock returns, we study the autocorrelation of the intraday return. The *intraday return* of a stock is defined as the price at the last trade of the day, minus the price at the first trade of the day, divided by the price at the first trade of the day. The use of intraday returns eliminates the nonsynchronous trading effect and eliminates or greatly reduces the bid-ask bounce effect in stock return autocorrelation. The use of intraday returns also

eliminates some of the partial price adjustment effect in stock return autocorrelation, but provides us with a direct measure of the remaining portion of the autocorrelation. We find that the *average* intraday return autocorrelation over the ten-year data period is positive and highly significant among small and medium firms; this is strong evidence of the existence of partial price adjustment, and that it is positive on average. For large firms, the average intraday return autocorrelation over the entire ten-year period is negative and statistically significant at the 1% level; among our five two-year subperiods, it is significant and negative in two subperiods, significant and positive in one subperiod, positive but not significant in one subperiod, and negative but not significant in the remaining subperiod. We argue that this pattern comes most plausibly from partial price adjustment which varies in sign from period to period. This could possibly be explained by variations over time in the number of traders using momentum strategies in trading large stocks.

For small and medium firms, we strongly reject the hypothesis that the intraday return autocorrelation of *each* firm is less than or equal to zero. For small, medium, and large firms, the hypothesis that the autocorrelation of intraday return of every stock is nonnegative is not rejected; indeed, the fraction of stocks with negative intraday return autocorrelation is systematically below the expected value of 2.5%. We conclude that the partial price adjustment effect is systematically positive for small and medium stocks. As with the average intraday return autocorrelation, the number of positive and negative return autocorrelations among large firms varies from period to period.

Our methods allow us to estimate a lower bound on the portion of the identifiable absolute autocovariance arising from partial price adjustment; we obtain values ranging from 52.6% to 60.7%.

B. Portfolio Returns

Bid-ask bounce predicts slightly negative autocorrelation in daily portfolio returns, while the nonsynchronous trading effect predicts positive autocorrelation; partial price adjustment can be either positive or negative. Thus, the finding that daily portfolio returns exhibit positive autocorrelation is not sufficient to establish the existence of partial price adjustment. A number of papers have carried out indirect tests that tend to support the presence of partial price adjustment, but the results have been controversial because of the indirect nature of the tests.

In this paper, we conduct two direct tests of nonsynchronous trading and partial price adjustment in explaining the positive autocorrelation in daily portfolio returns.

In the first test, we define the intraday return of a portfolio as the equally-weighted average of the intraday returns of the individual stocks in the portfolio. Since the intraday return of a portfolio on a given day depends only on trades that occur that day, the nonsynchronous trading effect is eliminated.

We find that the autocorrelation of conventional portfolio returns is positive and strongly significant for small, medium and large firms. The autocorrelation of intraday portfolio returns is positive and highly significant for small and medium firms, providing strong evidence of partial price adjustment. The autocorrelation of intraday portfolio returns is negative and not significant for large firms.

In the second test, we compute the cross-autocorrelation of daily returns on SPDRs up to the time of the last trade of a given stock, and that stock's next-day conventional return. In this setting, the nonsynchronous trading effect is eliminated, but bid-ask bounce remains and is negative. Our null hypothesis is that the autocorrelation is less than or equal to zero for every stock in the portfolio. This hypothesis is strongly rejected for all three portfolios and all three (Pearson, Andrews' modification of Pearson, and Kendall's tau) correlation tests.

Our method allows us to estimate a lower bound on the proportion of portfolio return autocorrelation arising from partial price adjustment: 54.58% for small, 59.54% for medium firms, and 36.82% for large firms.

The remainder of this paper is organized as follows. Section I details our methodology and null hypotheses. Section II describes the sampling of firms and provides descriptive statistics of our data. Section III presents and interprets the empirical results. Section IV provides a summary of our results and some suggestions for further research.

I. Methodology

As noted by Lo and MacKinlay (1990), the nonsynchronous trading effect arises from measurement error in calculating stock returns. If a security does not trade on a given day, its daily return is reported as zero; if it does not trade for several days, it is in effect accumulating several days of unreported gain or loss, which is captured in the data on the first subsequent day on which trade occurs. Think of the "true" price of the stock being driven by a positive trend plus a daily volatility term (with mean zero), with the reported price being updated only on those days on which trade occurs. On days on which no trade occurs, the reported return will be zero, which is below trend; on days on which trade occurs after one or more days without trade, the reported return represents several days worth of trend; this results in spurious negative autocorrelation. Even if a stock does trade on a given day, the reported "daily closing price" is the price at which the last transaction occurred, even if the last transaction occurred long before the market closed. Thus, a single piece of information that affects the underlying value of stocks i and j may be incorporated into the reported price of i today because i trades after the information is revealed,

but not incorporated into the reported price of j until tomorrow because j has no further trades today, resulting in a positive cross-autocorrelation between the prices of i and j . Hence, the nonsynchronous trading effect may cause spurious negative individual autocorrelation and positive individual cross-autocorrelation, resulting in positive autocorrelation of portfolios.

Since Fisher (1966) and Scholes and Williams (1977) first pointed out the nonsynchronous trading effect, the extent to which nonsynchronous trading can explain autocorrelation has been extensively studied, but remains very controversial.

The first main idea in this paper is to study stock returns over *disjoint* time intervals where a *trade occurs between the intervals*. More formally, we study the correlation of stock returns over intervals $[s,t]$ and $[u,v]$ with $s < t \leq u < v$ such that the stock trades at least once on the interval $[t,u]$. We apply this idea to derive tests in a number of different situations. Because these correlation calculations do not make use of stale prices, the nonsynchronous trading effect is, *by definition*, eliminated; if the correlation turns out to be nonzero, there must be a source, other than the nonsynchronous trading effect, for the correlation. This conclusion does not depend on any particular story of how the use of stale prices results in spurious correlation.

In addition to eliminating the nonsynchronous trading effect, our method of calculating correlations also eliminates or greatly reduces the bid-ask bounce effect in many of the situations in which we apply it.

We say that a stock exhibits partial price adjustment if there are trades at which the trade price does not fully reflect the information available at the time of the trade. Let r_{sti} denote the return on stock i ($i=1, \dots, J$) over the time interval $[s,t]$; in other words, $r_{sti} = \frac{S_i(t)}{S_i(s)} - 1$, where $S_i(t)$ is the price of stock i at the last trade occurring at or before time t . Let F_t denote the σ -algebra representing the information available at time t . Since the stock price at each trade is observable, $S_i(t)$ must be F_t -measurable. The absence of partial price adjustment in stock j implies the following:⁵

given times $s < t \leq u < v$ such that stock j trades at some time $w \in [t,u]$, r_{uj} is uncorrelated with every random variable which is F_w -measurable, and hence uncorrelated with r_{stj} .

Thus, we can test for the presence or absence of partial price adjustment by examining return correlations over time intervals $[s,t]$ and $[u,v]$ satisfying the condition just given.

Any single correlation can be tested by a variety of standard methods; the three methods we use are outlined below. Our method requires testing many correlations simultaneously. Depending on the test, we reject the hypothesis for an individual firm if the sample correlation falls outside either a symmetric 95%

confidence interval, or if the sample correlation lies in the upper (lower) 2.5% tail of the distribution. If the correlation tests were independent across firms, the number of rejections would have the binomial distribution. If the collection $\{r_{s_i t_i} : i = 1, \dots, I\}$ were a family of independent random variables, then X , the number of firms for which the zero-correlation hypothesis is rejected at the 5% (2.5%) level, would be binomially distributed, as $B(I, 0.05)$ ($B(I, 0.025)$), which has mean $0.05I$ ($0.025I$) and standard deviation $\sqrt{(0.05)(.95)I}$ ($\sqrt{(0.025)(.975)I}$). Since returns are not independent across stocks, X will not be binomial; its standard deviation is not readily ascertainable, and is likely higher than that of the binomial. However, we do know that X is a nonnegative, integer-valued, random variable with mean $0.05I$ ($0.025I$); this is our null hypothesis.

In all our tests, $I=100$, so X has mean $\mu=5$ or $\mu=2.5$. Since X is nonnegative, $P(X \geq \alpha\mu) \leq 1/\alpha$ for every $\alpha \geq 1$. Suppose that we compute X in each of n disjoint time periods. This provides us with n independent observations of X ; let X_1, \dots, X_n be the order statistics. Then for every $\alpha \geq 1$, $P(X_1 \geq \alpha\mu) \leq 1/\alpha^n$ and $P(X_2 \geq \alpha\mu) \leq 1/\alpha^n + n(1-1/\alpha)/\alpha^{n-1} = (n\alpha - (n-1))/\alpha^n$. Given particular realizations $x_1 \geq \mu$ and $x_2 \geq \mu$ of X_1 and X_2 , we obtain p -values of $p_1 = 1/(x_1/\mu)^n$ for x_1 and $p_2 = (n(x_2/\mu) - (n-1))/(x_2/\mu)^n$ for x_2 , respectively.

If we were to test using X_1 alone, the result could be strongly affected by a single outlier. In particular, if any single realization of X is less than μ , then $p_1=1$ and the null hypothesis will not be rejected. If we were to test using X_2 alone, the result would not be affected by a single outlier, but the test is less powerful than the test using X_1 in many situations. For this reason, we adopt the following combined test using both X_1 and X_2 . Compute the statistic $p_3 = 2 \min\{p_1, p_2\}$. Note that for any γ , $P(p_3 \leq \gamma) = P(2 \min\{p_1, p_2\} \leq \gamma) = P(p_1 \leq \gamma/2 \text{ or } p_2 \leq \gamma/2) \leq P(p_1 \leq \gamma/2) + P(p_2 \leq \gamma/2) = \gamma/2 + \gamma/2 = \gamma$. Thus, the p -value in the combined test is $p_3 = 2 \min\{p_1, p_2\}$. Note that p_3 depends on μ and n .

Two of our tests focus on what we call intraday returns; in these tests, bid-ask bounce is eliminated or greatly reduced. The intraday return of a stock on a given day is defined as the price of the last trade of the day, less the price at the first trade of the day, divided by the price at the first trade of the day. Thus the intraday return of stock i on day d is $r_{s_i t_i}$, where s_i and t_i are the times of the first and last trades of the stock on day d .⁶ We compute the correlation $\rho(r_{s_i t_i}, r_{u_i v_i})$, where u_i and v_i are the times of the first and last trades on day $d+1$. Note that $s_i < t_i < u_i < v_i$. Bid-ask bounce arises in conventional daily return autocorrelation because the correlation considered is $\rho(r_{w_i t_i}, r_{t_i v_i})$, where w_i is the time of the last trade prior to day d . Note that the end time in calculating $r_{w_i t_i}$ is the same as the starting time in calculating $r_{t_i v_i}$, resulting in negative

autocorrelation, as explained in Roll (1984); Roll’s model assumes that at each trade, the toss of a fair coin determines whether the trade occurs at the bid or ask price. In the calculation of the intraday autocorrelation, the end time t_i of the first interval is different from the starting time u_i of the second interval. Moreover, the trades at t_i and u_i are different trades, so the coin tosses for these trades are independent; if we apply Roll’s model to this situation, the autocorrelation resulting from bid-ask bounce is zero. If we extend Roll’s model to multiple stocks, and assume the coin tosses are independent across stocks, the autocorrelation and cross-autocorrelation of intraday stock returns are zero. Relaxing the independence assumption results in slightly negative autocorrelation and cross-autocorrelation of intraday returns.⁷

A. Individual Stock Returns

Studies of autocorrelation in individual stock returns have focused on the average autocorrelation of groups of firms, finding it to be statistically insignificant and usually positive; see Säfvenblad (2000) for a survey. This finding is consistent with two possibilities: either the autocorrelation of each individual stock is essentially zero; or some stocks exhibit positive autocorrelation and others exhibit negative autocorrelation, with the two largely canceling out when averaged over stocks. None of the previous studies analyzed the autocorrelation of individual stocks one by one. Doing so is essential for testing whether the autocorrelation arises from nonsynchronous trading or partial price adjustment, or both.

In this paper, we average the autocorrelation over groups of firms, segregated by firm size, and also consider the autocorrelations of individual firms. We calculate the autocorrelation in two different ways: the conventional daily return autocorrelation, and the intraday return autocorrelation.

A.1. Conventional Daily Return Autocorrelation

For each firm, we calculate the daily return on each day in the conventional way: the closing price on day d , minus the closing price on the last day prior to day d on which trade occurs, divided by the closing price on the last date prior to day d on which trade occurs. When we compute returns in the conventional way, the nonsynchronous trading effect is present; under the usual story for the nonsynchronous trading effect on the autocorrelation of individual firm returns, the nonsynchronous trading, like the bid-ask bounce effect, predicts negative autocorrelation of individual stock returns. Hence, in the absence of partial price adjustment, every firm should exhibit daily return autocorrelation less than or equal to zero, and consequently the average daily return autocorrelation in a group of stocks must be less than or equal to zero. Null Hypothesis I (IA, IB) is that the *average* daily return autocorrelation is zero (nonpositive, nonnegative) in each firm group; we test these hypotheses by comparing the average daily return autocorrelation to the associated standard error. Rejection of Null Hypothesis IA implies that there is partial price adjustment, and

it is positive on average. Rejection of Null Hypothesis IB implies that on average, the sum of the nonsynchronous trading, bid-ask bounce, and partial price adjustment effects is negative. Since we selected our sample firms separately in each of the five two-year subperiods, we compute the average autocorrelation and associated standard error in each subperiod. We then compute the average autocorrelation over the whole ten-year period as a weighted average of the period averages, weighting each period by the inverse square of the associated standard error. We report the averages and standard errors for the individual subperiods, as well as the weighted average and standard error over the whole ten-year period.

Null Hypothesis II is that *every* firm's daily returns exhibit zero autocorrelation. For each firm, we test whether daily returns exhibit zero autocorrelation, in each of $n=5$ disjoint two-year subperiods. In each subperiod, we test whether the sample autocorrelation of each stock lies in a symmetric 95% confidence interval that places 2.5% probability on each of the two tails, so $\mu=5$. Applying this test to 100 firms in each of the 5 subperiods, we reject Null Hypothesis II if $p_3 = p_3(\mu=5, n=5) < 0.05$. Null Hypothesis IIA is that each firm exhibits nonpositive autocorrelation. For each firm, we test whether daily returns exhibit nonpositive autocorrelation, in each of the 5 disjoint two-year subperiods. In each subperiod, we test whether the sample autocorrelation lies below the 97.5 %-ile, so $\mu=2.5$.⁸ Applying this test to 100 firms in each of the 5 disjoint subperiods, we reject Null Hypothesis II if $p_3 = p_3(\mu=2.5, n=5) < 0.05$. Rejection of Null Hypothesis IIA implies that in at least some firms, the partial price adjustment effect is positive and, indeed, is larger than the sum of the nonsynchronous trading and bid-ask bounce effects.

Null Hypothesis IIB is that each firm exhibits nonnegative autocorrelation. For each firm and each subperiod, we test whether the sample autocorrelation lies below the 2.5 %-ile; we reject Null Hypothesis IIB if $p_3 = p_3(\mu=2.5, n=5) < 0.05$. Rejection of Null Hypothesis IIB implies that in some firms, the sum of the nonsynchronous trading, bid-ask bounce, and partial price adjustment effects is less than or equal to zero. This is consistent with there being a negative partial price adjustment effect for this group of stocks, or a positive partial price adjustment effect which is outweighed by the nonsynchronous trading and bid-ask bounce effects.

A.2. Intraday Return Autocorrelation

As above, we define the intraday return on day d as the price at the final trade on day d , minus the price at the first trade on day d , divided by the price at the first trade on day d . If a given stock does not trade, or has only one trade, on a given day, we drop the observation of that stock for that day from our data set.⁹ If we compare intraday returns on day d and day $d+1$, there is *no* nonsynchronous trading effect: the intraday returns are computed over disjoint time intervals, with each interval beginning and ending with a trade, so stale prices never enter the calculation. Moreover, because the first trade on day $d+1$ is a different trade

from the last trade on day d , bid-ask bounce is eliminated or greatly reduced. If the nonsynchronous trading effect and bid-ask bounce are the sole sources of stock return autocorrelation, the theoretical autocorrelation of intraday returns on each stock must be less than equal to zero, and close to zero. This implies that the average autocorrelation of intraday returns on each group of stocks must be less than or equal to zero and close to zero. Null Hypothesis III (IIIA, IIIB) is that the *average* autocorrelation of intraday returns is zero (nonpositive, nonnegative) in each group of stocks; we test these hypotheses by comparing the average daily return autocorrelation to the associated standard error. Rejection of Null Hypothesis III implies that there is partial price adjustment. Rejection of Null Hypothesis IIIA implies that there is partial price adjustment, and it is positive on average. Rejection of Null Hypothesis IIIB implies that there is partial price adjustment, and that it is negative on average. As in the case of conventional returns, we report the average daily return autocorrelations in each of the five two-year subperiods, as well as the weighted average of the subperiod returns, along with the associated standard errors.

Our Null Hypothesis IV is that the autocorrelation of intraday returns on *each* stock is zero, Null Hypothesis IVA is that the autocorrelation of intraday returns on each stock is nonpositive, and Null Hypothesis IVB is that the autocorrelation of intraday returns on each stock is nonnegative. The criteria for rejection of Null Hypotheses IV, IVA, and IVB are identical to those of Null Hypotheses II, IIA and IIB, except that we use intraday returns rather than conventional daily returns. In particular, we divide our data period into five two-year subperiods, and test using $p_3(\mu=5, n=5)$ for Null Hypothesis IV, and $p_3(\mu=2.5, n=5)$ for Null Hypotheses IVA and IVB. Because intraday returns exhibit neither the nonsynchronous trading nor the bid-ask bounce effect, rejection of Null Hypothesis IV implies that partial price adjustment is a source of individual stock autocorrelation. Rejection of Null Hypothesis IVA implies that the autocorrelation arising from partial price adjustment is positive, arising from slow incorporation of information into prices more than from overshooting due to positive-feedback strategies.

A.3. Analysis of Autocovariance

In this section, we describe a method to obtain a lower bound on the portion of the individual stock autocovariance attributable to partial price adjustment. Conventional daily returns are calculated from the closing trade one day to the closing trade of the next day on which trade occurs; the union of these intervals, from one closing trade to the next, covers the whole time span of our data period. However, the intraday returns of the stocks are calculated over a portion of the total trading period, namely the union of the intervals of time beginning with the first trade of a stock on a day and the last trade of the same stock on that day; on each day, the time interval from the last trade of the day to the first trade on the next day is omitted. In all conventional models of stock pricing, the standard deviation of intraday return

should be lower than the standard deviation of conventional daily return. For example, if the stock price is any Itô Process, the price changes over the excluded intervals are uncorrelated with the price changes over the included intervals. Since the variance of a sum of uncorrelated random variables is the sum of the variances, the exclusion of the intervals must decrease the variance. Notice that this argument applies to the theoretical variance—the variance of the theoretical distribution of returns. The observed variance of returns for a given stock is the variance of a sample out of that theoretical distribution of returns, so the standard deviation of intraday return might be larger than the standard deviation of conventional daily return for a few stocks. In our sample, we find that only 21 of the 1500 stocks (300 stocks per subperiod times five subperiods) exhibits sample standard deviation of intraday returns greater than the sample standard deviation of conventional daily return.

For each stock, we can compute the conventional daily (intraday) return autocovariance by taking the product of the conventional daily (intraday) return autocorrelation times the conventional daily (intraday) return variance. Note that these autocovariances can be either positive or negative, so it is not appropriate to compute their ratio. However, we know that partial price adjustment is the only source of the intraday return autocovariance. If C_i and I_i denote the conventional daily and intraday return autocovariances of stock i , let $W_i = C_i - I_i$ denote the residual. W_i , C_i , and I_i may each be either

positive or negative. Thus, we consider $\frac{|I_i|}{|I_i| + |W_i|}$ as the fraction of the identifiable absolute autocovariance arising from intraday returns. This ratio is a lower bound on the portion of the identifiable return autocorrelation attributable to partial price adjustment. It understates the proportion of the autocorrelation attributable to partial price adjustment for two reasons. First, partial price adjustment can induce both negative and positive effects; these cancel, and we see only the net effect in this calculation. Second, partial price adjustment occurring between the last trade of a stock on a given day and the first trade on the next day is also omitted from this calculation.

B. Portfolio Returns

Atchison, Butler, and Simonds (1987) and Lo and MacKinlay (1990) find that nonsynchronous trading explains only a small part of the portfolio autocorrelation (16% for daily autocorrelation in Atchison, Butler, and Simonds, 0.07, a small part of the total autocorrelation, for weekly autocorrelation in Lo and MacKinlay). However, Boudoukh, Richardson, and Whitelaw (1994) find that the weekly autocorrelation attributed to nonsynchronous trading in a portfolio of small stocks is as high as 0.20 (56% of the total autocorrelation) when the standard assumptions by Lo and MacKinlay are loosened by considering heterogeneous nontrading

probabilities and heterogeneous betas.¹⁰ This leads them to conclude that “institutional factors are the most likely source of the autocorrelation patterns.”

The use of intraday data has led to renewed interest in this issue. For example, Ahn, Boudoukh, Richardson, and Whitelaw (2002) comment that Kadlec and Patterson (1999), using intraday data and simulation, find that “nontrading can explain 85%, 52%, and 36% of daily autocorrelations on portfolios of small, random, and large stocks, respectively. In other words, nontrading is important *but not the whole story* [italics added].” Ahn, Boudoukh, Richardson, and Whitelaw assert that the positive autocorrelation of portfolio returns “can most easily be associated with market microstructure-based explanations, as partial [price] adjustment models do not seem to capture these characteristics of the data.”

There has also been support for partial price adjustment. Chan (1993) provides a model in which there is a separate market-maker for each stock; each market-maker observes a signal of the value of his/her stock, and sets the price at the correct conditional expectation, given the signal, so that individual stock returns show no autocorrelation; and stock returns exhibit positive cross-autocorrelation, because the signals are correlated across stocks. He tests some predictions of this model, finding support for positive cross-autocorrelation, and for his prediction that the cross-autocorrelation is higher following large price movements. Chordia and Swaminathan (2000) compare portfolios of large, actively traded stocks, to portfolios of smaller, thinly traded stocks, arguing that the nonsynchronous trading effect should be more significant in the latter than in the former. The data they report on the autocorrelations of these portfolios “suggest that nontrading issues cannot be the sole explanation for the autocorrelations [...] and other evidence [concerning the rate at which prices of stocks adjust to information] to be presented.” Llorente, Michaely, Saar, and Wang (2002) relate the volume to the autocorrelation, arguing that the relative importance of hedging and speculative trading determines the direction of the relationship, with positive autocorrelation arising if speculative trading (in which informed agents slowly exercise their informational advantage) predominates.¹¹

While many papers have studied whether the nonsynchronous trading effect can fully explain positive portfolio autocorrelation, all of the tests have been indirect. In this paper, we propose and carry out two direct tests that eliminate the nonsynchronous trading effect. In both tests, we compute the correlation of returns of securities over disjoint time intervals separated by a trade, so that stale prices never enter the correlation calculation. If the nonsynchronous trading effect and bid-ask bounce are the sole explanations of stock return autocorrelation, the autocorrelation computed by our methods must be less than or equal to zero.

B.1. First Method, Intraday Returns

In the first method, we compute the intraday returns of each individual stock as defined in Section I.A.2. As noted there, intraday returns on different days do not exhibit the nonsynchronous trading effect, and bid-ask bounce should be eliminated or greatly reduced. We consider three portfolios, each containing 100 stocks, representing small, medium, and large market capitalization.

We define the intraday return of a portfolio on a given day as the equally-weighted average of the intraday returns for that day on all stocks in the portfolio, omitting those stocks which have fewer than two trades on that day. Note that the autocorrelation of the intraday return of the portfolio is just the average of the correlations of the intraday returns of the individual pairs (including the diagonal pairs) of stocks in the portfolio. Since 99% of these pairs are off-diagonal, the portfolio return autocorrelation is dominated by the cross-autocorrelations between pairs of stocks. In particular, the portfolio return autocorrelation is *not* the average of the individual return autocorrelations of the stocks in the portfolio.

If the nonsynchronous trading effect and bid-ask bounce are the sole sources of stock return autocorrelation, the autocorrelation of the intraday return of the portfolio must be less than or equal to zero, and close to zero. Thus, our Null Hypothesis V is that the autocorrelation of the intraday return of the portfolio is zero. Rejection of Null Hypothesis V implies that there is a nonzero partial price adjustment effect; the sign of the partial price adjustment effect is determined by the sign of the autocorrelation. As in the case of average individual stock autocorrelations, we test the portfolio return autocorrelation separately in each of our five two-year subperiods, and compute the portfolio return autocorrelation for the whole ten-year period as a weighted average of the subperiod results; we report both the subperiod results and the weighted average, along with the associated standard errors.

The computation of the autocorrelation of the intraday return of the portfolio allows us to obtain a lower bound on the portion of the conventional daily return autocorrelation attributable to partial price adjustment. As in Section I.A.3, all conventional models of stock pricing predict that the variance of intraday portfolio returns should be lower than the variance of conventional daily portfolio returns; we find that this is the case in each of the three portfolios and each of the five two-year subperiods in our data set. We calculate the autocovariance of conventional daily (intraday) portfolio returns by multiplying the conventional daily (intraday) autocorrelation of portfolio returns by the conventional daily (intraday) variance of portfolio returns. The residual is defined as the difference of the conventional and intraday autocovariances. The autocovariance of intraday portfolio returns can only come from partial price adjustment, so the ratio of the intraday autocovariance to the sum of the absolute values of the intraday and residual autocovariances gives a lower bound on the proportion of the autocorrelation that is attributable to partial price adjustment.

B.2. Second Method, ETFs

In the second method, we define our portfolio as the price of an ETF. ETFs are continuously-traded securities which represent ownership of the stocks in a particular mutual fund or index. Because a mutual fund is valued once a day, and an index is calculated at any given instant by averaging the most recent price of each stock in the index, and some of those prices are stale, the mutual funds and indices are themselves subject to the nonsynchronous trading effect. For example, the quoted value of the S&P 500 Index exhibits stale pricing because it is an average of the most recent trade price of the stocks in the Index (see Kimelman (2003)). ETFs are traded continuously and very actively, the value is updated continuously, rather than with lags arising from intervals between trades of the underlying stocks. At any instant, each stock price is somewhat stale because it has not been adjusted since the last trade, so the index exhibits staleness; however, each trade of the ETF represents an actual trade, which by definition is not stale at the time it occurs. In particular, each trade of the ETF occurs at a price different from the current value of the index; in the absence of partial price adjustment, the ETF price should reflect all the information in the market, in particular the “correct” price of the stocks in the index, even if many of those stocks have not traded for some time.

For this paper, the ETF we choose is SPDRs, an ETF based on the S&P 500 Index; each SPDR share represents a claim to one-tenth of the value of the S&P 500 Index. In our sample period, SPDRs exhibit weakly negative daily autocorrelation. The daily return of each individual stock on day d is computed in the conventional way: the price at the final trade on day d , minus the price at the last trade prior to day d , divided by the price at the last trade prior to day d . We compute the correlation between the return of stock i on day $d+1$ (in other words, the return from the final trade of the stock on day d to the final trade of the stock on day $d+1$) with the return of the SPDRs over the interval from the time of the last trade of the SPDRs on day $d-1$ through the time of the last trade of the *stock* on day d . If a stock does not trade on day d or the stock does not trade on day $d+1$, we omit the data from our calculation.¹² Note that each time we compute a correlation, it is the correlation of a stock return over a given interval with the return of a traded security, SPDRs, over a disjoint interval, with both the SPDRs and the stock trading at the common point of the two intervals. Thus, the calculation of the correlation does not use stale prices, and hence there is no nonsynchronous trading effect. There may be an effect due to bid-ask bounce, but if so, it should be negative. Thus, in the absence of partial price adjustment, the correlation between the return of the individual stock and the return of the SPDRs must be less than or equal to zero. Our Null Hypothesis VI (VIA, VIB) is that the correlation of each of the individual stock returns and the return of the SPDRs is zero (nonpositive, nonnegative). As in Null Hypotheses II (IIA, IIB) and IV (IVA, IVB), we divide our data period into five two-year subperiods, and test using $p_3(\mu=5, n=5)$ for Null Hypothesis VI, and $p_3(\mu=2.5, n=5)$ for Null Hypotheses VIA and VIB. Rejection of Null Hypothesis VIA implies that the partial price adjustment effect exists and is positive.

Ex ante, this seems an unlikely place to search for partial price adjustment. Partial price adjustment is usually discussed in the market microstructure literature, and is understood to mean the slow incorporation of private, firm-specific, information into the prices of individual securities. The current price of SPDRs is public, not private. Indeed, because of its link to the closely-watched S&P 500 Index, it is one of a handful of the most visible market statistics. Moreover, the information contained in the price of SPDRs describes the overall status of the market, rather than any firm-specific factors that could result in a large movement in the price of individual firms. In our test, the only detectable source of partial price adjustment is the slow incorporation into the price of individual firms of the very public, non-firm-specific, information contained in the price of SPDRs. The remainder, which presumably constitutes the vast majority of the total partial price adjustment present in the market, is not captured by these tests.

Because the definitions underlying Null Hypotheses VI, VIA, and VIB are somewhat complex, we here present a more formal statement of the model.

Assuming closing time is 4:00 p.m., we define the following notation:

$$\begin{aligned}
S_{(d,h)i} &= \text{Price of stock } i \text{ at hour } h \text{ on date } d, \\
\bar{S}_{(d,h)} &= \text{Price of SPDRs at hour } h \text{ on date } d, \\
h(d,i) &= \text{Hour of last trade of stock } i \text{ on date } d, \\
S_{di} &= S_{(d,h(d,i))i} \text{ (the closing price),} \\
\bar{S}_d &= \bar{S}_{d,4pm}, \\
r_{di} &= \frac{S_{di} - S_{(d-1)i}}{S_{(d-1)i}}, \\
\bar{r}_d &= \frac{\bar{S}_d - \bar{S}_{d-1}}{\bar{S}_{d-1}},
\end{aligned}$$

where ‘‘hour’’ means actual time of transaction; thus, it indicates transaction data down to the minute and second.

We decompose the daily return of the SPDRs, $\frac{\bar{S}_d - \bar{S}_{d-1}}{\bar{S}_{d-1}}$, into two components, $\frac{\bar{S}_d - \bar{S}_{d,h(d,i)}}{\bar{S}_{d-1}}$ and $\frac{\bar{S}_{d,h(d,i)} - \bar{S}_{d-1}}{\bar{S}_{d-1}}$. No stale prices are used in the calculation of $Corr\left(r_{(d+1)i}, \frac{\bar{S}_{d,h(d,i)} - \bar{S}_{d-1}}{\bar{S}_{d-1}}\right)$, the correlation between the return of stock i tomorrow and today’s return of SPDRs up to the time of the stock i ’s today’s last transaction. These two returns are computed on disjoint intervals separated by a

trade, as we can see in Figure 1, so there is no nonsynchronous trading effect; for details, see Appendix A. Consequently, in the absence of partial price adjustment and bid-ask bounce, this covariance must be zero. Since bid-ask bounce induces negative autocorrelation, in the absence of partial price adjustment, the correlation must be less than or equal to zero.

$$\text{Corr}\left(\eta_{(d+1)i}, \frac{\bar{S}_{d,h(d,i)} - \bar{S}_{d-1}}{\bar{S}_{d-1}}\right) \leq 0 \quad (1)$$

<Insert Figure 1>

C. Testing Hypothesis: Testing for Covariance

Testing each of our Null Hypotheses requires testing whether a correlation or a set of correlations is zero, positive, or negative. We use three test methods: the Pearson correlation test, the modified Pearson correlation test using Andrews' (1991) heteroskedasticity and autoregression consistent (HAC) covariance estimator, and Kendall's tau test.

- **Pearson correlation test (a parametric test):** This method tests whether the correlation between two variables is zero, positive, or negative using the Pearson product-moment correlation coefficient. Letting r_p be the Pearson sample correlation coefficient, the t -test statistics are

$$t = r_p \sqrt{\frac{n-2}{1-r_p^2}} \sim t(n-2)$$

This test assumes that the variables have a bivariate normal distribution.

- **Modified Pearson correlation test:** This test modifies the Pearson test, taking into account the possibility that the error terms exhibit heteroskedasticity or autocorrelation. We use Andrews' (1991) heteroskedasticity and autoregression consistent (HAC) covariance estimator to estimate the correlation coefficient and to test whether it is zero, positive, or negative. The test is based on the fact that the t -test statistic of the correlation coefficient of the two variables is numerically equal to the t -statistic on the regression coefficient of one variable with respect to the other. The HAC covariance is obtained using Andrews' quadratic spectral (QS) kernel with automatic bandwidth selection method.
- **Kendall's tau test (a nonparametric test):** This nonparametric test makes no assumptions on the joint distribution of the variables. Kendall's sample rank correlation coefficient is defined by

$$\hat{\tau} = \frac{2K}{n(n-1)}$$

where $K = \sum_{i=1}^{n-1} \sum_{j=i+1}^n Q((X_i, Y_i), (X_j, Y_j))$ and $Q((a, b), (c, d)) = \begin{cases} 1, & \text{if } (c-a)(d-b) > 0 \\ -1, & \text{if } (c-a)(d-b) < 0 \end{cases}$. Then

Kendall's tau test statistic is given by

$$T = 3\hat{\tau} \frac{\sqrt{n(n-1)}}{\sqrt{2(2n+5)}} \sim N(0,1)$$

which is asymptotically normal; the normal provides an excellent approximation provided that $n > 10$.¹³

II. Data

Our data period covers the ten calendar years 1993 through 2002. We divide this into five two-year subperiods: 1993-94, 1995-96, 1997-98, 1999-2000, and 2001-02. Within each two-year subperiod, we obtain a sample of 100 small, 100 medium, and 100 large firms. The samples are obtained using the following criteria:

- Since our analysis requires firms' market capitalization, we select the sample from the set of common stocks included in both the Trade and Quote (TAQ) database master file and the Center for Research in Security Price (CRSP) tapes for the subperiod; the trade data comes exclusively from TAQ, while we use CRSP data to establish market capitalization. We exclude closed end investment companies or trusts from our set of common stocks.
- We remove firms in which the total number of shares outstanding changes by more than ten percent during the two-year subperiod because changes in the number of shares outstanding could significantly affect the frequency of trade and in particular the time interval between the last trade of the day and the market close.¹⁴
- We remove firms for which no transaction occurs for 30 or more consecutive trading days within the two-year subperiod.
- We remove firms whose shares traded for less than \$5 at any point in the two-year subperiod. We did this in order to eliminate financially distressed firms. We verified manually that most of them did eventually become penny stocks.
- We eliminate any stock whose shares traded for more than \$1,000 at any point in the two-year subperiod; only one stock (Berkshire Hathaway) was removed by this process.
- We form three different groups of firms stratified by market capitalization. For each stock, we

calculate the market capitalization by multiplying the number of shares outstanding by the daily closing price (or the average of the bid-ask quotes) on the last day of trading preceding the subperiod. The large-firm sample consists of the 100 largest firms not eliminated by the criteria above; 97 (average for the five two-year subperiods) of the original firms were eliminated. The medium-firm sample consists of the 100 firms with market capitalizations closest to the median that were not eliminated by the criteria above; 120.2 of the original firms were eliminated. For the small-firm sample, following Bessembinder (1997), we eliminated the 50 smallest, in order to avoid including an unduly large number of financially distressed firms in the sample; we then took the smallest 100 firms not eliminated by the criteria above; 268 were eliminated at this stage.

For each of these 100 NYSE-listed firms in a given group and subperiod, we obtain transaction data from the TAQ database; we exclude trades that occurred on other exchanges. We manually cleaned the data to remove clearly erroneous prices.¹⁵ We also use transaction data of the SPDRs over the same sample data period. Since our goal is to understand the sources of daily autocorrelation in stock returns, and daily returns are calculated from the closing market prices, we exclude trade that occurs in the after-hours markets from our data set. Thus, for each individual stock, we compute the closing price in the conventional way: the last transaction price reported *before* 4:00 p.m.

Table I reports descriptive statistics of our 300 NYSE-listed sample firms stratified into each of three groups: small, medium, and large firms. The variables are the number of firms, the firms' market capitalization (in millions of dollars; max, mean, and min), average daily trading volume (in shares), average time interval between the last transaction and the market close (in seconds), and average number of days on which trade occurs. The figures reported are the averages over the five two-year subperiods, except for max and min where we report the max of the subperiod maxes and the min of the subperiod mins. The numbers of Table I reflect the diverse sample of securities used in this study. The firms' market values range from 23.9 million dollars to 475.0 billion dollars. The Table shows that as the firm size increases, both average trading volume and average trading days increase. Table I also reports the mean daily returns of each portfolio; 0.0604%, 0.0395%, and 0.0348% for small-, medium-, and large-firm portfolio respectively, reflecting a strong small-firm effect. For the first-order autocorrelation of conventional daily portfolio returns for each portfolio, see Table VI.

The closing trade for an individual stock is the last trade occurring before 4:00 p.m. To ensure that there is no overlap in the time interval before the closing trade of the individual stocks and the time interval after the closing trade of the SPDRs, we take the closing trade of the SPDRs to be the first trade occurring after 4:00 p.m., except on the seven days on which the market closed early, where we take the

closing trade of the SPDRs to be the last trade occurring before 4:00 p.m. For individual stocks as well as SPDRs, the closing price is defined to be the transaction price of the closing trade. On average, conditional on there being at least one transaction in a given stock, the last transaction of the small, medium, and large firms occurs 47.3 minutes, 11.4 minutes, and 1.2 minutes, respectively, before the closing trade of the SPDRs.¹⁶

For SPDRs, the first-order daily return autocorrelation is slightly negative (-0.0199) but statistically insignificant. Since SPDRs are traded continuously, we do not expect to find positive daily return autocorrelation arising from the nonsynchronous trading effect, as we see in portfolios; however, SPDRs are subject to bid-ask bounce.

<Insert Table I>

III. Empirical Results and Their Implications

A. Individual Stock Returns

A.1. Conventional Daily Return Autocorrelation

Table II reports results on the conventional daily return autocorrelation of individual firms, stratified into the three groups. Null Hypothesis IA, that the average conventional daily return autocorrelation is less than or equal to zero, is strongly rejected among the small firms. The average is not significant among medium and large firms. Since both the nonsynchronous trading and bid-ask bounce effects result in negative autocorrelation, this finding provides compelling evidence that partial price adjustment is the main source of daily return autocorrelation among small firms. It also provides evidence that partial price adjustment is as important as nonsynchronous trading and bid-ask bounce combined for medium and large firms.

In the model of Chan (1993), the daily return autocorrelation of each individual firm is zero, which implies our Null Hypothesis I. In Chan's Table I, he reports that the average autocorrelation for all NYSE and AMEX firms was positive and highly significant in the period 1980-84, negative and highly significant in the period 1985-89, and not significant over the entire period 1980-89. He also found that average daily return autocorrelation was negative and highly significant for small firms, not significant for medium firms, and positive and highly significant for large firms in the period 1980-89. Although he does not note this, Chan's results are not consistent with zero daily return autocorrelation in each firm; in particular, the results reported in his Table I imply rejection of all three of our Null Hypotheses I, IA and IB. While the average return autocorrelation of all stocks is insignificant over his entire data period, the

systematic effects of firm size and year on the autocorrelation provide clear evidence that there is more going on in the market than in Chan's model.

<Insert Table II>

Table III reports the results for Null Hypotheses II, IIA, and IIB. Null Hypothesis II, that each firm's daily return autocorrelation is zero, is rejected at the 1% level, for all three correlation tests, for both small and medium firms. Null Hypothesis IIA, that each firm's daily return autocorrelation is nonpositive, is rejected at the 1% level, for all three correlation test, for small firms; it is rejected at the 5% level, for two of the three correlation tests, and the 1% level for the remaining test, for medium firms. The strong rejection of Null Hypothesis IIA shows that the nonsynchronous trading and bid-ask bounce effects, which predict negative daily return autocorrelation for individual firms, cannot be the only source, or even the main source, of autocorrelation in small and medium firms.

Null Hypothesis IIB, that each firm's daily return autocorrelation is nonnegative, is rejected at the 5% level, for all three correlation tests, in medium firms. It is not rejected among large firms, and is rejected by only one of the three correlation tests among small firms. The rejection among medium firms could indicate that the partial price adjustment effect is negative for some stocks, due to overshooting. Alternatively, it could result from negative nonsynchronous trading and bid-ask bounce effects that more than offset the positive partial price adjustment effect.

These results present a clear picture. Partial price adjustment must be a significant source of positive autocorrelation of daily returns of individual stocks among small and medium firms. Nonsynchronous trading and bid-ask bounce may partially offset partial price adjustment in some stocks, and more than offset it in other stocks; or the partial price adjustment effect may be negative in some stocks, due to overshooting. On balance, average daily return autocorrelation is positive and significant among small firms, and is statistically insignificant among medium and large firms. The lack of significance among medium and large firms could indicate that the partial price adjustment is roughly as important as the nonsynchronous trading and bid-ask bounce effects for medium and large firms, or that the partial price adjustment effect is negative for some firms and positive for others, or that tests based on the average autocorrelation over stocks have less power than our tests based on individual stocks.

<Insert Table III>

A.2. Intraday Return Autocorrelation

Table IV reports results on average intraday return autocorrelation of individual firms, stratified into the three groups. As noted above, intraday returns are calculated so as to eliminate the nonsynchronous trading effect and eliminate or greatly reduce the bid-ask bounce effect. The only plausible source of the

remaining autocorrelation is partial price adjustment. Averages are reported for each of the five two-year subperiods, along with an average for the whole ten-year period obtained as a weighted average of the subperiod averages.

Null Hypothesis III, that the average intraday return autocorrelation is zero, is rejected at the 1% level for small, medium and large firms. The average autocorrelation coefficients are positive for small and medium firms, so Null Hypothesis IIIA is also rejected, indicating that the partial price adjustment is on average positive among these firms. For small firms, the average is positive and highly significant in all five subperiods, while for medium firms, the average is positive and highly significant in four of the five subperiods, negative and highly significant in the fifth subperiod.

The average autocorrelation coefficient for large firms is small and negative (-0.0064), but this is still significant at the 1% level; it is significant and negative in two of the subperiods, significant and positive in one subperiod, positive but insignificant in one subperiod, and negative but insignificant in the remaining subperiod. The negative values cannot be explained by the nonsynchronous trading effect, since there *is* no nonsynchronous trading in the calculation of intraday returns. Because the effect over the whole ten-year sample period is small, it could conceivably result from the remaining remnant of the bid-ask bounce effect remaining in intraday returns. However, bid-ask bounce should be relatively uniform over the period, and decreasing somewhat as tick size declined. Thus, bid-ask bounce cannot explain the pattern that occurs over the subperiods. The most likely explanation is that the intraday autocorrelation among large stocks comes primarily from partial price adjustment, but the sign of the partial price adjustment varies from period to period. For example, this variation could result from variation in the number of traders who use momentum strategies in trading large stocks, with overshooting (and hence negative autocorrelation) when momentum strategies predominate.

<Insert Table IV>

Table V reports our results on individual intraday stock returns. Null Hypothesis IV, that each firm's intraday return autocorrelation is zero, is rejected at the 1% level, for all three correlation tests, for small and medium firms; it is not rejected among large firms.

As noted above, Null Hypothesis IVA, that each firm's intraday return autocorrelation is nonpositive, is the most important hypothesis in this part of our study. Null Hypothesis IVA is rejected at the 1% level for all three correlation tests for small firms. It is rejected at the 1% level for two of the three correlation tests for medium firms, barely missing the 1% level in the remaining test. This shows that partial price adjustment is a significant source of daily return autocorrelation in small and medium stocks.

Null Hypothesis IVB, that each firm's intraday return autocorrelation is nonnegative, is not rejected, for any of the correlation tests, in any of the three size groups. This indicates that the partial price

adjustment effect is systematically positive among small and medium firms: the negative autocorrelation resulting from positive-feedback strategies and overshooting is systematically smaller than the positive autocorrelation resulting from the slow incorporation of information into prices.

The numbers of positive and negative rejections among large firms vary significantly among subperiods, preventing us from rejecting either Null Hypothesis IVA or IVB. This is consistent with our conjecture that variation in the number of traders using momentum strategies explains the autocorrelation pattern among individual large stock returns.

Finally, we use the methodology described above to provide a lower bound of the identifiable absolute autocovariance of individual stocks arising from partial price adjustment. Over the five two-year subperiods, our estimates range from 48.0% to 61.8% (average 56.2%) for small stocks, 50.5% to 64.6% (average 60.7%) for medium stocks, and 48.8% to 56.7% (average 52.6%) for large stocks. In all three firm-size groups, more than half of the autocovariance of individual stocks comes from partial price adjustment.

<Insert Table V>

A.3. Summary—Individual Stock Return Autocorrelation

Partial price adjustment must be an important source of the autocorrelation of daily returns of individual stocks, among small and medium firms. The partial price adjustment effect is systematically positive among small and medium, indicating that the positive autocorrelation arising from slow incorporation of information into prices outweighs the negative autocorrelation arising from positive-feedback strategies and consequent overshooting. The most likely explanation of our findings among large firms is that partial price adjustment is also an important source of autocorrelation of daily returns, but the sign of partial price adjustment varies among subperiods, possibly as a result of variation in the popularity of momentum strategies among traders of large stocks.

B. Portfolio Returns

B.1. First Method, Intraday Returns

Tables VI and VII report our results concerning conventional and intraday portfolio returns. Results are presented for each of the five two-year subperiods, along with a value for the whole ten-year period computed as a weighted average of the subperiod returns. Because conventional portfolio returns do not provide a test of partial price adjustment, we did not formalize null hypotheses on conventional portfolio returns. Table VI presents the results for conventional portfolio returns. The conventional daily return autocorrelations of small-, medium-, and large-firm portfolios are positive and significant at the 1% level for all three correlation tests. As the firm size becomes larger, the first-order autocorrelation of portfolio

return becomes smaller. This result is consistent with those of the previous studies (e.g., Chordia and Swaminathan (2000, Table I on page 917)). Table VII presents the results for intraday portfolio returns, and our tests of Null Hypothesis V. The intraday portfolio return autocorrelation is positive and significant at the 1% level, for all three correlation tests, for small and medium firms. This provides strong evidence that partial price adjustment is an important source of portfolio return autocorrelation, and that it is on balance positive, in small and medium firms. Among large firms, the conventional portfolio return autocorrelation is positive and significant, although smaller than among medium and small firms. By contrast, the intraday portfolio return is not significant and slightly negative. Thus, our intraday portfolio returns do not provide evidence of partial price adjustment in portfolios of large stocks; however, as we shall see, we do find such evidence in our test involving SPDRs.

Table VIII shows the autocovariances of conventional (intraday) returns in each of the three portfolios, obtained by multiplying the conventional (intraday) return autocorrelations by the conventional (intraday) variances. The ratio of the intraday autocovariance to the sum of the absolute values of the intraday and residual autocovariances ranges from 44.77% to 65.54% over the five subperiods, with an average of 54.58%, for small firms; from 5.03% to 90.84%, with an average of 59.54%, for medium firms, and from 18.41% to 64.02%, with an average of 36.82%, for large firms. As noted above, these figures represent lower bounds of the portion of the autocorrelation attributable to partial price adjustment.

<Insert Table VI>

<Insert Table VII>

<Insert Table VIII>

B.2. Second Method, ETFs

Table IX shows the results of our tests of Null Hypotheses VI (VIA, VIB): that the correlation of individual stock and SPDRs returns is zero (nonpositive, nonnegative). As explained above, these correlations are calculated in a way that eliminates the nonsynchronous trading effect, but not necessarily bid-ask bounce. Consequently, if the nonsynchronous trading and bid-ask bounce effects are the only sources of stock return autocorrelation, the correlation must be less than or equal to zero. Thus, Null Hypothesis VIA is the most important of this group of hypotheses. For all three correlation tests, it is rejected at the 1% level for small and medium firms, and at the 5% level for large firms. The rejection of Null Hypothesis VIA provides strong existence of partial price adjustment among small, medium and large firms. As noted above, finding strong evidence of partial price adjustment in this setting is surprising. The information contained in the SPDR price is very public, and underlying S&P 500 index is arguably the most closely-watched indicator of the state of the whole U.S. market. The information in the SPDR

price concerns the state of the whole market, not the particular prospects of individual firms. Although the evidence is stronger among small and medium firms, it is statistically significant among large firms, for which the markets should be particularly efficient. The fact that we find partial price adjustment, even in this setting, indicates that it must be very pervasive. It is important to note, also, that Null Hypothesis VIB, that the correlation of individual stock and SPDRs returns, is nonnegative, is not rejected for any of the correlation tests or size groups of firms.

When we look at the five two-year subperiods, a somewhat more nuanced story emerges. The ratio of positive to negative rejections varies substantially among subperiods, particularly among large stocks. In some subperiods, the number of negative rejections is very low, and this results in the failure to reject Null Hypothesis VIB. However, there are subperiods where the negative rejections outnumber the positive rejections, indicating again that the correlation pattern between the SPDRs and firms varies somewhat over time.

Our finding of partial price adjustment for large firms in the SPDRs tests contrasts with the failure to find such evidence for large firms in the portfolio intraday return tests. Here is a possible explanation. In the SPDRs tests, the time interval between the final trade of a stock and the last previous SPDR trade is very short, several seconds toward the end of our data period. By contrast, in the portfolio intraday return tests, the time interval between the first trade on day $d+1$ and the last trade on day d includes an entire overnight period, as well as some time when the markets are open. Chordia, Roll, and Subrahmanyam (2005) estimate the time it takes for the prices of large stocks to adjust to information (in their case, the imbalance between buy and sell orders in the order book), finding it to be more than five but less than sixty minutes. If the time it takes large stocks to adjust to information contained in the prices of other stocks, and of market indices, is similar, this could explain the finding of partial price adjustment in the SPDR tests, along with the failure to find it in the intraday portfolio return tests for large stocks. If so, the finding of partial price adjustment in the intraday portfolio return tests for small and medium stocks indicates that the prices of small and medium stocks adjust to information much more slowly than to large stocks.

<Insert Table IX>

IV. Concluding Remarks

We investigate whether stock return autocorrelation is spurious. We find compelling evidence that it is not. In particular, we find strong evidence that partial price adjustment is an important source, and

in some cases the main source, of stock return autocorrelation.

Previous attempts to measure the roles of the nonsynchronous trading and bid-ask bounce effects and partial price adjustment in generating individual stock and portfolio return autocorrelation have been indirect. We formulate several direct tests for the presence of partial price adjustment. These direct tests eliminate the nonsynchronous trading effect by considering correlations over disjoint time intervals, separated by a trade. Once the nonsynchronous trading effect is eliminated, the only plausible sources of autocorrelation are bid-ask bounce and partial price adjustment. Since bid-ask bounce results in negative autocorrelation, the absence of partial price adjustment implies that the correlation must be less than or equal to zero for every stock or pair of stocks. The repeated strong rejections, in a variety of situations, of the hypothesis that these correlations are all less than or equal to zero provides compelling evidence for the presence of partial price adjustment.

Previous studies of individual stock return autocorrelation computed the average autocorrelation over all stocks, producing inconclusive results. We explain why the relationship between daily return autocorrelation and firm size reported in Chan (1993) provides strong evidence of partial price adjustment; we also test average autocorrelation by firm size, and find strong evidence of partial price adjustment. We test the autocorrelation of each individual stock, and find that the hypothesis that the autocorrelations are all less than or equal to zero is strongly rejected among small and medium firms, and conclude that partial price adjustment must be an important source of the autocorrelation.

Previous tests for partial price adjustment in portfolio return autocorrelation have been indirect. By computing intraday portfolio returns, we are able to eliminate the nonsynchronous trading effect and eliminate or greatly reduce bid-ask bounce, allowing us to measure partial price adjustment. We find strong evidence of partial price adjustment in portfolios of small and medium stocks. We even find evidence of partial price adjustment in an unlikely place. The information contained in the price of SPDRs is very public, and is not firm-specific, so we expect the autocorrelation of SPDRs and individual stocks capture only a small portion of the total partial price adjustment. The fact that we find evidence for partial price adjustment, even in this setting, and even for large firms, indicates how pervasive partial price adjustment must be.

Some of our tests allow us to establish lower bounds on the proportion of the autocorrelation that comes from partial price adjustment. In each case, we find that the proportion is very substantial.

We use two methods to eliminate the nonsynchronous trading effect. The first method computes correlations of intraday returns; this method can be applied to eliminate the nonsynchronous trading effect with other types of securities, and on other exchanges. The second method, used in computing the correlation of individual stock returns and SPDRs, computes the return of the SPDRs separately in the

periods before and after the final trade of the stock; this method can be used to eliminate the nonsynchronous trading effect for any security which, like the SPDRs, is traded nearly continuously.

By dividing our data period into disjoint subperiods, we are able to work around the problem of correlation of returns across stocks. Our test, based on the order statistics of the results from the subperiods, can be applied to other types of securities and other exchanges.

Further research is needed on the following questions:

- to what extent do these findings extend to other markets involving different institutional structures?
- among large firms, we find strong evidence of partial price adjustment among portfolios in our test using SPDRs, but not our tests involving intraday returns. The use of intraday returns allows us to measure only a portion of the partial price adjustment; that portion is large enough to generate statistical significance among small and medium firms, but evidently not among large firms. The test involving SPDRs captures a different portion of the partial price adjustment. Is there some other way to capture more of the partial price adjustment in a single test?
- our tests seem to indicate that partial price adjustment among large firms is positive in certain periods and negative in other periods. This might be explained by variations in the number of traders playing momentum strategies in these stocks. Is there a way to test this?
- we find strong evidence of partial price adjustment among individual small and medium firm returns, but not among individual large firm returns. The number of positive autocorrelations of individual large-firm intraday returns in the various subperiods is usually substantially above the expected value, suggesting that there is partial price adjustment in this setting, but our test involving the first two order statistics is too weak to detect it, at least using five two-year subperiods or ten one-year subperiods. Would a longer overall data period, and/or a different statistical test, allow one to establish partial price adjustment in this setting?

Appendix A: Derivation of Equation (1)

As we can see in Figure 1, the daily return of SPDRs at day d , $\frac{\bar{S}_d - \bar{S}_{d-1}}{\bar{S}_{d-1}}$, consists of two components, $\frac{\bar{S}_{d,h(d,i)} - \bar{S}_{d-1}}{\bar{S}_{d-1}}$ (the return of the SPDRs from 4:00 p.m. yesterday (day $d-1$), to time $h(d, i)$ today (day d)) and $\frac{\bar{S}_d - \bar{S}_{d,h(d,i)}}{\bar{S}_{d-1}}$ (the return from time $h(d, i)$ today to 4:00 p.m. today). Here $h(d, i)$ is the time of the individual stock i 's last transaction on day d . We have the following identity:

$$\frac{\bar{S}_d - \bar{S}_{d-1}}{\bar{S}_{d-1}} = \frac{\bar{S}_d - \bar{S}_{d,h(d,i)}}{\bar{S}_{d-1}} + \frac{\bar{S}_{d,h(d,i)} - \bar{S}_{d-1}}{\bar{S}_{d-1}}$$

For individual stock i , the last transaction occurs at time $h(d, i)$ of day d and $h(d+1, i)$ of day $d+1$. The usual story for correlation arising from the nonsynchronous trading effect goes as follows: Suppose that information affecting the value of the stock i becomes known between $h(d, i)$ and 4:00 p.m. of day d (interval B in Figure 1). This information will not be reflected in stock i 's closing price on day d , but will be reflected in price on day $d+1$, and thus in the return, $r(d+1, i)$, on day $d+1$. However, the SPDRs trade very frequently, and will usually trade at many times between $h(d, i)$ and 4:00 p.m. of day d . Consequently, the information will be reflected in the SPDRs' price and return on day d . This induces a spurious positive correlation between the SPDRs' return on day d and the stock return on day $d+1$.

Our analysis, however, is not dependent on the particular mechanism by which the nonsynchronous trading effect induces spurious correlation. The contribution of the nonsynchronous trading effect to the correlation between the SPDRs return and the stock return comes solely from the interval B in Figure 1, where there is an overlap between the time intervals on which the day d return of the SPDRs and the day $d+1$ return of stock i are computed. Said slightly differently, the return of the stock on day $d+1$ is computed using the price of the stock at the time of its last trade on day d , and that price is stale on the interval B in Figure 1, but is *not* stale at the time $h(d,i)$.

$r_{(d+1),i}$ is the return over the intervals B and C. $\frac{\bar{S}_d - \bar{S}_{d-1}}{\bar{S}_{d-1}}$ is the return over the intervals A and B.

The correlation comes only from the overlap, interval B. If we eliminate interval B from our return calculation for the SPDRs, the return becomes $\frac{\bar{S}_{d,h(d,i)} - \bar{S}_{d-1}}{\bar{S}_{d-1}}$. In the correlation

$Corr\left(r_{(d+1)i}, \frac{\bar{S}_{d,h(d,i)} - \bar{S}_{d-1}}{\bar{S}_{d-1}}\right)$, no stale prices are used; if the correlation is not zero, it must be coming from something other than the nonsynchronous trading effect. Since bid-ask bounce induces negative correlation, in the absence of partial price adjustment, the correlation must be less than or equal to zero, so Equation (1) holds.

Appendix B: Time-Varying Risk Premia

In this Appendix, we analyze the effect of time-varying risk premia on stock return autocorrelation. This allows us to estimate the magnitude of the effect on our autocorrelation estimates, and thereby validate our estimates of partial price adjustment.

Under the assumption that stock prices follow one of the standard processes in finance (such as a geometric Itô or geometric Lévy Process), rejection of the hypothesis that stock return autocorrelation is zero is equivalent to rejection of the hypothesis that the expected rate of return is constant. In other words, if we impose the assumption that the return in each period is composed of an expected return plus a volatility term, where the volatility term is uncorrelated with the returns in disjoint periods, then returns are uncorrelated if and only if the expected return is constant. As noted by Campbell, Lo, and McKinlay (1997, page 66), the “ R^2 of a regression of returns on a constant and its first lag is the square of the slope coefficient, which is simply the first-order autocorrelation.” As a consequence, if the first-order autocorrelation coefficient of return is α , the proportion of the variation in return that “is predictable using the preceding day’s ... return” is α^2 . Thus, time-varying expected rates of return and return autocorrelation are simply different faces of a single phenomenon.

However, time-varying expected rates of return and time-varying risk premia are distinct phenomena. To illustrate, suppose that stock prices follow Itô processes of the form $\frac{dS}{S} = \mu dt + \sigma dW$; a similar analysis holds if prices follow other standard processes. The absence of arbitrage is equivalent to the existence of a vector process λ of prices of risk such that $\mu - r = \sigma \lambda'$; here, μ is the vector process of expected returns rates and r is the risk-free rate. The expected rates of return μ will vary as a result of changes in r , σ , and λ , and the resulting variation in μ cannot be exploited by arbitrage; this is the variation attributable to time-varying risk premia. Other variation in μ constitutes time-varying expected rates of return, not time-varying risk premia; it can be exploited by arbitrage.

Standard autocorrelation tests are designed to test for time-varying expected returns, but cannot distinguish time-varying risk premia from other forms of time-varying expected returns. In particular, if a stock has a run of positive returns, autocorrelation tests will conclude that the stock had a high expected rate of return over that period, but cannot distinguish whether or not this is the result of a high risk premium. If it is, then it cannot be exploited by arbitrage by informed traders.

If the high expected rate of return is not the result of a high risk premium, then it *can* be exploited by arbitrage. If no one knew the expected rate of return was high, there would be nothing pushing the stock higher, and it would stay relatively stable until the good news underlying the high expected rate of return were announced, at which point the stock price would rise abruptly. If it were widely known that the expected rate of return was high, then many traders would buy the stock, forcing the price to rise abruptly until the future expected returns were reduced to the appropriate risk-adjusted level. These abrupt rises in price would be captured econometrically as volatility, and not as autocorrelation. Thus, if we see autocorrelation in stock returns after eliminating the nonsynchronous trading and bid-ask bounce effects, it can only come from two sources: time-varying risk premia, or the strategic decision of a small group of informed traders to exercise their informational advantage slowly. In short, it must either be time-varying risk premia or partial price adjustment.

To the best of our knowledge, no paper has asserted that time-varying risk premia are a significant source of autocorrelation in daily returns of individual stocks or portfolios over periods of length two years.¹⁷ Nonetheless, time-varying risk premia will induce some autocorrelation, and hence induce some bias in our measures of the partial price adjustment effect. However, we can put an upper bound on the potential impact of time-varying risk premia on the autocorrelation. This allows us to validate our claim that the autocorrelation we find could come only from partial price adjustment.

The analysis we give can be applied to different periods of return (daily, weekly, monthly, quarterly, annual) and different time horizons (one month, six months, one year, two years, five years, ten years, several decades). The analysis shows that the bias in the measured autocorrelation resulting from time-varying risk premia depends on both the return period and the time horizon. The bias becomes larger as the time horizon increases (because the variation of return resulting from time-varying risk premia is larger over longer time horizons) and larger as the period of return increases (daily returns are much noisier than yearly returns, so the bias represents a larger fraction of the total return variance). We find that the bias in daily returns over a two-year time horizon is very small; however, the bias in annual returns over a horizon of decades could be substantial.

All of our correlations are calculated over two-year subperiods of the period 1993-2002. How much might the expected daily risk-adjusted return of a well-diversified stock portfolio vary over one of these two-year subperiods? Stocks and portfolios could conceivably have expected rates of return below the risk-free rate, but only if they were negatively correlated with undiversifiable risks and thus provided insurance against those risks. Since broadly diversified stock portfolios are

positively correlated with two important undiversifiable risks (the market as a whole and aggregate income), it is implausible that investors would hold the portfolio if it had an expected rate of return below the risk-free rate. To maintain equilibrium, stock prices would have to fall to raise the future expected return sufficiently to induce stockholders to retain their holdings. On the other hand, if the expected return of the portfolio exceeded the risk-free rate by 15% per annum, investors would surely choose to substantially increase their stockholdings: taking the volatility into account, there is very little chance of a substantial decline in stock prices and a very good chance of a substantial gain. Over the two-year periods we consider, the average variation of the risk-free rate (as measured by the three-month Treasury Bill Rate) is 2.49%.¹⁸ Thus, we assume that time-varying risk premia will induce variation in the return on a well-diversified portfolio of no more than 18% per annum over a two-year period.

In the case of an individual stock, the expected return should reflect the risk premia of the factors underlying its pricing. Some stocks may have low—even negative—risk premia, while others may have large risk premia. However, the correlation of any given stock with the main risk factors should be relatively stable over time periods of two years.¹⁹ Thus, we assume that time-varying risk premia induce a variation of no more than 18% per annum in the expected return of each of our securities, and each of the portfolios we consider, in any of our two-year periods. This assumption does not restrict the variation in expected rates of return across securities; our assumption limits only the variation across time for a given security.

We now turn to the effect of time-varying risk premia on our correlation estimates. For the sake of simplicity, we assume that securities prices follow geometric Itô processes; we believe similar estimates would hold for geometric Lévy processes. Suppose the security (or portfolio) price S follows the stochastic differential equation

$$\frac{dS}{S} = \mu dt + \sigma dW$$

where W is a standard Wiener process, and μ and σ are continuous deterministic functions of time. Assume there are 250 trading days per year. Let

$$\begin{aligned} \sigma_k &= \sigma \left(\frac{k}{250} \right) \\ \mu_k &= \mu \left(\frac{k}{250} \right) \\ \bar{\mu} &= \frac{1}{498} \left(\frac{\mu_1}{2} + \sum_{k=2}^{498} \mu_k + \frac{\mu_{499}}{2} \right) \\ &\simeq \frac{1}{2} \int_0^2 \mu(t) dt \end{aligned}$$

$$\begin{aligned}
\Delta W_k &= W\left(\frac{k+1}{250}\right) - W\left(\frac{k}{250}\right) \\
\bar{v} &= \frac{1}{498} \left(\frac{\sigma_1 \Delta W_1}{2} + \sum_{k=2}^{498} \sigma_k \Delta W_k + \frac{\sigma_{499} \Delta W_{499}}{2} \right) \\
&\simeq \frac{1}{498} \int_0^2 \sigma dW \\
r_k &= \frac{\mu_k}{250} + \sigma_k \Delta W_k \\
\bar{r} &= \frac{1}{498} \left(\frac{r_1}{2} + \sum_{k=2}^{498} r_k + \frac{r_{499}}{2} \right) \\
&= \frac{\bar{\mu}}{250} + \bar{v} \\
\sigma_r^2 &= \frac{1}{498} \left(\frac{(r_1 - \bar{r})^2}{2} + \sum_{k=2}^{498} (r_k - \bar{r})^2 + \frac{(r_{499} - \bar{r})^2}{2} \right) \\
&= \frac{1}{498} \left(\frac{1}{2} \left(\frac{\mu_1 - \bar{\mu}}{250} + \sigma_1 \Delta W_1 - \bar{v} \right)^2 + \sum_{k=2}^{498} \left(\frac{\mu_k - \bar{\mu}}{250} + \sigma_k \Delta W_k - \bar{v} \right)^2 \right. \\
&\quad \left. + \frac{1}{2} \left(\frac{\mu_{499} - \bar{\mu}}{250} + \sigma_{499} \Delta W_{499} - \bar{v} \right)^2 \right) \\
&\simeq \frac{1}{2} \int_0^2 \left(\frac{\mu(t) - \bar{\mu}}{250} \right)^2 dt + \frac{2}{498} \int_0^2 \left(\frac{\mu(t) - \bar{\mu}}{250} \right) \sigma(t) dW - \bar{v} \int_0^2 \frac{\mu(t) - \bar{\mu}}{250} dt \\
&\quad + \frac{1}{498} \left(\frac{\sigma_1^2 (\Delta W_1)^2}{2} + \sum_{k=2}^{498} \sigma_k^2 (\Delta W_k)^2 + \frac{\sigma_{499}^2 (\Delta W_{499})^2}{2} \right) - \bar{v}^2 \\
&\simeq \frac{1}{2} \int_0^2 \left(\frac{\mu(t) - \bar{\mu}}{250} \right)^2 dt + \frac{2}{498} \int_0^2 \left(\frac{\mu(t) - \bar{\mu}}{250} \right) \sigma(t) dW \\
&\quad + \frac{1}{498} \int_0^2 \sigma^2(t) dt - \bar{v}^2
\end{aligned}$$

Then the Pearson sample autocorrelation coefficient is given by

$$\begin{aligned}
r_p &= \frac{\sum_{k=1}^{498} \left(r_k - \frac{1}{498} \sum_{j=1}^{498} r_j \right) \left(r_{k+1} - \frac{1}{498} \sum_{j=2}^{499} r_j \right)}{\sqrt{\sum_{k=1}^{498} \left(r_k - \frac{1}{498} \sum_{j=1}^{498} r_j \right)^2} \sqrt{\sum_{k=2}^{499} \left(r_k - \frac{2}{499} \sum_{j=2}^{499} r_j \right)^2}} \\
&\simeq \frac{\sum_{k=1}^{498} (r_k - \bar{r})(r_{k+1} - \bar{r})}{\sqrt{\sum_{k=1}^{498} (r_k - \bar{r})^2} \sqrt{\sum_{k=2}^{499} (r_k - \bar{r})^2}} \\
&\simeq \frac{\sum_{k=1}^{498} (r_k - \bar{r})(r_{k+1} - \bar{r})}{498 \sigma_r^2} \\
&= \frac{\sum_{k=1}^{498} \left(\frac{\mu_k - \bar{\mu}}{250} + \sigma_k \Delta W_k - \bar{v} \right) \left(\frac{\mu_{k+1} - \bar{\mu}}{250} + \sigma_{k+1} \Delta W_{k+1} - \bar{v} \right)}{498 \sigma_r^2} \\
&\simeq \frac{1}{\sigma_r^2} \left(\frac{1}{2} \int_0^2 \left(\frac{\mu(t) - \bar{\mu}}{250} \right)^2 dt + \frac{2}{498} \int_0^2 \left(\frac{\mu(t) - \bar{\mu}}{250} \right) \sigma(t) dW \right)
\end{aligned}$$

$$\begin{aligned}
& -\bar{v} \int_0^2 \left(\frac{\mu(t) - \bar{\mu}}{250} \right) dt + \frac{1}{498} \sum_{k=1}^{498} \sigma_k \sigma_{k+1} \Delta W_k \Delta W_{k+1} - \bar{v}^2 \Big) \\
= & \frac{1}{\sigma_r^2} \left(\frac{1}{2} \int_0^2 \left(\frac{\mu(t) - \bar{\mu}}{250} \right)^2 dt + \frac{2}{498} \int_0^2 \left(\frac{\mu(t) - \bar{\mu}}{250} \right) \sigma(t) dW \right. \\
& \left. + \frac{1}{498} \sum_{k=1}^{498} \sigma_k \sigma_{k+1} \Delta W_k \Delta W_{k+1} - \bar{v}^2 \right) \\
\approx & \frac{A + Z}{B + Z}
\end{aligned}$$

where

$$\begin{aligned}
A &= \frac{1}{498} \sum_{k=1}^{498} \sigma_k \sigma_{k+1} \Delta W_k \Delta W_{k+1} - \bar{v}^2 \\
B &= \frac{1}{498} \int_0^2 \sigma(t)^2 dt - \bar{v}^2 \\
Z &= Z_1 + Z_2 \\
Z_1 &= \frac{1}{2} \int_0^2 \left(\frac{\mu(t) - \bar{\mu}}{250} \right)^2 dt \\
Z_2 &= \frac{2}{498} \int_0^2 \left(\frac{\mu(t) - \bar{\mu}}{250} \right) \sigma(t) dW
\end{aligned}$$

In the absence of time-varying risk premia, we would have $\mu(t) = \bar{\mu}$ for all t , hence Z_1 and Z_2 would be identically zero. In the presence of time-varying risk premia, Z_1 and Z_2 induce a bias into our measurement of partial price adjustment. Notice that neither Z_1 nor Z_2 depends on the rate at which μ changes, only on the distribution of μ and (in the case of Z_2) the correlation between μ and σ . Z_1 is a positive constant. Assuming that μ is distributed uniformly over an interval of length $18\% = 0.18$ per annum, $Z_1 = 4.320 \times 10^{-8}$. Z_2 is normally distributed with mean zero and standard deviation

$$\sigma_Z = \frac{2}{498} \sqrt{\int_0^2 \left(\frac{\mu(t) - \bar{\mu}}{250} \right)^2 \sigma^2(t) dt}$$

Assuming that μ is uniformly distributed over an interval of length .18 per annum and σ is constant, $\sigma_Z = 1.180 \times 10^{-6} \sigma$. Moreover, the conditional distribution of Z_2 , conditional on A and B , is asymptotically normal.

Our tests of portfolio autocorrelation (putting aside the tests involving SPDRs) compute the daily conventional and intraday return autocorrelations of portfolios. The rejection of the null hypotheses (conventional return autocorrelation is less than or equal to zero in small, medium and large firm portfolios, intraday return autocorrelation is less than or equal to zero in small and medium firm portfolios) is overwhelming, and it is easy to see that the small bias induced by time-varying risk premia—the bias induced by Z —cannot make any difference in those results. However,

the effect of the bias on the tests involving individual stock autocorrelations and the tests involving SPDRs requires more careful analysis. Since those tests are based on comparing the actual number of rejections to the expected number of rejections, we need to carefully estimate the effect of the bias on the expected number of rejections.

Notice that A and B are quadratic in σ (i.e. if we double the function $\sigma(t)$ at all times, then A and B are quadrupled), while Z_1 is independent of σ and Z_2 is linear in σ . Thus, the bias induced in r_p by Z is maximized when σ is minimized. The returns on individual stocks in a portfolio are more volatile than the returns of the portfolio, and the returns of smaller stocks are more volatile than the returns of larger stocks. We take the volatility of the S&P 500 Index as a lower bound on the volatility of the individual stocks in our analysis. For the S&P 500 Index, the average value of σ_r , over our five two-year subperiods, is .01041. Assuming σ is constant, we obtain the estimate $\sigma = \sqrt{249} \ln(1.01041) \simeq .16341$. With probability $2 \times (1 - N(4)) > 1 - 2 \times 10^{-4}$, $|B - (\sigma_r)^2| = |Z| \leq 4.320 \times 10^{-8} + 4 \times 1.18 \times 10^{-6} \sigma = 8.145 \times 10^{-7}$, so $B = (1.084 \pm .008) \times 10^{-4}$; since the bias is maximized when B is minimized, we assume $B = 1.076 \times 10^{-4}$, $\sigma = .1600$, and $\sigma_Z = 1.888 \times 10^{-7}$.

The Pearson test compares $\frac{\sqrt{498} r_p}{\sqrt{1-r_p^2}}$ to the standard normal. We have

$$\begin{aligned} \sqrt{498} \frac{r_p}{\sqrt{1-r_p^2}} &= \sqrt{498} \frac{\frac{A+Z}{B+Z}}{\sqrt{1 - \frac{(A+Z)^2}{(B+Z)^2}}} \\ &= \sqrt{498} \frac{A+Z}{\sqrt{(B+Z)^2 - (A+Z)^2}} \end{aligned}$$

Let

$$g_{AB}(Z) = \sqrt{498} \frac{A+Z}{\sqrt{(B+Z)^2 - (A+Z)^2}}$$

When $|A| \leq B$ and $|Z| \leq B$, $\frac{A+Z}{B+Z}$ is concave in Z . It follows that on the relevant range of values ($g_{AB}(Z) \simeq 1.96$, so $\frac{A+Z}{B+Z} \simeq .09$), g_{AB} is concave, so letting $h_{AB}(Z)$ be the first-order Taylor series of g_{AB} , we have

$$g_{AB}(Z) \leq h_{AB}(Z) = \sqrt{\frac{498}{B^2 - A^2}} \left(A + \frac{B}{B+A} Z \right)$$

on the relevant range of values, and $g_{AB}(0) = h_{AB}(0)$.

The correct test would measure only the autocorrelation coming from partial price adjustment. Our actual test measures the autocorrelation coming from partial price adjustment and time-varying risk premia. Our most important null hypothesis is that the autocorrelation is less than or equal

to zero. Since $Z_1 \geq 0$, and $|A| \leq B$, the presence of Z_1 leads our test to reject the null hypothesis in situations in which the correct test would fail to reject; in any situation in which the correct test rejects, our test will also reject. Thus, the presence of Z_1 increases the expected number of rejections.

$|Z_2|$ is typically larger than Z_1 . However, we shall see that Z_2 induces a smaller bias in the expected number of rejections because Z_2 can be either positive or negative. The presence of Z_2 leads our test to reject in some cases in which the correct test does not reject, and to fail to reject in some cases in which the correct test does reject. The symmetry of Z_2 will imply that the two effects very nearly cancel.

Since $g_{AB}(Z) \leq h_{AB}(Z)$ on the relevant values, the probability of rejection in our test, which compares $g_{AB}(Z)$ to the standard normal, is lower than the probability of rejection in a hypothetical test comparing $h_{AB}(Z)$ to the standard normal. Noting that A and B are random variables, let $Y = \frac{\sqrt{498}A}{\sqrt{B^2-A^2}}$ be the random variable $h_{AB}(0)$. Let N denote the cumulative distribution function of the standard normal. The probability that Z changes an insignificant value to a significant value, using the critical value α , is the probability that $Y = h_{AB}(0) < \alpha$ and $h_{AB}(Z) \geq \alpha$, which equals

$$\begin{aligned} & \frac{1}{\sqrt{2\pi}\sigma_Z} \int_0^\infty e^{-z^2/2\sigma_Z^2} \left(N(\alpha) - N\left(\alpha - \frac{\sqrt{498}B}{\sqrt{B^2-A^2}(B+A)}(Z_1+z)\right) \right) dz \\ & \leq 1 - N(4) + \frac{1}{\sqrt{2\pi}\sigma_Z} \int_0^{4\sigma_Z} e^{-z^2/2\sigma_Z^2} (N(\alpha) - N(\alpha - \gamma_1(Z_1+z))) dz \end{aligned}$$

where γ_1 is the value of $\frac{\sqrt{498}B}{\sqrt{B^2-A^2}(B+A)}$ corresponding to $Y = \alpha - \gamma_1(Z_1 + 4\sigma_Z)$. Although it is hard to solve for γ_1 exactly, we can estimate it as follows:

$$\begin{aligned} & \frac{\partial}{\partial A} \left(\frac{\sqrt{498}B}{\sqrt{B^2-A^2}(B+A)} \right) \\ & = \sqrt{498}B \left(\frac{\partial}{\partial A} \left((B^2-A^2)^{-1/2} (B+A)^{-1} \right) \right) \\ & = \sqrt{498}B \left(A(B^2-A^2)^{-3/2} (B+A)^{-1} - (B^2-A^2)^{-1/2} (B+A)^{-2} \right) \\ & > -\sqrt{498}B (B^2-A^2)^{-1/2} (B+A)^{-2} \\ & \geq -\frac{.92\sqrt{498}}{B^2} \end{aligned}$$

on the relevant range (which is included in $1 \leq Y \leq \alpha = 1.96$, so $.044B \leq A \leq .09B$). It follows that

$$\gamma_1 \leq \gamma_0 + \frac{.92\sqrt{498}}{B^2} (Z_1 + 4\sigma_Z)$$

where γ_0 is the value of $\frac{\sqrt{498}B}{\sqrt{B^2-A^2}(B+A)}$ corresponding to $Y = \alpha$; for $\alpha = 1.96$, we have $\gamma_0 \simeq \frac{.9231\sqrt{498}}{B}$.

Similarly, the probability that Z changes an insignificant value to a significant value is bounded below by

$$\frac{1}{\sqrt{2\pi}\sigma_Z} \int_0^{4\sigma_Z} e^{-z^2/2\sigma_Z^2} (N(\alpha + \gamma_2(z - Z_1)) - N(\alpha)) dz$$

where

$$\gamma_2 \geq \gamma_0 - \frac{1.01\sqrt{498}}{B^2}(4\sigma_Z - Z_1)$$

Thus,

$$\begin{aligned} 0 \leq \gamma_1 - \gamma_2 &\leq \frac{8\sqrt{498}\sigma_Z}{B^2} \\ 0 \leq \gamma_1 + \gamma_2 &\leq 2\gamma_0 \\ 0 \leq \gamma_1^2 - \gamma_2^2 &= (\gamma_1 - \gamma_2)(\gamma_1 + \gamma_2) \\ &\leq \frac{16\sqrt{498}\sigma_Z}{B^2}\gamma_0 \\ 0 \leq \gamma_1^2 + \gamma_2^2 &\leq (\gamma_1 + \gamma_2)^2 \\ &\leq 4\gamma_0^2 \end{aligned}$$

Using the second order Taylor Expansion for the normal cumulative distribution function, the increase in the probability of rejection resulting from Z is bounded above by

$$\begin{aligned} &\frac{1}{\sqrt{2\pi}\sigma_Z} \int_0^{4\sigma_Z} e^{-z^2/2\sigma_Z^2} (2N(\alpha) - N(\alpha - \gamma_1(Z_1 + z)) - N(\alpha + \gamma_2(z - Z_1))) dz \\ &\quad + 1 - N(4) \\ &\leq \frac{1}{\sqrt{2\pi}\sigma_Z} \int_0^\infty e^{-z^2/2\sigma_Z^2} (N'(\alpha)(\gamma_1 + \gamma_2)Z_1 + (\gamma_1 - \gamma_2)z) dz \\ &\quad - \frac{1}{\sqrt{2\pi}\sigma_Z} \int_0^\infty e^{-z^2/2\sigma_Z^2} \frac{N''(\alpha)}{2} ((\gamma_1(Z_1 + z))^2 + (\gamma_2(z - Z_1))^2) dz \\ &\quad + \frac{1}{\sqrt{2\pi}\sigma_Z} \int_0^{4\sigma_Z} e^{-z^2/2\sigma_Z^2} \frac{N'''(\xi(z))}{6} (\gamma_1(Z_1 + z))^3 dz + 10^{-4} \\ &\quad \text{for some measurable function } \xi : [0, \infty) \rightarrow \mathbf{R} \\ &\leq N'(\alpha) \left(\frac{\gamma_1 + \gamma_2}{2} Z_1 + \frac{(\gamma_1 - \gamma_2)\sigma_Z}{\sqrt{2\pi}} \right) \\ &\quad - \frac{N''(\alpha)}{2} \left(\frac{\gamma_1^2 + \gamma_2^2}{2} Z_1^2 + \frac{2\sigma_Z}{\sqrt{2\pi}} (\gamma_1^2 - \gamma_2^2) Z_1 + \frac{\gamma_1^2 + \gamma_2^2}{2} \sigma_Z^2 \right) \\ &\quad + \frac{N'''(\xi(z))\gamma_1^3}{6} \left(\frac{Z_1^3}{2} + \frac{3Z_1^2\sigma_Z}{\sqrt{2\pi}} + \frac{3Z_1\sigma_Z^2}{2} + \frac{2\sigma_Z^3}{\sqrt{2\pi}} \right) \\ &\quad + 10^{-4} \\ &\leq \frac{e^{-\alpha^2/2}}{\sqrt{2\pi}} \left(\gamma_0 Z_1 + \frac{8\sqrt{498}\sigma_Z^2}{\sqrt{2\pi}B^2} \right) \end{aligned}$$

$$\begin{aligned}
& + \frac{\alpha e^{-\alpha^2/2}}{2\sqrt{2\pi}} \left(2\gamma_0^2 Z_1^2 + \frac{32\sqrt{498}\sigma_Z^2}{\sqrt{2\pi}B^2} \gamma_0 Z_1 + 2\gamma_0^2 \sigma_Z^2 \right) \\
& + \frac{e^{-3/2}\gamma_1^3}{3\sqrt{2\pi}} \left(\frac{Z_1^3}{2} + \frac{3Z_1^2\sigma_Z}{\sqrt{2\pi}} + \frac{3Z_1\sigma_Z^2}{2} + \frac{2\sigma_Z^3}{\sqrt{2\pi}} \right) + 10^{-4}
\end{aligned}$$

For the reasons explained above, with probability $1 - 2 \times 10^{-4}$, we may take $\alpha = 1.96$, $\sigma = 0.1600$, $B = 1.076 \times 10^{-4}$, $\sigma_z = 1.888 \times 10^{-7}$, $Z_1 = 4.320 \times 10^{-8}$, $\gamma_0 = \frac{.9231\sqrt{498}}{B} = 1.914 \times 10^5$, $\gamma_1 \leq \gamma_0 + \frac{.92\sqrt{498}}{B^2} (Z_1 + 4\sigma_z) = 1.929 \times 10^5$. Then the increase in the probability of rejection resulting from Z is at most

$$\begin{aligned}
& \frac{e^{-\alpha^2/2}}{\sqrt{2\pi}} \left(\gamma_0 Z_1 + \frac{8\sqrt{498}\sigma_Z^2}{\sqrt{2\pi}B^2} \right) \\
& + \frac{\alpha e^{-\alpha^2/2}}{2\sqrt{2\pi}} \left(2\gamma_0^2 Z_1^2 + \frac{32\sqrt{498}\sigma_Z^2}{\sqrt{2\pi}B^2} \gamma_0 Z_1 + 2\gamma_0^2 \sigma_Z^2 \right) \\
& + \frac{e^{-3/2}\gamma_1^3}{3\sqrt{2\pi}} \left(\frac{Z_1^3}{2} + \frac{3Z_1^2\sigma_Z}{\sqrt{2\pi}} + \frac{3Z_1\sigma_Z^2}{2} + \frac{2\sigma_Z^3}{\sqrt{2\pi}} \right) + 3 \times 10^{-4} \\
& = .0584 \left(8.271 \times 10^{-3} + 2.193 \times 10^{-4} \right) \\
& + .0573 \left(1.368 \times 10^{-4} + 7.254 \times 10^{-6} + 2.613 \times 10^{-3} \right) \\
& + .0297 \left(7.174 \times 10^{15} \right) \left(4.031 \times 10^{-23} + 4.217 \times 10^{-22} \right. \\
& \quad \left. + 2.310 \times 10^{-21} + 5.370 \times 10^{-21} \right) \\
& + 3 \times 10^{-4} \\
& = 4.958 \times 10^{-4} + 1.580 \times 10^{-4} + 1.735 \times 10^{-5} + 3 \times 10^{-4} \\
& = 9.555 \times 10^{-4}
\end{aligned}$$

Thus, the bias induced by time-varying risk premia increases the probability of rejection from 0.0250 by at most .001 to 0.026, so the expected number of rejections in 100 autocorrelations increases by at most 0.1 from 2.5 to 2.6. Our tests of individual stock autocorrelations, as well as our tests using SPDRs, are all based on comparing the number of rejections to 2.5. Changing 2.5 to 2.6 to adjust for the bias increases the p -values slightly but makes no qualitative change in those findings. Since the estimates we have just given very likely substantially overstate the bias, we feel confident in the results reported in the tables.

REFERENCES

- Ahn, Dong-Hyun, Jacob Boudoukh, Matthew P. Richardson, and Robert F. Whitelaw, 2002, Partial adjustment or stale prices?: Implications from stock index and futures return autocorrelations, *Review of Financial Studies* 15, 655-689.
- Andrews, Donald W. K., 1991, Heteroskedasticity and autoregression consistent covariance estimation, *Econometrica* 59, 817-858.
- Atchison, Michael D, Kirt C. Butler, and Richard R. Simonds, 1987, Nonsynchronous security trading and market index autocorrelation, *Journal of Finance* 42, 533-553.
- Badrinath, Swaminathan G., Jayant R. Kale, and Thomas H. Noe, 1995, Of shepherds, sheep, and the cross-autocorrelations in equity returns, *Review of Financial Studies* 8, 410-430.
- Bessembinder, Hendrik, 1997, The degree of price resolution and equity trading costs, *Journal of Financial Economics* 45, 9-34.
- Boudoukh, Jacob, Matthew P. Richardson, and Robert F. Whitelaw, 1994, A tale of three schools: Insights on autocorrelations of short-horizon stock returns, *Review of Financial Studies* 7, 539-573.
- Brennan, Michael J., Narasimhan Jegadeesh, and Bhaskaran Swaminathan, 1993, Investment analysis and the adjustment of stock prices to common information, *Review of Financial Studies* 6, 799-824.
- Campbell, John Y., Sanford Grossman, and Jiang Wang, 1993, Trading volume and serial correlation in stock returns, *Quarterly Journal of Economics* 108, 905-939.
- Campbell, John Y., Andrew W. Lo, and Craig A. MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press, Princeton, NJ).
- Chan, Kalok, 1993, Imperfect information and cross-autocorrelation among stock prices, *Journal of Finance* 48, 1211-1230.
- Choi, Jong Y., Dan Salandro, and Kuldeep, Shastri, 1988, On the estimation of bid-ask spreads: Theory and evidence, *Journal of Financial and Quantitative Analysis* 23, 219-230.
- Chordia, Tarun, and Bhaskaran Swaminathan, 2000, Trading volume and cross-autocorrelations in stock returns, *Journal of Finance* 55, 913-935.
- Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam, 2005, Evidence on the speed of convergence to market efficiency, *Journal of Financial Economics*, forthcoming.
- Connolly, Robert, and Chris Stivers, 2003, Momentum and reversals in equity-index returns during periods of abnormal turnover and return dispersion, *Journal of Finance* 58, 1521-1555.
- Conrad, Jennifer, and Gautam Kaul, 1988, Time-variation in expected returns, *Journal of Business* 61, 409-425.
- Conrad, Jennifer, and Gautam Kaul, 1989, Mean reversion in short-horizon expected returns, *Review of Financial Studies* 2, 225-240.
- Conrad, Jennifer, Gautam Kaul, and M. Nimalendran, 1991, Components of short-horizon individual security returns, *Journal of Financial Economics* 29, 365-384.

- Eom, Kyong Shik, Sang Buhm Hahn, and Sangyong Joo, 2004, Partial price adjustment and autocorrelation in foreign exchange markets, Preprint, University of California at Berkeley.
- Fisher, Lawrence, 1966, Some new stock market indexes, *Journal of Business* 39, 191-225.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns on buying winners and selling losers: Implications for Market Efficiency, *Journal of Finance* 48, 64-91.
- Kadlec, Gregory B., and Douglas M. Patterson, 1999, A transactions data analysis of nonsynchronous trading, *Review of Financial Studies* 12, 609-630.
- Kimelman, John, 2003, Fresh pricing is a draw for Exchange-Traded Funds, *New York Times*, (November 9).
- Lee, Charles M. C., Belinda Mucklow, and Mark J. Ready, 1993, Spreads, depths, and the impact of earnings information: An intraday analysis, *Review of Financial Studies* 6, 345-374.
- Llorente, Guillermo, Roni Michaely, Gideon Saar, and Jiang Wang, 2002, Dynamic volume-return relation of individual stocks, *Review of Financial Studies* 15, 1005-1047.
- Lo, Andrew W., and Craig A. MacKinlay, 1990, An econometric analysis of nonsynchronous trading, *Journal of Econometrics* 45, 181-212.
- McQueens, Grant, Michael Pinegar, and Steven Thorley, 1996, Delayed reaction to good news and cross-autocorrelation of portfolio returns, *Journal of Finance* 51, 889-920.
- Mech, Timothy S., 1993, Portfolio return autocorrelation, *Journal of Financial Economics* 34, 307-344.
- Roll, Richard, 1984, A simple implicit measure of the effective bid-ask spread in an efficient market, *Journal of Finance* 39, 1127-1140.
- Säfvenblad, Patrik, 2000, Trading volume and autocorrelation: Empirical evidence from the Stockholm Stock Exchange, *Journal of Banking and Finance* 24, 1275-1287.
- Scholes, Myron, and Joseph Williams, 1977, Estimating betas from nonsynchronous data, *Journal of Financial Economics* 5, 309-328.
- Sheskin, David J., 1997, *Handbook of Parametric and Nonparametric Statistical Procedures* (CRC Press, Boca Raton, FL).

FOOTNOTES

¹ The momentum effect has been cited as an explanation of medium-term (3 to 12 months) autocorrelation (see Jegadeesh and Titman (1993)); however, the focus of this paper is on daily autocorrelation of portfolios.

² In a portfolio of stocks, the individual stocks are traded; the portfolio itself is not traded, and its price is obtained by averaging the prices of the individual stocks it contains. Thus, while the price of an individual stock may bounce between the bid and ask, there is no bid or ask between which the portfolio price jumps. If the bounce process, which determines whether a given trade occurs at the bid or ask price, were independent across different stocks, bid-ask bounce would produce a slight negative autocorrelation in portfolio returns coming from the negative autocorrelation of the individual stocks in the portfolio; the cross bid-ask bounce effects would be zero. In practice, the bounce process probably shows positive correlation across stocks; if stock prices generally rise (fall) just before the close, then most stocks final trade will be at the ask (bid) price, inducing negative autocorrelation in the daily portfolio return. Thus, bid-ask bounce should cancel some of the positive autocorrelation in daily portfolio returns that results from the nonsynchronous trading effect and partial price adjustment.

³ There is a trade-off between the number of subperiods and the lengths of the subperiods. Because stock returns are very noisy, it is much easier to detect autocorrelation in longer subperiods than in shorter subperiods. Some of our results are statistically stronger when the analysis is done with five two-year periods, while others are statistically stronger with ten one-year periods.

⁴ When we ran our tests with ten one-year subperiods, the number of small and medium firms with negative return autocorrelation is strongly significant.

⁵ More precisely, the absence of partial price adjustment and time-varying risk premia imply the stated conclusion. See Appendix B for a detailed analysis of the magnitude of the potential bias resulting from time-varying risk premia.

⁶ Since there is no trade in the stock after time t_i , the intraday return also equals $r_{s,ti}$, where t is the time when the market closes.

⁷ The assumption in Roll's model that the coin tosses are independent across trades is restrictive. Choi, Salandro, and Shastri (1988) showed that serial correlation of either sign in the coin tosses affects the magnitude, but not the sign, of the autocorrelation in conventional daily returns induced by bid-ask bounce. Positive (negative) serial correlation in the coin tosses of a given stock induces negative (positive) autocorrelation of intraday returns, but it appears that the magnitude is much smaller than that of the autocorrelation of conventional daily returns. It seems likely the serial correlation of the coin tosses is positive, so we expect bid-ask bounce to induce slight negative autocorrelation of individual stock intraday returns. If we extend Roll's model to multiples stocks, and assume that the coin tosses are independent *across stocks*, the cross-autocorrelations induced by bid-ask bounce will be zero. It is unclear how restrictive the assumption of independence of the coin tosses across stocks is. If the coin tosses are correlated across stocks, it appears that the correlation should be positive: if the market as a whole is

rising, this seems likely to cause buyers to raise their bids to match the current ask; if the market as a whole is falling, this seems likely to cause sellers to lower their asks to match the current bid. Positive correlation of the coin tosses across stocks would result in negative cross-autocorrelation in daily returns, and slight negative cross-autocorrelation in intraday returns.

⁸ The specific tests for autocorrelation will be described in Section I.C.

⁹ The reader might have expected us to set the intraday return of that stock to be zero for that day. Doing so could introduce a nonsynchronous trading bias for essentially the same reason that imputing a zero return on days on which a given stock does not trade induces negative autocorrelation in individual daily stock returns. The results when the observations are included and set to zero are essentially the same.

¹⁰ Boudoukh, Richardson, and Whitelaw (1994) report first-order autocorrelation of 0.23 for weekly returns of an equally-weighted index and 0.36 for weekly returns of a small-stock portfolio.

¹¹ Using weekly data, Connolly and Stivers (2003) “find substantial momentum (reversals) in consecutive weekly returns when the latter week has unexpectedly high (low) turnover.” In contrast, Chordia and Swaminathan (2000) use turnover and return shocks for their tests, using a model specification similar to those of Campbell, Grossman, and Wang (1993) and Llorente, Michaely, Saar, and Wang (2002). Even though the model specification of Connolly and Stivers differs from that of Chordia and Swaminathan, the results from Connolly and Stivers are along the same line as those of Chordia and Swaminathan; and Llorente, Michaely, Saar, and Wang; supporting the partial price adjustment hypothesis. For other literature on the partial price adjustment hypothesis, see Brennan, Jegadeesh, and Swaminathan (1993), Mech (1993), Badrinath, Kale, and Noe (1995), MacQueen, Pinegar, and Thorley (1996), among others.

¹² As above, the reader might have expected us to set the return to zero on days on which the stock does not trade. We chose instead to omit the data for the reasons explained above. Setting the return to zero and including it in the data makes little difference in the results.

¹³ See Sheskin (1997), page 633.

¹⁴ This method is suggested in Lee, Mucklow, and Ready (1993). To test the robustness of our results, we also used a twenty percent criterion, but we found no significant difference between the two rules.

¹⁵ For example, if the TAQ dataset reported successive trades in a stock at prices of \$10, \$41, \$11, we would eliminate the transaction with a reported price of \$41.

¹⁶ Kadlec and Patterson (1999) report that the average small stock trades within 3 hours of the close and the average large stock trades within 2 minutes of the close on each day.

¹⁷ Conrad and Kaul (1988, 1989) and Conrad, Kaul, and Nimalendran (1991) (hereafter collectively abbreviated as CKN) estimate that predictable time-varying rates of return can explain 25% of the variance in weekly and monthly portfolio returns. They do not apply their methodology to daily returns; if they had, they presumably would have found a somewhat smaller percentage. As noted in the text, predictable time-varying rates of return are simply autocorrelation by another name, and are not necessarily attributable to time-varying rates of

return. CKN invoke a strong form of the Efficient Markets Hypothesis to assert that, since anyone could in principle exploit any knowledge of the time-varying rate of return, there cannot be any exploitable information. Since we are, in effect, testing a version of the Efficient Markets Hypothesis, we are unwilling to impose the Efficient Markets Hypothesis as an assumption. The predictable expected rates of return documented by CKN vary substantially from week to week, and we find it implausible that time-varying risk premia vary this much over the span of a week or two; see Ahn, Boudoukh, Richardson, and Whitelaw (2002, page 656), who note that “time variation in [risk premia] is not a high-frequency phenomenon: asset pricing models link expected returns with changing investment opportunities, which, by nature, are low-frequency events” (the original says “expected returns,” but it is clear from the context that by this, they meant risk premia as we use the terms in this paper).

¹⁸ The variations (max - min) in the three-month Treasury Bill rates for our five two-year subperiods are as follows: 3.34% (6.39%-3.05%) (93-94), 1.28% (6.40%-5.12%) (95-96), 0.86% (5.83%-4.97%) (97-98), 1.99% (6.84%-4.85%) (99-00), and 4.96% (6.27%-1.31%) (01-02). The average of the subperiod variations is 2.49%.

¹⁹ The most likely reason for a major change in the correlation of a stock with the risk factors is diversification into a new line of business, or the sale or spin-off of a line of business. In order to significantly change the correlations, the divested or acquired line of business would have to be a reasonably large fraction of the business of the firms as a whole. We eliminate in each data period any firm for which the number of outstanding shares changes by more than 10% in that data period. This should eliminate most firms that acquire substantial new business lines through acquisition, and many of those that divest substantial business lines through sale or spin-off.




Figure 1. Time diagram for Null Hypotheses VI, VIA, and VIB. Our Null Hypotheses VI (VIA, VIB) are that the correlation of each of the individual stock returns and the return of the SPDRs is zero (nonpositive, nonnegative). $r_{(d+1,i)}$ is the daily return of each individual stock on day $d+1$, computed in the conventional way: the price at the final trade on day $d+1$, minus the price at the last trade prior to day $d+1$, divided by the price at the last trade prior to day $d+1$. We compute the correlation between the return of stock i on day $d+1$ (in other words, the return from the final trade of the stock on day d to the final trade of the stock on day $d+1$, corresponding to the intervals B and C) with the return of the SPDRs over the interval from the time of the last trade of the SPDRs on day $d-1$ through the time of the last trade of the *stock* on day d , corresponding to the interval A. If a stock does not trade on day d or the stock does not trade on day $d+1$, we omit the data from our calculation.

Table I
Descriptive Statistics of Data

This Table provides descriptive statistics for our 300 NYSE-listed sample firms, stratified in three groups by market capitalization, and for the SPDRs. We sampled the firms every two-year, five times over the ten years spanned from 1993 to 2002. The Table presents the number of firms, the market capitalization (in millions of dollars; max, mean, and min), daily portfolio returns (mean and standard deviation), average daily trading volume (in shares), average time interval between the closing trade of the individual stock and the closing trade of the SPDRs (in seconds), and average number of days on which trade occurs. All statistics except max and min of market capitalization are averages over five two-year intervals. Max and min of market capitalization denote max of max and min of min. The closing trade of an individual stock is the last trade occurring before 4:00 p.m.; the closing trade of the SPDRs is the first trade reported after 4:00 p.m., except on 26 days immediately before holidays on which the market closed early, where we take the last trade occurring before 4:00 p.m. We also report the first-order autocorrelation of SPDRs over ten years using conventional daily returns. All daily returns are calculated in the conventional way: the price at the closing trade on day d , minus the price at the last trade prior to day d , divided by the price at the last trade prior to day d . We use transaction data from TAQ database over the sample period from January 4, 1993 to December 31, 2002. ¶ denotes the statistics of portfolios, not average of individual firms of each group. † denotes that SPDRs introduced to the American Stock Exchange on January 29, so that its trading day are short of 19 trading days compared to other stocks.

	Small Firm	Medium Firm	Large Firm	SPDRs
Number of observations		504.8		
Number of firms	100	100	100	-
Market capitalization (in mil. of dollars)				
Max	346.9	1,225.3	475,003.2	-
Mean	135.9	755.9	26,304.6	-
Min	23.9	454.1	4,716.8	-
Daily <i>portfolio</i> returns¶				
Mean (%)	0.0604	0.0395	0.0348	0.0004
Std. dev. (%)	0.6273	0.8251	0.8990	0.0108
Average daily trading volume (in shares)	22,200.3	105,637.4	1,458,878.2	42,603.8
Average time interval between the closing trade of individual stock and the closing trade of SPDRs (in seconds)	-2,839.4	-682.5	-73.5	-
Average number of days on which trade occurs	490.6	503.6	504.0	500.2†
First-order autocorrelation of SPDRs				
Conventional daily returns (Standard error)	-	-	-	-0.0199 (0.0558)

Table II
Average *Individual* Daily Return Autocorrelations: Conventional Daily Returns
(Null Hypotheses I, IA, IB)

This Table reports the average *individual* daily returns autocorrelations using conventional daily returns. The conventional daily return of each individual stock on day d is computed in the usual way: the price at the final trade on day d , minus the price at the final trade on day $d-1$, divided by the price at the final trade on day $d-1$. Our sample consists of 300 NYSE-listed sample firms, stratified into the three groups. We use transaction data from the TAQ database over the sample period from January 4, 1993 to December 31, 2002. Numbers in parenthesis are standard errors. ** and * denote positive significance at the 1% and 5% level, respectively; ++ denotes negative significance at the 1% level. The last column shows p -value of the average individual daily returns autocorrelation over ten years from 1993 to 2002.

Portfolio	Number of Firms	Average <i>Individual</i> Daily Returns Autocorrelations:						p -value
		Conventional Daily Returns						
		93~94	95~96	97~98	99~00	01~02	1993~2002	
Small firm	100	0.0084 (0.0102)	0.0034 (0.0112)	0.0664** (0.0093)	0.0412** (0.0085)	0.0330** (0.0098)	0.0335** (0.0043)	0.0000
Med. firm	100	0.0173* (0.0082)	-0.0058 (0.0076)	0.0208* (0.0097)	0.0119 (0.0080)	-0.0354++ (0.0075)	-0.0009 (0.0036)	0.5938
Large firm	100	0.0115* (0.0058)	-0.0011 (0.0058)	-0.0082 (0.0057)	0.0217** (0.0061)	-0.0076++ (0.0055)	0.0026 (0.0026)	0.8413

Table III
Autocorrelation of Daily *Individual* Stock Returns: Conventional Daily Returns
(Null Hypotheses II, IIA, IIB)

This Table reports the results of our tests of individual-stock daily return autocorrelations using conventional daily returns. The conventional daily return of each individual stock on day d is computed in the usual way: the price at the final trade on day d , minus the price at the final trade on day $d-1$, divided by the price at the final trade on day $d-1$. Our sample consists of 300 NYSE-listed sample firms, stratified into the three groups. We use three methods to test whether the correlation is zero: the Pearson correlation test (Panel A), a modified Pearson correlation test (Panel B), and the Kendall's tau test (Panel C). Our modified Pearson test uses Andrew's (1991) heteroskedasticity and autoregression consistent (HAC) covariance estimator for the estimation and test of the correlation coefficient. We obtain the Andrew's HAC covariance using the quadratic spectral (QS) kernel with automatic bandwidth selection method. We use transaction data from the TAQ database over the sample period from January 4, 1993 to December 31, 2002. + and - denote the numbers of stocks with statistically significant positive and negative autocorrelation, at the 2.5% level; +- denotes the numbers of stocks with statistically significant autocorrelation in a two-sided test, at the 5% level. X_1 denotes the first order statistic (minimum) of the observations for the five two subperiods, while X_2 denotes the second order statistic (second smallest). p_1 and p_2 denote the probability that X_1 and X_2 , respectively, would exceed the observed value in a nonparametric test using only the fact that the numbers in + and - are nonnegative random variables with expectation 2.5. $p_3 = 2 \min\{p_1, p_2\}$ is an upper bound on the probability that either X_1 or X_2 would exceed the observed value. ** and * denote significance at the 1% and 5% level using p_3 .

Portfolio	Autocorrelation of Daily <i>Individual</i> Stock Returns: Conventional Daily Returns									
	93~94	95~96	97~98	99~00	01~02	1993~2002				
						X_1	p_1	X_2	p_2	p_3
Panel A: Pearson Correlation Test										
Small firm: +	23	22	42	30	26	22	0.0000	23	0.0006	0.0000**
-	16	23	5	7	10	5	0.0313	7	0.0581	0.0625
+-	39	45	47	37	36	36	0.0001	37	0.0015	0.0001**
Med. firm: +	19	12	22	16	4	4	0.0954	12	0.0078	0.0157*
-	12	17	11	7	22	7	0.0058	11	0.0109	0.0116*
+-	31	29	33	23	26	23	0.0005	26	0.0058	0.0010**
Large firm: +	7	4	6	15	1	1	1.0000	4	0.3815	0.7629
-	4	8	9	5	6	4	0.0954	5	0.1875	0.1907
+-	11	12	15	20	7	7	0.1859	11	0.1358	0.2717
Panel B: Modified Pearson Correlation Test										
Small firm: +	18	16	35	23	23	16	0.0001	18	0.0017	0.0002**
-	13	15	5	7	6	5	0.0313	6	0.1005	0.0625
+-	31	31	40	30	29	29	0.0002	30	0.0033	0.0003**
Med. firm: +	19	8	19	13	3	3	0.4019	8	0.0358	0.0072**
-	9	15	9	4	21	4	0.0954	9	0.0232	0.0463*
+-	28	23	28	17	24	17	0.0022	23	0.0092	0.0044**
Large firm: +	7	5	5	11	1	1	1.0000	5	0.1875	0.3750
-	3	5	6	2	6	2	1.0000	3	0.8038	1.0000
+-	10	10	11	13	7	7	0.1859	10	0.1875	0.3719
Panel C: Kendall's Tau Test										
Small firm: +	9	16	36	26	26	9	0.0017	16	0.0026	0.0033**
-	25	32	13	9	11	9	0.0017	11	0.0109	0.0033**
+-	34	48	49	35	37	34	0.0001	35	0.0018	0.0001**

Med. firm: +	20	10	23	15	3	3	0.4019	10	0.0156	0.0313*
-	17	26	12	6	27	6	0.0126	12	0.0078	0.0157*
+-	37	36	35	21	30	21	0.0008	30	0.0033	0.0015**
Large firm: +	5	4	5	6	5	4	0.0954	5	0.1875	0.1907
-	3	7	2	2	10	2	1.0000	2	1.0000	1.0000
+-	8	11	7	8	15	7	0.1859	8	0.3815	0.3719

Table IV
Average Individual Daily Returns Autocorrelations: Intraday Returns
(Null Hypotheses III, IIIA, IIIB)

This Table reports the average *individual* daily returns autocorrelations using intraday returns. The intraday return on day d of each stock in the portfolio is defined as the price at the final trade on day d , minus the price at the first trade on day d , divided by the price at the first trade on day d . Our sample consists of 300 NYSE-listed sample firms, stratified into the three groups. We use transaction data from the TAQ database over the sample period from January 4, 1993 to December 31, 2002. Numbers in parenthesis are standard errors. ** and * denote positive significance at the 1% and 5% level, respectively; ++ and + denote negative significance. The last column shows p -value of the average individual daily returns autocorrelation over ten years from 1993 to 2002.

Portfolio	Number of Firms	Average Individual Daily Returns Autocorrelations:						
		Intraday Returns						p -value
		93~94	95~96	97~98	99~00	01~02	1993~2002	
Small firm	100	0.0389** (0.0074)	0.0394** (0.0084)	0.0729** (0.0090)	0.0441** (0.0083)	0.0240** (0.0092)	0.0435** (0.0038)	0.0000
Med. firm	100	0.0543** (0.0071)	0.0321** (0.0065)	0.0326** (0.0096)	0.0265** (0.0067)	-0.0376** (0.0073)	0.0218** (0.0032)	0.0000
Large firm	100	0.0064 (0.0062)	-0.0133+ (0.0063)	-0.0324** (0.0066)	0.0126* (0.0058)	-0.0103 (0.0054)	-0.0064** (0.0027)	0.9918

Table V
Autocorrelation of Daily *Individual* Stock Returns: Intraday Returns
(Null Hypotheses IV, IVA, IVB)

This Table reports the results of our tests of individual-stock daily return autocorrelations using intraday returns. The intraday return on day d of each stock in the portfolio is defined as the price at the final trade on day d , minus the price at the first trade on day d , divided by the price at the first trade on day d ; Our sample consists of 300 NYSE-listed sample firms, stratified into the three groups. We use three methods to test whether the correlation is zero: the Pearson correlation test (Panel A), a modified Pearson correlation test (Panel B), and the Kendall's tau test (Panel C). Our modified Pearson test uses Andrew's (1991) heteroskedasticity and autoregression consistent (HAC) covariance estimator for the estimation and test of the correlation coefficient. We obtain the Andrew's HAC covariance using the quadratic spectral (QS) kernel with automatic bandwidth selection method. We use transaction data from the TAQ database over the sample period from January 4, 1993 to December 31, 2002. + and - denote the numbers of stocks with statistically significant positive and negative autocorrelation, at the 2.5% level; +- denotes the numbers of stocks with statistically significant autocorrelation in a two-sided test, at the 5% level. X_1 denotes the first order statistic (minimum) of the observations for the five two subperiods, while X_2 denotes the second order statistic (second smallest). p_1 and p_2 denote the probability that X_1 and X_2 , respectively, would exceed the observed value in a nonparametric test using only the fact that the numbers in + and - are nonnegative random variables with expectation 2.5. $p_3 = 2 \min\{p_1, p_2\}$ is an upper bound on the probability that either X_1 or X_2 would exceed the observed value. ** and * denote significance at the 1% and 5% level using p_3 .

Portfolio	Autocorrelation of Daily <i>Individual</i> Stock Returns: Intraday Returns									
	93~94	95~96	97~98	99~00	01~02	1993~2002				
						X_1	p_1	X_2	p_2	P_3
Panel A: Pearson Correlation Test										
Small firm: +	28	28	41	32	23	23	0.0000	28	0.0003	0.0000**
-	3	7	4	8	11	3	0.4019	4	0.3815	0.7629
+-	31	35	45	40	34	31	0.0001	34	0.0021	0.0002**
Med. firm: +	36	23	27	16	2	2	1.0000	16	0.0026	0.0052**
-	3	4	10	5	23	3	0.4019	4	0.3815	0.7629
+-	39	27	37	21	25	21	0.0008	25	0.0067	0.0015**
Large firm: +	11	7	6	11	3	3	0.4019	6	0.1005	0.2009
-	4	12	20	6	5	4	0.0954	5	0.1875	0.1907
+-	15	19	26	17	8	8	0.0954	15	0.0453	0.0905
Panel B: Modified Pearson Correlation Test										
Small firm: +	24	25	33	23	21	21	0.0000	23	0.0006	0.0000**
-	1	5	3	2	8	1	1.0000	2	1.0000	1.0000
+-	25	30	36	25	29	25	0.0003	25	0.0067	0.0006**
Med. firm: +	33	21	21	13	3	3	0.4019	13	0.0058	0.0116*
-	2	3	7	3	19	2	1.0000	3	0.8038	1.0000
+-	35	24	28	16	22	16	0.0030	22	0.0109	0.0060**
Large firm: +	7	5	4	9	1	1	1.0000	4	0.3815	0.7629
-	4	8	14	3	3	3	0.4019	3	0.8038	0.8038
+-	11	13	18	12	4	4	1.0000	11	0.1358	0.2716
Panel C: Kendall's Tau Test										
Small firm: +	28	27	40	39	25	25	0.0000	27	0.0003	0.0000**
-	4	5	1	2	10	1	1.0000	2	1.0000	1.0000
+-	32	32	41	41	35	32	0.0001	32	0.0026	0.0002**
Med. firm: +	43	24	28	20	4	4	0.0954	20	0.0011	0.0022**

	-	2	3	6	3	24	2	1.0000	3	0.8038	1.0000
	+/-	45	27	34	23	28	23	0.0005	27	0.0050	0.0010**
Large firm:	+	11	5	5	5	3	3	0.4019	5	0.1875	0.3750
	-	3	5	8	1	8	1	1.0000	3	0.8038	1.0000
	+/-	14	10	13	6	11	6	0.4019	10	0.1875	0.3750

Table VI

First-Order Autocorrelation of Daily *Portfolio* Returns: Conventional Daily Returns

This Table reports the first-order autocorrelation of conventional daily *portfolio* returns using conventional daily returns. The conventional daily return on each stock on day d is defined as the closing price (the price at the final trade of the day) on day d , less the closing price on day $d-1$, divided by the closing price on day $d-1$. The conventional daily return of the portfolio on day d is an equally-weighted average of the conventional daily returns of the stocks in the portfolio, omitting stocks that do not trade on day d . Our sample consists of 300 NYSE-listed sample firms, stratified into the three groups. We use three methods to test whether the correlation is zero: the Pearson correlation test (Panel A), a modified Pearson correlation test (Panel B), and the Kendall's tau test (Panel C). Our modified Pearson test uses Andrews' (1991) heteroskedasticity and autoregression consistent (HAC) covariance estimator for the estimation and test of the correlation coefficient. We obtain Andrews' HAC covariance using the quadratic spectral (QS) kernel with automatic bandwidth selection method. We use transaction data from the TAQ database over the sample period from January 4, 1993 to December 31, 2002. Numbers in parenthesis are standard errors. ** and * denote significance at the 1% and 5% level, respectively.

Portfolio	Number of Firms	First-Order Autocorrelation of Daily <i>Portfolio</i> Returns: Conventional Daily Returns					
		93~94	95~96	97~98	99~00	01~02	1993~2002
Panel A: Pearson Correlation Test							
Small firm	100	0.2722**	0.2299**	0.3854**	0.2750**	0.1109*	0.2550**
Med. firm	100	0.2244**	0.1825**	0.2353**	0.1677**	0.0136	0.1651**
Large firm	100	0.0441	0.1009*	0.0057	0.1060*	0.0219	0.0558**
(Std. error)		(0.0445)	(0.0445)	(0.0445)	(0.0446)	(0.0448)	(0.0199)
Panel B: Modified Pearson Correlation Test							
Small firm	100	0.2881**	0.2514**	0.3977**	0.2736**	0.1333	0.2893**
		(0.0499)	(0.0397)	(0.0422)	(0.0435)	(0.0745)	(0.0208)
Med. firm	100	0.2250**	0.1888**	0.2321**	0.1618**	0.0165	0.1664**
		(0.0536)	(0.0478)	(0.0647)	(0.0454)	(0.0586)	(0.0236)
Large firm	100	0.0426	0.1130*	-0.0074	0.1006	0.0237	0.0550**
		(0.0412)	(0.0469)	(0.0540)	(0.0533)	(0.0476)	(0.0214)
Panel C: Kendall's Tau Test							
Small firm	100	0.1770**	0.1434**	0.2522**	0.1951**	0.0472	0.1632**
Med. firm	100	0.1360**	0.1384**	0.1853**	0.0997**	0.0128	0.1147**
Large firm	100	0.0264	0.0561	0.0387	0.0352	0.0102	0.0334**
(Std. error)		(0.0298)	(0.0298)	(0.0298)	(0.0299)	(0.0300)	(0.0134)

Table VII
First-Order Autocorrelation of Daily *Portfolio* Returns: Intraday Returns
(Null Hypothesis V)

This Table reports the first-order autocorrelation of daily *portfolio* returns using intraday returns. The intraday return on day d of each stock in the portfolio is defined as the price at the final trade on day d , minus the price at the first trade on day d , divided by the price at the first trade on day d ; the intraday return of the portfolio is defined as the equally-weighted average of the intraday returns of the stocks in the portfolio, omitting any stocks that do not trade at least twice on the day. Our sample consists of 300 NYSE-listed sample firms, stratified into the three groups. We use three methods to test whether the correlation is zero: the Pearson correlation test (Panel A), a modified Pearson correlation test (Panel B), and the Kendall's tau test (Panel C). Our modified Pearson test uses Andrew's (1991) heteroskedasticity and autoregression consistent (HAC) covariance estimator for the estimation and test of the correlation coefficient. We obtain the Andrew's HAC covariance using the quadratic spectral (QS) kernel with automatic bandwidth selection method. We use transaction data from the TAQ database over the sample period from January 4, 1993 to December 31, 2002. Numbers in parenthesis are standard errors. ** and * denote positive significance at the 1% and 5% level, respectively; + denotes negative significance at the 5% level.

Portfolio	Number of Firms	First-Order Autocorrelation of Daily <i>Portfolio</i> Returns: Intraday Returns					
		93~94	95~96	97~98	99~00	01~02	1993~2002
Panel A: Pearson Correlation Test							
Small firm	100	0.2046**	0.2118**	0.3623**	0.2541**	0.0579	0.2184**
Med. firm	100	0.1643**	0.2031**	0.2117**	0.1781**	-0.0009	0.1515**
Large firm	100	-0.0616	-0.0384	-0.1072 ⁺	0.0921*	0.0055	-0.0220
(Std. error)		(0.0445)	(0.0445)	(0.0445)	(0.0446)	(0.0447)	(0.0199)
Panel B: Modified Pearson Correlation Test							
Small firm	100	0.2274**	0.2749**	0.3696**	0.2513**	0.0790	0.2555**
		(0.0551)	(0.0386)	(0.0611)	(0.0462)	(0.0738)	(0.0228)
Med. firm	100	0.1617**	0.2217**	0.2127**	0.1821**	0.0020	0.1534**
		(0.0513)	(0.0478)	(0.0874)	(0.0480)	(0.0533)	(0.0240)
Large firm	100	-0.0608	-0.0284	-0.1017	0.0857	0.0054	-0.0187
		(0.0402)	(0.0501)	(0.0714)	(0.0543)	(0.0473)	(0.0224)
Panel C: Kendall's Tau Test							
Small firm	100	0.1471**	0.1209**	0.2561**	0.1917**	0.0143	0.1462**
Med. firm	100	0.0864**	0.1513**	0.1896**	0.0904**	0.0070	0.1051**
Large firm	100	-0.0434	-0.0197	-0.0333	0.0267	0.0077	-0.0124
(Std. error)		(0.0298)	(0.0298)	(0.0298)	(0.0298)	(0.0299)	(0.0133)

Table VIII

Proportion of Partial Price Adjustment in the Autocorrelation of the Portfolio Returns

This Table presents the standard deviation, autocorrelation, autocorrelation weighted by ratio of standard deviation, and autocovariance of conventional and intraday returns in each of the small-, medium-, and large-firm portfolios. The conventional daily return of each individual stock on day d is computed in the usual way: the price at the final trade on day d . Our sample consists of 300 NYSE-listed sample firms, stratified into the three groups. We use transaction data from the TAQ database over the sample period from January 4, 1993 to December 31, 2002. The intraday return of each individual stock on day d is defined in the following way: the price at the final trade on day d , minus the price at the first trade on day d , divided by the price at the first trade on day d . Our sample consists of 300 NYSE-listed sample firms, stratified into the three groups. We use transaction data from the TAQ database over the sample period from January 4, 1993 to December 31, 2002.

Portfolio	93~94	95~96	97~98	99~00	01~02
Panel A: Standard Deviation					
Small firm: Conventional daily return (%)	0.5008	0.4544	0.6232	0.5793	0.9789
Intraday return (%)	0.3950	0.3549	0.4963	0.4879	0.8920
Med. firm: Conventional daily return (%)	0.4757	0.4706	0.9388	0.9102	1.3303
Intraday return (%)	0.4146	0.3971	0.8099	0.8418	1.2184
Large firm: Conventional daily return (%)	0.5498	0.5845	0.9368	1.0846	1.3394
Intraday return (%)	0.4831	0.5114	0.8354	0.9310	1.1693
Panel B: Autocorrelation					
Small firm: Conventional daily return	0.2722	0.2299	0.3854	0.2750	0.1109
Intraday return	0.2046	0.2118	0.3623	0.2541	0.0589
Med. firm: Conventional daily return	0.2244	0.1825	0.2353	0.1677	0.0135
Intraday return	0.1643	0.2031	0.2117	0.1781	-0.0009
Large firm: Conventional daily return	0.0441	0.1009	0.0057	0.1060	0.0219
Intraday return	-0.0616	-0.0384	-0.1072	0.0921	0.0055
Panel C: Autocovariance (Autocorrelation times Variance)					
Small firm: Conventional daily return (%)	0.0683	0.0475	0.1497	0.0923	0.1063
Intraday return (%)	0.0319	0.0267	0.0892	0.0605	0.0476
Residual (%)	0.0363	0.0208	0.0604	0.0318	0.0587
Intraday return autocovariance as percentage of intraday plus residual return autocovariance	46.76	56.20	59.62	65.54	44.77
Med. firm: Conventional daily return (%)	0.0508	0.0404	0.2074	0.1389	0.0239
Intraday return (%)	0.0282	0.0320	0.1389	0.1262	-0.0013
Residual (%)	0.0225	0.0084	0.0685	0.0127	0.0252
Intraday return autocovariance as percentage of intraday plus residual return autocovariance	55.62	79.24	66.96	90.84	5.03
Large firm: Conventional daily return (%)	0.0133	0.0345	0.0050	0.1247	0.0393
Intraday return (%)	-0.0144	-0.0100	-0.0748	0.0798	0.0075
Residual (%)	0.0277	0.0445	0.0798	0.0499	0.0318
Intraday return autocovariance as percentage of conventional daily return autocovariance	34.16	18.41	48.38	64.02	19.14

Table IX
ETFs
(Null Hypotheses VI, VIA, VIB)

This Table reports the number of rejections of Null Hypotheses VI (VIA, VIB). If the nonsynchronous trading effect and bid-ask bounce are the sole sources of stock return autocorrelation, the correlation between the return of the individual stock and the return of the SPDRs must be less than or equal to zero. Our Null Hypotheses VI (VIA, VIB) are that the correlation of each of the individual stock returns and the return of the SPDRs is zero (nonpositive, nonnegative). The daily return of each individual stock on day d is computed in the conventional way: the price at the final trade on day d , minus the price at the last trade prior to day d , divided by the price at the last trade prior to day d . We compute the correlation between the return of stock i on day $d+1$ (in other words, the return from the final trade of the stock on day d to the final trade of the stock on day $d+1$) with the return of the SPDRs over the interval from the time of the last trade of the SPDRs on day $d-1$ through the time of the last trade of the *stock* on day d . If a stock does not trade on day d or the stock does not trade on day $d+1$, we omit the data from our calculation. Our sample consists of 300 NYSE-listed sample firms, stratified into the three groups. We use three methods to test whether the correlation is zero: the Pearson correlation test (Panel A), a modified Pearson correlation test (Panel B), and the Kendall's tau test (Panel C). Our modified Pearson test uses Andrew's (1991) heteroskedasticity and autoregression consistent (HAC) covariance estimator for the estimation and test of the correlation coefficient. We obtain the Andrew's HAC covariance using the quadratic spectral (QS) kernel with automatic bandwidth selection method. We use transaction data from the TAQ database over the sample period from January 4, 1993 to December 31, 2002. + and - denote the numbers of stocks with statistically significant positive and negative correlation with the SPDRs, at the 2.5% level; +- denotes the numbers of stocks with statistically significant correlation in a two-sided test, at the 5% level. X_1 denotes the first order statistic (minimum) of the observations for the five two subperiods, while X_2 denotes the second order statistic (second smallest). p_1 and p_2 denote the probability that X_1 and X_2 , respectively, would exceed the observed value in a nonparametric test using only the fact that the numbers in + and - are nonnegative random variables with expectation 2.5. $p_3 = 2 \min\{p_1, p_2\}$ is an upper bound on the probability that either X_1 or X_2 would exceed the observed value. ** and * denote significance at the 1% and 5% levels using p_3 .

Portfolio	Correlation of Daily <i>Individual</i> Stock Returns and SPDRs									
	93~94	95~96	97~98	99~00	01~02	X_1	p_1	X_2	p_2	P_3
Panel A: Pearson Correlation Test										
Small firm: +	17	21	41	15	17	15	0.0001	17	0.0021	0.0003**
-	0	0	1	0	10	0	1.0000	0	1.0000	1.0000
+-	17	21	42	15	27	15	0.0041	17	0.0286	0.0082**
Med. firm: +	35	22	35	15	7	7	0.0058	15	0.0033	0.0067**
-	1	1	7	7	13	1	1.0000	1	1.0000	1.0000
+-	36	23	42	22	20	20	0.0010	22	0.0109	0.0020**
Large firm: +	10	23	13	5	9	5	0.0313	9	0.0231	0.0463*
-	2	1	22	25	6	1	1.0000	2	1.0000	1.0000
+-	12	24	35	30	15	12	0.0126	15	0.0453	0.0251*
Panel B: Modified Pearson Correlation Test										
Small firm: +	15	22	40	13	16	13	0.0003	15	0.0033	0.0005**
-	0	0	1	1	10	0	1.0000	0	1.0000	1.0000
+-	15	22	41	14	26	14	0.0058	15	0.0453	0.0116*
Med. firm: +	34	21	33	14	7	7	0.0058	14	0.0044	0.0087**
-	1	1	7	9	15	1	1.0000	1	1.0000	1.0000
+-	35	22	40	23	22	22	0.0006	22	0.0109	0.0012**

Large firm: +	11	24	12	6	9	6	0.0126	9	0.0232	<i>0.0251*</i>
-	2	1	23	25	6	1	1.0000	2	1.0000	<i>1.0000</i>
+ -	13	25	35	31	15	13	0.0084	15	0.0453	<i>0.0168*</i>
Panel C: Kendall's Tau Test										
Small firm: +	16	20	41	9	16	9	0.0017	16	0.0026	<i>0.0033**</i>
-	0	0	1	0	14	0	1.0000	0	1.0000	<i>1.0000</i>
+ -	16	20	42	9	30	9	0.0529	16	0.0358	<i>0.0715</i>
Med. firm: +	33	24	49	14	6	6	0.0126	14	0.0044	<i>0.0087**</i>
-	0	1	5	7	17	0	1.0000	1	1.0000	<i>1.0000</i>
+ -	33	25	54	21	23	21	0.0008	23	0.0092	<i>0.0015**</i>
Large firm: +	9	19	26	7	10	7	0.0058	9	0.0232	<i>0.0116*</i>
-	4	1	15	25	10	1	1.0000	4	0.3815	<i>0.7629</i>
+ -	13	20	41	32	20	13	0.0084	20	0.0156	<i>0.0168*</i>