Stock return autocorrelation is not spurious

Robert M. Anderson[a,*], Kyong Shik Eom[b], Sang Buhm Hahn[c],
Jong-Ho Park[d]

[a] *University of California at Berkeley, Department of Economics, 549 Evans Hall #3880, Berkeley, CA 94720-3880, USA*
[b] *KSRI, 45-2 Yoido-dong, Youngdeungpo-gu, Seoul, 150-974, Korea, and University of California at Berkeley, Department of Economics, 549 Evans Hall #3880, Berkeley, CA 94720-3880, USA*
[c] *KSRI, 45-2 Yoido-dong, Youngdeungpo-gu, Seoul, 150-974, Korea*
[d] *Sunchon National University, Department of Business Administration, 315 Maegok-dong, Sunchon, Chonnam 540-742, Korea*

Abstract

We decompose stock return autocorrelation into spurious components—the nonsynchronous trading effect (NT) and bid-ask bounce (BAB)—and genuine components—partial price adjustment (PPA) and time-varying risk premia (TVRP), using four key ideas: theoretically signing or bounding the components; computing returns over disjoint subperiods separated by a trade to eliminate NT and greatly reduce BAB; dividing the data period into disjoint subperiods to obtain independent measures of autocorrelation; and computing the portion of the autocorrelation that can be unambiguously attributed to PPA. We analyze daily individual and portfolio return autocorrelations in ten years' NYSE transaction data and find compelling evidence that the PPA is a major source of the autocorrelation.

First Version: *February*, *2003*
This Version: *May 26*, *2008*

JEL classification: G12; G14; D40; D82
Keywords: Stock return autocorrelation; Nonsynchronous trading; Partial price adjustment; Market microstructure; Open-to-close return; SPDRs

---

We are grateful to Dong-Hyun Ahn, Jonathan Berk, Greg Duffee, Bronwyn Hall, Joel Hasbrouck, Rich Lyons, Ulrike Malmendier, Mark Rubinstein, Paul Ruud, Jacob Sagi, and Adam Szeidl for helpful comments. This work was supported by the Korea Research Foundation Grant funded by the Korean Government (MOEHRD) (KRF-2005-042-B00081). Anderson's research was also supported by Grant SES-0214164 from the U.S. National Science Foundation and the Coleman Fung Chair in Risk Management at UC Berkeley; he is grateful for the hospitality of the Korea Securities Research Institute.
* Corresponding author. University of California at Berkeley, Department of Economics, 549 Evans Hall #3880, Berkeley, CA 94720-3880, USA. Tel.: +1 510 642 5248; fax: +1 510 642 6615.

E-mail addresses: anderson@econ.berkeley.edu (R.M. Anderson), kseom@ksri.org (K.S. Eom), sbhahn@ksri.org (S.B. Hahn), schrs@sunchon.ac.kr (J.-H. Park).

Stock return autocorrelation is not spurious

1. Introduction

One of the most visible stylized facts in empirical finance is the autocorrelation of stock returns at fixed intervals (daily, weekly, monthly). This autocorrelation presented a challenge to the main models in continuous-time finance, which rely on some form of the random walk hypothesis. Consequently, there is an extensive literature on stock return autocorrelation; it occupies 55 pages of Campbell, Lo, and MacKinlay (1997).[1] The results of this literature were, however, inconclusive; see the Literature review in Section 2.

Over the last fifteen years, as increasing computer power and new statistical methods have permitted the analysis of very large datasets using intraday data, the focus has shifted from autocorrelation at fixed intervals to the varying speed of price discovery across various assets. This new literature clearly shows that price adjustment occurs at varying speeds in different asset classes, and that trades occur as prices are adjusting to reflect new information; thus, the new literature refutes at least the more stringent forms of the random walk hypothesis. The price discovery literature has not been particularly concerned with daily return autocorrelation.

Our goal is to show that simple methods, applied to intraday data, allow us to resolve the questions concerning *daily return autocorrelation* left unanswered by the literature. It is not our goal to compete with the price discovery literature. Moreover, it is not our goal to revisit the random walk hypothesis, except to study the role that the failure of the random walk hypothesis plays in daily return autocorrelation.

Daily return autocorrelation has been attributed to four main sources: spurious autocorrelation arising from market microstructure biases, including the nonsynchronous trading effect (NT) (in which correlations are calculated using stale prices) and bid-ask bounce (BAB), and genuine autocorrelation arising from partial price adjustment (PPA) (i.e., trade takes place at prices that do not fully reflect the information possessed by traders) and time-varying risk premia (TVRP).[2] Our use of the terms "spurious" to describe the NT and BAB effects and "genuine" to describe PPA and TVRP follows the terminology of Campbell, Lo, and MacKinlay (1997).[3] The term "spurious" indicates that NT and BAB arise from microstructure sources which bias the autocorrelation tests.[4] This bias would produce the appearance of autocorrelation even if the underlying "true" securities price process were a process such as geometric Brownian motion with constant drift.

The price discovery literature clearly establishes PPA. However, because that literature has paid little attention to daily return autocorrelation, it does not tell us whether or not PPA plays a significant role in daily return autocorrelation. Since daily return autocorrelation remains one of the most visible stylized facts in empirical finance, it is desirable to have a clear understanding of its sources and their respective magnitudes.

In this paper, we analyze ten years' worth of transaction data and find that daily individual stock and portfolio returns on the NYSE *are* serially correlated. Anderson (2006) shows that, in our setting, TVRP is sufficiently small that it can be ignored in our tests. The hypothesis that PPA makes no contribution to the autocorrelation of individual stock and portfolio returns is strongly rejected in all of our tests involving small and medium firms, and in some of our tests involving large firms. We conclude that PPA must be a significant source, and in some cases the main source, of autocorrelation in individual stock and portfolio returns.

We propose several methods to decompose individual stock and portfolio return autocorrelation into the four components. Our methods rely on established properties of the spurious components (BAB and NT) and on bounding TVRP to identify a portion of the stock return autocorrelation that can only be attributed to PPA; they make no assumptions whatsoever on the causes or form of PPA, and are thus direct. Each of our methods uses one or a combination of the following four key ideas:

Sign and/or Bound the Sources of Autocorrelation Theoretically.

- NT is negative for individual stock returns, and is generally positive for portfolio returns.
- BAB is negative for both individual stock and portfolio returns, and is generally considered to be very small for portfolio returns.[5] Some of our tests greatly reduce BAB. Thus, we shall assume that for our portfolio tests, the contribution of BAB to portfolio return autocorrelation is zero.
- PPA can be either positive or negative for both individual stock and portfolio returns.[6]
- TVRP is positive for both individual stock and portfolio returns. To the best of our knowledge, no paper has asserted that TVRP is a significant source of autocorrelation in the empirical setting (daily returns of individual stocks or portfolios and autocorrelations calculated over two-year time

horizons) considered in this paper. Anderson (2006) calculates an upper bound on the size of TVRP, depending on the return period (daily, weekly, monthly, quarterly, annual returns), the time horizon over which the autocorrelations are calculated, and the variation in risk premia. For plausible values of the variation in risk premia, TVRP is too small to affect the tests in this paper; consequently, in the discussion of our tests, we shall assume there are only three sources, (NT, BAB, and PPA) for daily return autocorrelation of individual stocks, and only two sources (NT and PPA) for daily portfolio return autocorrelation.[7] If we find statistically significant positive autocorrelation in individual stock returns, it can only come from PPA. If we find statistically significant negative autocorrelation in portfolio returns, it can only come from PPA.

Eliminate NT by Computing Returns over Disjoint Return Subperiods, Separated by a Trade.—NT arises when autocorrelations are computed using stale prices. If we compute stock returns in a way that stale prices are never used, NT will be eliminated.

- For individual daily stock returns, NT arises when a stock is not traded for several days, so the return is taken to be zero, then the stock trades and incorporates several days' worth of trend all in one day. Conventional daily return on a given day is defined as the price of the last trade on that day, minus the price of the last trade prior to the day, divided by the price of the last trade prior to the day. We define the *open-to-close return* on a given day to be the price of the last trade on that day, minus the price of the first trade on that day, divided by the price of the first trade on that day. Any catch-up to trend will be reflected in the first trade of the day, and consequently will not be incorporated into the open-to-close return. On days in which fewer than two trades occur, we treat the data as missing rather than zero. This eliminates NT, and greatly reduces BAB. Since TVRP can be ignored in our setting, we conclude that statistically significant open-to-close return autocorrelation for individual stocks can only come from PPA.
- Portfolio return autocorrelation is not the average of the individual return autocorrelations of stocks in the portfolio. Instead, it is dominated by cross-autocorrelations of the returns of the individual stocks in the portfolio. NT arises when information arrives during a given day. For stocks that trade after the information arrives, the information is reflected in the return that day; for stocks that do not trade after the information arrives, the information is reflected in the return the next day, so we see positive cross-autocorrelation in the individual stocks in the portfolio, generating positive portfolio return autocorrelation. We define the open-to-close portfolio return

on a given day to be the average of the open-to-close returns for that day of the individual stocks in the portfolio. We then define the open-to-close return autocorrelation of the portfolio to be the autocorrelation of the open-to-close portfolio returns. This eliminates NT, and greatly reduces whatever small amount of BAB there may be in portfolio returns. Since TVRP can be ignored in our setting, we conclude that statistically significant autocorrelation in open-to-close portfolio returns can only come from PPA.

- Exchange-traded funds (ETFs) are mutual funds that trade in real-time during the market day, rather than only once per day at net asset value. Many ETFs, especially those based on broad market indices, trade essentially continuously. SPDRs are an ETF based on the Standard and Poor's 500 index (S&P 500). Since the return autocorrelation of a portfolio is essentially the average of the return *cross*-autocorrelations of the pairs of stocks in the portfolio, we compute the cross-autocorrelation of tomorrow's conventional return on a stock with today's SPDR return up to the SPDR trade immediately preceding today's last trade of the stock. This eliminates NT, and greatly reduces whatever small BAB effect exists in portfolio returns. Since TVRP can be ignored in our setting, PPA is the only remaining source of autocorrelation. If the autocorrelation is statistically significantly different from zero, it can only come from PPA.

Measure Autocorrelation over Disjoint Time-Horizon Subperiods to Obtain Independence for Statistical Power.—Binomial models cannot be used to aggregate tests of autocorrelation of individual stocks in a single data period, because returns are correlated across stocks; thus, the sample autocorrelations of individual stock returns are correlated across stocks. However, under the assumption that the theoretical autocorrelation is zero in all subperiods, the sample return autocorrelations are independent across disjoint time periods.[8] Consequently, if we divide our data period into disjoint time-horizon subperiods, then tests within the various subperiods *can* be aggregated using binomial models or other methods that require independence:

- Within each time-horizon subperiod, compute the individual return autocorrelations of each stock in a group, then average across stocks in the group and test whether the average return is statistically significant. Count the number of subperiods in which the average is statistically significant, and test this using a binomial model.
- Within each time-horizon subperiod, the number of stocks with statistically significant individual return autocorrelation is a nonnegative random variable with known mean. The probability that

this number exceeds any given level can be bounded by the Markov Inequality. Compute the number in each subperiod, then compute the order statistics of the subperiod counts. The probability that the first (smallest) order statistic exceeds a certain level, and the probability that the second (second smallest) order statistic exceeds a given level, can be determined from the binomial distribution.

- Within each time-horizon subperiod, compute the sample portfolio return autocorrelation, and test whether it is statistically significant. Count the number of subperiods in which the sample autocorrelation is statistically significant, and test this using a binomial model.

Compute the Proportion of the Autocorrelation Clearly Attributable to PPA.—Using the three key ideas above, we can identify a portion of the autocorrelation that can only be attributed to PPA, a portion that cannot be PPA, and a portion that may or may not be PPA. Take the portion that can only be attributed to PPA, and compute the residual (total autocorrelation less the portion that can only be attributed to PPA; equivalently, the sum of the portion that cannot be PPA and the portion that may or may not be PPA). Compute the absolute value of the portion that can only be attributed to PPA, divided by the sum of the absolute values of the portion that can only be attributed to PPA and the absolute value of the residual. This gives a lower bound on the portion of the autocorrelation arising from PPA.

Dataset and Findings.—We examine ten years (1993-2002) of transaction data of stocks listed on the NYSE. As noted above, the specific tests we use are based on combinations of the four key ideas we have just explained. The following two subsections outline our tests and findings, for individual stock returns and portfolio returns.

Individual Stock Returns.—NT and BAB both produce 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; see Chan (1993).

Our analysis focuses on the five two-year time-horizon subperiods of our ten-year data period 1993-2002. Under the hypothesis that the theoretical stock return autocorrelation is zero, the *average* (within each size group) of the individual *conventional* sample stock return autocorrelations in one time-horizon subperiod is independent of the average in the other time-horizon subperiods. We are interested in measuring autocorrelation arising from PPA. The presence of negative NT and BAB means that the sample autocorrelation is a downward-biased estimator of the autocorrelation arising from PPA, so we do a one-sided test of the hypothesis that the average autocorrelation is zero in each subperiod. The hypothesis is rejected at the 1% level for small, medium and large firms. We conclude that PPA is an important source of the autocorrelation in stock returns among firms of all sizes.

Under the assumption that the theoretical conventional return autocorrelation of each individual stock is zero in every time-horizon subperiod, the sample conventional return autocorrelations of the stocks are independent across disjoint time-horizon. We test the hypothesis that the conventional daily return autocorrelation of *each* individual stock is zero in each subperiod. As above, the sample autocorrelation is a downward-biased estimator of the autocorrelation arising from PPA, so we do a one-sided test. We find that this hypothesis is rejected at the 1% level among small and (depending on which test of autocorrelation we use) at the 1% or 5% level for medium firms. We conclude that PPA is present in these firms, and that its effect is larger than the combined effect of NT and BAB.

In addition to studying the autocorrelation of conventional daily stock returns, we study the autocorrelation of the *open-to-close* returns. The use of open-to-close returns eliminates NT and greatly reduces BAB, so BAB and TVRP are small enough to be ignored in our autocorrelation tests. The use of open-to-close returns also eliminates some of the PPA in stock return autocorrelation, but provides us with a direct measure of the remaining portion of the autocorrelation. Since there is no bias of consequence arising from NT, BAB or TVRP, we do two-sided tests of the hypothesis that the open-to-close return is zero.

Breaking our data period into five two-year time-horizon subperiods, we find that the *average* open-to-close return autocorrelation is positive and overwhelmingly significant among small and medium firms; this is strong evidence of the existence of PPA as a source of autocorrelation, and that it is positive on average. For large firms, we find strong evidence that the contribution of PPA to autocorrelation is negative in at least some subperiods, with the overall sign of PPA varying between periods. The variation in the sign of PPA most plausibly reflects changes in the ratio of informed to momentum traders, with positive autocorrelation when informed traders predominate, and negative autocorrelation when momentum traders predominate.

We reject at the 1% level for small and medium firms, for all three correlation tests, the hypothesis that the open-to-close return autocorrelation of *each* firm is zero. This is strong evidence that PPA is an important source of stock return autocorrelation among small and medium firms. The rejections come overwhelmingly from positive autocorrelations; indeed, the number of stocks with negative open-to-close return autocorrelation is systematically below the expected value. We conclude that PPA is systematically positive for small and medium stocks. As with the average open-to-close 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 PPA: 56.2% for small firms, 60.7% for medium firms, and 52.6% for large firms.

Portfolio Returns.—BAB produces slightly negative autocorrelation in daily portfolio returns, while NT produces positive autocorrelation; PPA can be either positive or negative. Thus, the finding that daily portfolio returns exhibit positive autocorrelation is not sufficient to establish that PPA is a source of the autocorrelation. A number of papers have carried out indirect tests that tend to support the role of PPA in autocorrelation, but the results have been controversial because of the indirect nature of the tests. In this paper, we conduct two direct tests of the role of PPA in explaining the positive autocorrelation in daily portfolio returns.

In the first test, we define the open-to-close return of a portfolio as the equally-weighted average of the open-to-close returns of the individual stocks in the portfolio.[9] Since the open-to-close return of a portfolio on a given day depends only on trades that occur that day, NT 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 open-to-close portfolio returns is positive and highly significant for small and medium firms, providing strong evidence that PPA is a major source of the autocorrelation. The autocorrelation of open-to-close portfolio returns is negative and not significant for large firms.

For the second test, note that the return autocorrelation of a portfolio of 100 stocks is the average of 10,000 autocorrelations: 9,900 cross-autocorrelations and 100 own-autocorrelations. Thus, the portfolio return autocorrelation is essentially the average of the return cross-autocorrelations of the pairs of stocks. In the second test, we compute the cross-correlation 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, NT is eliminated, and BAB is greatly reduced. Our main null hypothesis is that the cross-correlation is equal to zero for every stock. This hypothesis is strongly rejected for all three portfolios and all three (Pearson, the Andrews modification of Pearson, and Kendall tau) correlation tests.

Our method provides a lower bound on the proportion of portfolio return autocorrelation arising from PPA: 54.6% (small firms), 59.5% (medium firms), and 36.8% (large firms).

The remainder of this paper is organized as follows. Section 2 reviews the literature on daily return autocorrelation. Section 3 details our methodology and null hypotheses. Section 4 describes the sampling of firms and provides descriptive statistics of our data. Section 5 presents and interprets the empirical results. Section 6 provides a summary of our results and some suggestions for further research.

2. Literature review

In this section, we review the literature on daily stock return autocorrelation. There has been considerable controversy over the proportion of the autocorrelation that should be attributed to each of the four components: NT, BAB, PPA and TVRP. In part, the controversy has arisen because previous tests have depended on particular market microstructure models, and thus were joint tests of those models as well as of the source of the autocorrelation; we refer to such tests as indirect. There have been direct tests of the speed of price adjustment and PPA, but not of the role of PPA in stock return autocorrelation.[10]

Since Fisher (1966) and Scholes and Williams (1977) first pointed out NT, the extent to which it can explain autocorrelation has been extensively studied, but remains very controversial. Atchison, Butler, and Simonds (1987) and Lo and MacKinlay (1990) find that NT explains only a small part of the portfolio autocorrelation (16% for daily autocorrelation in Atchison, Butler, and Simonds (1987), 0.07, a small part of the total autocorrelation, for weekly autocorrelation in Lo and MacKinlay (1990)). Bernhardt and Davies (2005) find that the impact of NT on portfolio return autocorrelation is negligible. However, Boudoukh, Richardson, and Whitelaw (1994) find that the weekly autocorrelation attributed to NT in a portfolio of small stocks is up to 0.20 (56% of the total autocorrelation) when the standard assumptions by Lo and MacKinlay (1990) are loosened by considering heterogeneous nontrading probabilities and heterogeneous betas;[11] they 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 (2002) 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."

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 cross-country survey.

Chan (1993) provides a model which addresses individual stock return autocorrelation and cross-autocorrelation. In Chan's model, 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. Chan 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. 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 inconsistent with one prediction of his model, namely that daily return autocorrelation in each firm is zero; see also the test of our Null Hypothesis I, below.

Chordia and Swaminathan (2000) compare portfolios of large, actively traded stocks, to portfolios of smaller, thinly traded stocks, arguing that NT 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.[12]

There has even been controversy over whether the autocorrelation still exists: Chordia, Roll, and Subrahmanyam (2005) write: "Daily returns for stocks listed on the New York Stock Exchange (NYSE) are not serially correlated."[13]

3. Methodology

As noted by Lo and MacKinlay (1990), NT arises from measurement error in calculating stock returns. If an individual stock does not trade on a given day, its daily return is reported as zero;[14] 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 (negative) drift component, the equilibrium mean return, plus a daily mean-zero volatility term, 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 (above) 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, which might be several hours 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, NT causes spurious negative individual autocorrelation and positive individual cross-autocorrelation, resulting in positive autocorrelation of portfolios. The first key idea in this paper is to theoretically sign and/or bound the various sources of autocorrelation, so that we may draw inferences about the source from the sign of the observed autocorrelation.

The second key 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, NT is, *by definition*, eliminated; if the correlation turns out to be nonzero, there must be a source, other than NT, 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 NT, our method of calculating correlations greatly reduces BAB, and we can ignore it as a source of stock return autocorrelation.

We say that a stock exhibits PPA 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,…,I) 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 PPA in stock j implies the following:[15]

> given times $s < t \leq u < v$ such that stock j trades at some time $w \in [t,u]$, r_{uvj} 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 PPA by examining return correlations over time intervals [s,t] and [u,v] satisfying the condition just given.

Two of our tests focus on what we call open-to-close returns; in these tests, NT is eliminated, and BAB is greatly reduced. The open-to-close 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 open-to-close return of stock i on day d is $r_{s_i t_i i}$, where s_i and t_i are the times of the first and last trades of the stock on day d.[16] We compute the correlation $\rho(r_{s_i t_i i}, r_{u_i v_i i})$, where u_i and v_i are the times of the first and last trades on day d+1. Note that $s_i < t_i < u_i < v_i$, so NT is eliminated. BAB arises in conventional daily return autocorrelation because the correlation considered is $\rho(r_{q_i t_i i}, r_{t_i v_i i})$, where q_i is the time of the last trade prior to day d. Note that the end time in calculating $r_{q_i t_i i}$ is the same as the starting time in calculating $r_{t_i v_i i}$, 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 open-to-close 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 BAB 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 open-to-close stock returns are zero. Relaxing the independence assumption results in slightly negative autocorrelation and cross-autocorrelation of open-to-close returns.[17] In this paper, we assume that BAB does not contribute to autocorrelation in open-to-close returns of individual stocks or portfolios.[18] This seems completely innocuous in the context of portfolio returns, since the consensus is that BAB plays no significant role in the autocorrelation of conventional portfolio returns, and its role in open-to-close portfolios would be even smaller. For individual stock open-to-close returns, note that the relevant coin tosses are those for the last trade one day and the first trade the next day. A lot happens overnight: a considerable amount of information comes in from news stories, corporate and governmental information releases, and foreign markets. Limit orders can be set to expire at the close of trade one day, allowing the trader to place new limit orders the next day, taking any new information into account. It seems to us that the overnight information flow amounts to a thorough randomization that should pretty much eliminate correlation in the value of the two coin tosses used in our analysis.[19] However, a reader who is still concerned by the assumption that open-to-close individual stock return BAB is zero should note that BAB will result in negative bias in our autocorrelation estimates. It could thus possibly invalidate our negative autocorrelation and two-sided autocorrelation tests. However, it makes it harder to find statistical significance in our positive autocorrelation tests. Our statistical significance comes from the positive autocorrelation tests; if we simply ignore the negative and two-sided tests, none of our main conclusions is affected; see Table 5. This issue is discussed in the results section, for the particular null hypotheses where it arises.

Each correlation is tested by the three methods outlined below.

The third key idea is to divide the data period into disjoint time-horizon subperiods, and note that under the assumption that the theoretical autocorrelation is zero, the sample return autocorrelations within disjoint time-horizon subperiods are independent, so we may derive tests using the binomial distribution. This idea is applied in two settings. In the first setting, we compute a single sample autocorrelation in each subperiod; in one case, we compute the average of individual stock return sample autocorrelations over each subperiod, while in another case, we compute the sample portfolio return autocorrelation over each subperiod. We count the number of time-horizon subperiods in which the single autocorrelation value is statistically significant, and use the binomial distribution to test for significance overall. The probability of a one-sided rejection in any given subperiod using a symmetric 5% rejection criterion is 2.5%, so the probability of obtaining two rejections in five subperiods is $(5!/(2!3!))(.025)^2 = 0.00625$, leading to overall rejection at the 1% level; the probability of a two-sided rejection in any given subperiod using a symmetric 5% rejection criterion is 5%, so the probability of obtaining three rejections in five subperiods is $(5!/(2!3!))(.05)^3 = 0.00125$, leading to overall rejection at the 1% level.

In the second setting, we apply the binomial distribution to counts of stocks with statistically significant return autocorrelations in each of the time-horizon subperiods. Depending on the test, we reject the hypothesis for an individual firm in a given time-horizon subperiod using either a symmetric 5% rejection criterion or a one-sided 2.5% rejection criterion. If the correlation tests were independent across firms, the number of rejections would have the binomial distribution. If the collection $\{ r_{s_i t_i i} : i = 1, \ldots, I \}$ 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 $\sqrt{(.05)(.95)I}$ ($\sqrt{(.025)(.975)I}$). Since returns are not independent across stocks, X will not be binomial. The standard deviation of X is not readily ascertainable, and is likely higher than that of the binomial. However, the failure of independence does not change the mean of X, so X is a nonnegative, integer-valued, random variable with mean $0.05I$ ($0.025I$).

In all of these tests, there are I=100 firms, so X has mean μ=5 or μ=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-horizon subperiods. This provides us with n independent observations of X; let $X_1, \ldots, X_n$ be the order statistics, i.e., X_1 is the smallest observation, X_2 the second smallest, and so forth. Then using the binomial distribution, 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. The test for the k^{th} order statistic X_k involves the combinatorial coefficient $n!/((k\text{-}1)!(n\text{-}k+1)!)$ as well as the factor $(\mu/x_k)^{n-k+1}$, both of which grow rapidly with k. Thus, the power of the test for X_k declines rapidly with k, suggesting the test be based on X_1 alone. However, the test for X_1 can 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. For these reasons, we adopt a combined test using both X_1 and X_2, and not using the higher order statistics. Compute the statistic $p_3 = 2\, min\{p_1,p_2\}$. Note that for any γ, $P(p_3 \leq \gamma) = P(2\, min\{p_1,p_2\} \leq \gamma) = P(p_1 \leq \gamma/2$ 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.

We analyzed ten years' worth of transaction data. There is a trade-off between the number of time-horizon subperiods and the lengths of the time-horizon subperiods. Because stock returns are very noisy, for a given return period, it is much easier to detect autocorrelation in longer time-horizon subperiods than in shorter time-horizon subperiods. On the other hand, the statistical power of the order statistic tests increases when the number of independent observations (the number of time-horizon subperiods) increases. 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; on the whole, the results are qualitatively similar.

3.1. *Individual stock returns*

Previous studies of individual stock return autocorrelation have focused on the average autocorrelation of groups of firms, finding it to be statistically insignificant and usually positive (Säfvenblad (2000)). This finding does not rule out the possibility that 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. This is the focus of our analysis, because it allows us to test whether the autocorrelation arises from PPA; as a comparison to the previous literature, we also compute the average autocorrelation over groups of firms, segregated by firm size. We calculate the autocorrelation in two different ways: the conventional daily return autocorrelation, and the open-to-close return autocorrelation.

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 day prior to day *d* on which trade occurs. When we compute individual stock returns in the conventional way, NT and BAB are both present, and both generate negative autocorrelation. Null Hypothesis I is that the *average* daily return autocorrelation is zero in each firm group in each of our five two-year time-horizon subperiods; we test this hypothesis by comparing the average sample daily return autocorrelation for each subperiod to the associated standard error. As a result of NT and BAB, the sample return autocorrelation is a downward-biased estimator of the autocorrelation arising from PPA; thus, we use a one-sided test of Null Hypothesis I. Rejection of Null Hypothesis I implies that PPA contributes to conventional daily stock return autocorrelation, and that it is positive in at least some subperiods.

Because the average sample autocorrelations are independent across disjoint time-horizon subperiods, the number of subperiods on which a hypothesis is rejected has the binomial distribution. In each subperiod, we use a one-sided rejection criterion at the 2.5% level. As noted above, we reject Null Hypothesis I at the 1% level if the average daily return autocorrelation is statistically positive in two or more of the five subperiods.

We also compute the average autocorrelation over the whole ten-year period as a weighted average of the averages in the five subperiods, weighting each period by the inverse square of the associated standard error.

We report the averages and standard errors for the individual time-horizon subperiods, as well as the weighted average and standard error over the whole ten-year period.

Null Hypothesis II is that *every* firm's conventional daily return exhibits zero autocorrelation. For each firm, we test whether daily returns exhibit zero autocorrelation, in each of n=5 disjoint two-year time-horizon subperiods. Because the sample return autocorrelation of each stock is a downward-biased estimator of the return autocorrelation arising from PPA, we use a one-sided test. For each firm, we test whether daily returns exhibit nonpositive autocorrelation, in each of the five disjoint two-year subperiods. In each subperiod, we test whether the sample autocorrelation lies above the 97.5 %-ile, so μ=2.5.[20] 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 II implies that in at least some firms, the contribution of PPA to conventional return autocorrelation is positive and is larger than the sum of the NT and BAB effects.

Open-to-Close Return Autocorrelation.—As above, we define the open-to-close 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 dataset.[21] If we compare open-to-close returns on day *d* and day *d*+1, there is *no* NT effect: the open-to-close 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*, BAB is sufficiently reduced so that it can be ignored; see the extended comments on this point above. If PPA makes no contribution to stock return autocorrelation, the theoretical autocorrelation of open-to-close returns on each stock must be zero. Since the average sample open-to-close return autocorrelation is an unbiased estimator of the open-to-close return autocorrelation arising from PPA, we use two-sided tests. Null Hypothesis III is that the *average* autocorrelation of open-to-close returns is zero in each of the five two-year time-horizon subperiods in each group of stocks. In each subperiod, we test whether the subperiod autocorrelation is significantly nonzero at the 5% level by comparing the autocorrelation to its standard error. We reject Null Hypothesis III at the 1% level (p=0.00125) if the autocorrelation is statistically different from zero in three or more of the five subperiods. Rejection of Null Hypothesis III implies that PPA contributes to stock return autocorrelation. As in the case of conventional returns, we also compute the average autocorrelations over the entire ten-year period.

Our Null Hypothesis IV is that the autocorrelation of open-to-close returns on *each* stock is zero in each two-year subperiod. Because the sample open-to-close return autocorrelation is an unbiased estimator of the open-to-close return autocorrelation arising from PPA, we use a two-sided test, but also report positive and negative rejections separately. We divide our data period into five two-year time-horizon subperiods, and test using $p_3(\mu=5,n=5)$; where we consider positive and negative rejections separately, we report $p_3(\mu=2.5,n=5)$. Rejection of Null Hypothesis IV implies that PPA is a source of individual stock return autocorrelation. Where PPA is positive, it arises from slow incorporation of information into prices more than from overshooting due to positive-feedback strategies.

Analysis of Autocovariance.—The fourth key idea in this paper is to use the decomposition of autocorrelation into its various components to estimate the fraction of the autocorrelation arising from PPA. In this section, we describe a method to obtain a lower bound on the portion of the individual stock autocovariance attributable to PPA. 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 our data period 24 hours per day, 7 days per week. However, the open-to-close returns of the stocks are calculated over a portion of the data 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. A portion of the period when the markets are open, and the entire period when the markets are closed, are omitted. In all conventional models of stock pricing, the standard deviation of open-to-close return should be lower than the standard deviation of conventional daily return. For example, if the stock price is any Itô Process, the realizations of the volatility term over the excluded intervals are uncorrelated with the realizations 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 open-to-close 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 1,500 stocks (300 stocks per time-horizon subperiod times five subperiods) exhibits sample standard deviation of open-to-close returns greater than the sample standard deviation of conventional daily return.

For each stock, we can compute the conventional daily (open-to-close) return autocovariance by taking the product of the conventional daily (open-to-close) return autocorrelation times the conventional daily (open-to-close) 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 PPA is the only source of the open-to-close return autocovariance. If C_i and I_i denote the conventional daily and open-to-close 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 open-to-close returns. This ratio is a lower bound on the portion of the identifiable return autocorrelation attributable to PPA. It understates the proportion of the autocorrelation attributable to PPA for two reasons. First, PPA can induce both negative and positive effects; these cancel, and we see only the net effect in this calculation. Second, PPA 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.

3.2. Portfolio returns

While many papers have studied whether NT 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 NT. 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 NT and BAB are the sole explanations of portfolio return autocorrelation, the autocorrelation computed by our methods must be less than or equal to zero.

First Method, Open-to-Close Returns.—In the first method, we compute the open-to-close returns of each individual stock as defined in section 3.1. As noted there, open-to-close returns on different days do not exhibit NT, and BAB should be essentially eliminated. We consider three portfolios, each containing 100 stocks, representing small, medium, and large market capitalization.

We define the open-to-close return of a portfolio on a given day as the equally-weighted average of the open-to-close 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 open-to-close return of the portfolio is just the average of the correlations of the open-to-close 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 own autocorrelations of the stocks in the portfolio.

If PPA makes no contribution to portfolio return autocorrelation, the theoretical autocorrelation of the open-to-close return of the portfolio must be zero. Thus, our Null Hypothesis V is that the autocorrelation of the open-to-close return of the portfolio is zero in each of the five two-year subperiods. Since the sample open-to-close portfolio return autocorrelation is an unbiased estimator of the contribution of PPA to portfolio return autocorrelation, we use a two-sided test. Null Hypothesis V is rejected at the 1% level ($p = 0.00125$) if the portfolio return autocorrelation is statistically nonzero in three of the five subperiods. Rejection of Null Hypothesis V implies that there is a nonzero PPA contribution to portfolio return autocorrelation; the sign of the PPA effect is determined by the sign of the autocorrelation. As in the case of average individual stock return autocorrelations, we test the portfolio return autocorrelation separately in each of our five two-year time-horizon 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 open-to-close return of the portfolio allows us to obtain a lower bound on the portion of the conventional daily return autocorrelation attributable to PPA. As in section 3.1, all conventional models of stock pricing predict that the variance of open-to-close 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 time-horizon subperiods in our dataset. We calculate the autocovariance of conventional daily (open-to-close) portfolio returns by multiplying the conventional daily (open-to-close) autocorrelation of portfolio returns by the conventional daily (open-to-close) variance of portfolio returns. The residual is defined as the difference of the conventional and open-to-close autocovariances. The autocovariance of open-to-close portfolio returns can only come from PPA, so the ratio of the open-to-close autocovariance to the sum of the absolute values of the open-to-close and residual autocovariances gives a lower bound on the proportion of the autocorrelation that is attributable to PPA.

Second Method, ETFs.—In the second method, we take our portfolio to be 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 NT. 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. 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 PPA, 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. Since SPDRs are a single security, NT arises only from days on which no trade occurs; since SPDRs are traded extremely actively, NT makes no contribution to the SPDRs' return autocorrelation. As a single security priced on a grid, SPDRs in principle exhibit BAB; however, in our sample period, SPDRs exhibit insignificant daily return autocorrelation, which confirms that BAB is not a significant source of 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 conventional return of stock i on day $d+1$ (in other words, the return over the return period 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 return period 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.[22] Note that each time we compute a correlation, it is the correlation of a stock return over a given return period with the return of a traded security, SPDRs, over a disjoint return period, with both the SPDRs and the stock trading in the interval separating the two return periods. Thus, the calculation of the correlation does not use stale prices, and hence there is no NT effect. Since BAB is not a significant source of return autocorrelation for the SPDRs, since the stock is a different security from the SPDRs, and since the stock trades occur at different times from the SPDR trades, any BAB between the SPDRs and the individual stock will be virtually eliminated and can be ignored. Thus, in the absence of PPA, the correlation between the return of the individual stock and the return of the SPDRs must zero. Our Null Hypothesis VI is that the correlation of each of the individual stock returns and the return of the SPDRs is zero in each of the five two-year time-horizon subperiods. Because the sample return autocorrelation is an unbiased estimator of the contribution of PPA to the autocorrelation, we use two-sided tests, but report positive and negative rejections separately. We test using $p_3(\mu=5,n=5)$ for Null Hypothesis VI, and $p_3(\mu=2.5,n=5)$ when considering the positive and negative rejections separately. Rejection of Null Hypothesis VI implies that the PPA contributes to portfolio return autocorrelation.

In our test, the only detectable source of PPA 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 PPA present in the market, is not captured by these tests.

Ex ante, this seems to us an unlikely place to search for PPA. PPA 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. Other studies of the incorporation of public information into securities prices, and their relationship to this paper, are discussed in footnote 8. The two studies most closely related to our SPDR tests are Hasbrouck (1996, 2003). Hasbrouck (2003) showed that index futures lead SPDRs in incorporating new information. Hasbrouck (1996) showed that order flows for large individual stocks generated by index futures-based program trading and index arbitrage have price impacts beyond the index futures themselves, but that these are very quickly incorporated; as indicated in his Figure 1, the impact is substantially achieved within two minutes, and completely achieved within three minutes. Indeed, the two-minute effect is quite close to the permanent effect, and the three-minute effect overshoots. Since index futures lead SPDRs, individual stocks should lag SPDRs by a shorter interval.[23] These relatively short lags seem to us unlikely to result in daily return autocorrelation, for two reasons. First, stock prices are very volatile and it seems likely that the autocorrelation resulting from a three-minute lag would be swamped by the variability over the other 387 minutes of the trading day. Second, the R^2 of overall market returns on the returns of individual stocks is quite low. Since the correlation of SPDRs with individual stocks induced by the lag will be the product of the information revealed in the lag times the R^2, the effect should be very small indeed. To test whether the lags documented by Hasbrouck (1996, 2003) can explain our findings, we rerun our tests for large firms, calculating the SPDR return over a return period ending three minutes before the last trade of the stock.

Because the definitions underlying Null Hypothesis VI 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:

$S_{(d,h)i}$ = Price of stock *i* at hour *h* on date *d*,

$\overline{S}_{(d,h)}$ = Price of SPDRs at hour *h* on date *d*,

$h(d,i)$ = Hour of last trade of stock *i* on date *d*,

$S_{di} = S_{(d,h(d,i))i}$ (the closing price),

$\overline{S}_d = \overline{S}_{d,4pm}$,

$$r_{di} = \frac{S_{di} - S_{(d-1)i}}{S_{(d-1)i}},$$

$$\overline{r}_d = \frac{\overline{S}_d - \overline{S}_{d-1}}{\overline{S}_{d-1}},$$

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, $\frac{\overline{S}_d - \overline{S}_{d-1}}{\overline{S}_{d-1}}$, into two components, $\frac{\overline{S}_d - \overline{S}_{d,h(d,i))}}{\overline{S}_{d-1}}$ and $\frac{\overline{S}_{d,h(d,i)} - \overline{S}_{d-1}}{\overline{S}_{d-1}}$. No stale prices are used in the calculation of $Corr\left(r_{(d+1)i}, \frac{\overline{S}_{d,h(d,i)} - \overline{S}_{d-1}}{\overline{S}_{d-1}}\right)$, 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 return periods separated by trades, as we can see in Fig. 1, so there is no NT effect; for details, see the Appendix. For the reasons given above, BAB can be ignored. In the absence of PPA, the covariance must be zero.

$$Corr\left(r_{(d+1)i}, \frac{\overline{S}_{d,h(d,i)} - \overline{S}_{d-1}}{\overline{S}_{d-1}}\right) = 0. \quad (1)$$

<Insert Fig. 1>

3.3. *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 the Kendall 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, assuming the variables have a bivariate normal distribution. Letting r_p be the Pearson sample correlation coefficient, the *t*-test statistics are

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

- ***Modified Pearson correlation test*:** We use the Andrews modification of the Pearson test, taking into account the possibility that the error terms exhibit heteroskedasticity or autocorrelation. We use the Andrews (1991) 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 tau test* (*a nonparametric test*):** Stochastic volatility biases the standard errors in the Pearson correlation test. The Kendall tau test is a nonparametric test that makes no assumptions on the joint distribution of the variables, and is completely immune to the effects of stochastic volatility. Kendall's sample rank correlation coefficient is

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

where $K = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} Q\left((X_i, Y_i), (X_j, Y_j)\right)$ and $Q((a,b),(c,d)) = \begin{cases} 1, & if\,(c-a)(d-b) > 0 \\ -1, & if\,(c-a)(d-b) < 0 \end{cases}$. Then the Kendall 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 $n > 10$.[24]

4. Data

Our data period covers the ten calendar years 1993-2002. We divide this into five two-year time-horizon 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.

Because our data period contains the market bubble that burst in 2000, one might be concerned about "regression to the max" (Ross (1987), Brown, Goetzmann, and Ross (1995)). These papers point out that if researchers deliberately focus on periods immediately before and after a big local maximum or minimum in securities prices, a sample selection bias arises in the computation of return autocorrelation. For example, the paths of Geometric Brownian Motion exhibit substantial local maxes and mins. If one takes a long history of Geometric Brownian Motion, then computes the autocorrelation only in the time-horizon subperiods in which a large local max or min occurred, one is likely to find positive autocorrelation that appears statistically significant in those subperiods, as a result of the bias. However, if one selects the data period without regard to local maxes and mins, there is no bias.

We selected our ten-year data period based on the data available at the time we began the study, and did not deliberately choose to analyze the period including the bubble.[25] Of course, the inadvertent inclusion of the bubble might have meant that our results would not extend to other periods. Could the inclusion of the bubble have affected our results? No. All of our autocorrelations are computed separately in each of the two-year time-horizon subperiods 1993-94, 1995-96, 1997-98, 1999-2000, 2001-02. Only one of those subperiods, 1999-2000, contains a local max or min that stands out in comparison to other time periods. As a robustness check, we recomputed our *p*-values using only the other four subperiods. The *p*-values rise somewhat because less data is being used, but virtually all of our results retain statistical significance.

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 Prices (CRSP) tapes for the subperiod; the trade data comes exclusively from TAQ, while CRSP data is used solely 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 time-horizon 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.[26]
- We remove firms for which no transaction occurs for 30 or more consecutive trading days within the two-year time-horizon subperiod.
- We remove firms whose shares traded for less than $5 at any point in the two-year time-horizon 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 time-horizon subperiod; only one stock (Berkshire Hathaway) was removed.
- 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 time-horizon subperiod. The large-firm sample consists of the 100 largest firms not eliminated by the criteria above; 97 (average for the five two-year time-horizon 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 time-horizon 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.[27] 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 dataset. Thus, for each individual stock, we compute the closing price as the price at the last transaction occurring *before* 4:00 p.m.[28]

Table 1 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 time-horizon 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 1 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 1 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 6.

The closing trade for an individual stock is the last trade occurring before 4:00 p.m. In order to ensure that the closing trade of the SPDRs occurs after the closing trade of each stock, we take the closing trade of the SPDRs to be the first trade occurring after 4:00 p.m., except on the 26 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.[29]

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 NT, as we see in portfolios; however, SPDRs are subject to BAB.

<Insert Table 1>

5. Empirical results and their implications

5.1. Individual stock returns

Conventional Daily Return Autocorrelation.—Table 2 reports results on the conventional daily return autocorrelation of individual firms, stratified into the three groups. Null Hypothesis I, that the average conventional daily return autocorrelation is zero in each of the five time-horizon subperiods, is rejected at the 1% level among small, medium, and large firms. As noted in section 3.1, we reject Null Hypothesis I at the 1% level if the autocorrelation is statistically positive at the 2.5% (one-sided) level in two or more of the five subperiods. The fact that medium and large firms exhibit insignificant or negative autocorrelation in some subperiods suggests that the relative importance of PPA, NT, and BAB varies somewhat from period to period. Since both the NT and BAB effects result in negative autocorrelation, this finding provides compelling evidence that PPA is the main source of daily return autocorrelation among small firms, and an important source of daily return autocorrelation among 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. Chan does not test Null Hypothesis I, but his results imply its rejection.

<Insert Table 2>

Table 3 reports the results for Null Hypothesis II, that each firm's daily return autocorrelation is zero. Recall that we use a one-sided test, since the sample return autocorrelation is a downward-biased estimator of the role of PPA in daily return autocorrelation, due to NT and BAB. Null Hypothesis II is rejected at the 1% level, for all three correlation tests, for small firms; it is rejected at the 5% level for two of the three correlation tests, and at the 1% level for the remaining test, for medium firms. The strong rejection of Null Hypothesis II shows that the NT and BAB effects cannot be the main source of autocorrelation in small and medium firms, and that PPA must play an important role.

<Insert Table 3>

Open-to-Close Return Autocorrelation.—Table 4 reports results on average open-to-close return autocorrelation of individual firms, stratified into the three groups. Recall that we present results on averages mainly to link our analysis to previous work; we consider the individual tests in Table 5 to be our main tests for individual stock open-to-close autocorrelation. As noted above, open-to-close returns are calculated so as to eliminate NT and greatly reduce BAB, so that PPA is the only plausible source of autocorrelation and we use two-sided correlation tests. Averages are reported for each of the five two-year time-horizon subperiods, along with an average for the whole ten-year period obtained as a weighted average of the subperiod averages.

Null Hypothesis III is overwhelmingly rejected among small and medium firms,[30] and is rejected at the 1% level (p = 0.00125) among large firms, indicating that PPA plays a significant role in stock return autocorrelation. The average PPA-induced autocorrelation is significant and positive in all five subperiods for small firms and four out of five subperiods for medium firms. The average PPA-induced autocorrelation among large firms is significant and negative in two of the subperiods, positive and significant in a third subperiod. The negative autocorrelation presumably arises from overshooting due to momentum traders, while the positive autocorrelation presumably arises from slow price adjustment as informed traders exercise their informational advantage; the overall sign of the autocorrelation is presumably determined by the relative numbers of informed and momentum traders. If the reader remains skeptical of our assumption that open-to-close individual stock returns are free of BAB, note that the statistical significance for small and medium firms comes from the positive autocorrelation tests, which do not depend on the assumption; only the results for large firms could potentially be compromised.

<Insert Table 4>

Table 5 reports our results on individual open-to-close stock returns. Null Hypothesis IV, that each firm's open-to-close 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. This shows that PPA is a significant source of individual daily return autocorrelation among small and medium firms. There are far more positive rejections than negative rejections among small and medium firms, indicating that the contribution of PPA to individual stock return autocorrelation is generally positive, indicating that 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 lack of statistical significance of negative autocorrelations supports our assumption that BAB is not a factor in individual stock open-to-close returns. If the reader remains skeptical of that assumption, note that only the negative and two-sided autocorrelation tests could possibly be compromised; if we simply ignore those tests, we still obtain the same statistical significance from the positive autocorrelation tests.

The numbers of positive and negative rejections among large firms vary significantly among time-horizon subperiods, 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 in section 3.1 to provide a lower bound of the identifiable absolute autocovariance of individual stocks arising from PPA. Over the five two-year time-horizon 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 PPA.

<Insert Table 5>

Summary: Individual Stock Return Autocorrelation.—PPA must be an important source of the autocorrelation of daily returns of individual stocks, among small and medium firms. PPA is systematically positive among small and medium firms, 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 PPA is also an important source of autocorrelation of daily returns, but the sign of PPA varies among time-horizon subperiods, possibly as a result of variation in the popularity of momentum strategies among traders of large stocks.

5.2. Portfolio returns

First Method, Open-to-Close Returns.—Tables 6 and 7 report our results concerning conventional and open-to-close portfolio returns. Results are presented for each of the five two-year time-horizon subperiods, along with a value for the whole ten-year period computed as a weighted average of the subperiod returns. Table 6 presents the results for conventional portfolio returns. Since the sample autocorrelation of conventional portfolio returns is an upward-biased estimator of the PPA-induced autocorrelation, the correct test would be a one-sided test with rejection if the sample autocorrelation is sufficiently negative. Since the effect of PPA on autocorrelation is usually positive, rejection seemed unlikely and we did not formalize a Null Hypothesis. Table 6 confirms that there are no one-sided negative rejections. Many of the positive results would be significant in a two-sided test, indicating that the combined contribution of PPA and NT to the autocorrelation is positive. As the firm size becomes larger, the first-order autocorrelation of portfolio return becomes smaller. This is consistent with the previous studies (e.g., Chordia and Swaminathan (2000, Table I on page 917)).

Table 7 presents the results for open-to-close portfolio returns, and our tests of Null Hypothesis V. The results are statistically significant and positive at the 1% level in four of the five data subperiods, indicating overwhelming rejection for small and medium firms. Moreover, over the whole ten-year data period, the open-to-close 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 PPA 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 open-to-close portfolio return is not significant and slightly negative. Thus, our open-to-close portfolio returns do not provide evidence of PPA in portfolios of large stocks; however, as we shall see, we do find such evidence in our test involving SPDRs.

Table 8 shows the autocovariances of conventional (open-to-close) returns in each of the three portfolios, obtained by multiplying the conventional (open-to-close) return autocorrelations by the conventional (open-to-close) variances. The ratio of the open-to-close autocovariance to the sum of the absolute values of the open-to-close and residual autocovariances ranges from 44.8% to 65.5% over the five time-horizon subperiods, with an average of 54.6%, for small firms; from 5.0% to 90.8%, with an average of 59.5%, for medium firms, and from 18.4% to 64.0%, with an average of 36.8%, for large firms. As noted above, these figures represent lower bounds of the portion of the autocorrelation attributable to PPA.

<Insert Table 6>

<Insert Table 7>

<Insert Table 8>

Second Method, ETFs.—Table 9 shows the results of our tests of Null Hypothesis VI, that the correlation of individual stock and SPDRs returns is zero. As explained above, these correlations are calculated in a way that eliminates NT and virtually eliminates BAB, so we do two-sided tests, but also report the positive and negative rejections separately, including the *p*-values for the associated one-sided tests. For small, medium and large firms, the positive rejections significantly outnumber the negative rejections, so we obtain greater statistical significance from the one-sided tests. For all three correlation tests, Hypothesis VI is rejected in a one-sided positive test at the 1% level for small and medium firms, and at the 5% level for large firms. Hypothesis VI is rejected in the two-sided test at either the 5% or 1% level for small, medium and large firms and for all three correlation tests, with one exception.[31] The rejection of Null Hypothesis VI provides strong evidence of PPA among small, medium, and large firms.

It is important to note also that Null Hypothesis VI is not rejected by any of the one-sided negative tests. When we look at the five two-year time-horizon subperiods, a 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. However, in some subperiods the negative rejections outnumber the positive rejections, indicating again that the correlation pattern between the SPDRs and firms varies somewhat over time, perhaps with the prevalence of program trading to arbitrage discrepancies between the prices of the ETF and its underlying stocks.

<Insert Table 9>

As noted above, Hasbrouck (1996, 2003) has documented that prices of ETFs lead prices of the constituent stocks, but the lags are small, not more than three minutes. To test whether these short lags could be explaining our results for large stocks, we reran the tests of Null Hypothesis VI, computing the SPDR return up through the last SPDR trade at least three minutes before the closing trade of the stock. The results are reported in Table 10. The counts of large stocks with statistically significant correlation change very little, and the *p*-values either stay the same or increase very slightly. Null Hypothesis VI is still rejected at the 5% level for all three two-sided tests. For the positive one-sided tests, it is rejected at the 5% level by the modified Pearson and Kendall tau tests, and just misses (p = 0.0625) for the Pearson test. We conclude that the lags documented by Hasbrouck cannot explain our findings.

Our finding of PPA for large firms in the SPDRs tests contrasts with the failure to find such evidence for large firms in the portfolio open-to-close return tests. In the SPDR tests, the time interval between the final trade of a stock and the last previous SPDR trade is just over three minutes. By contrast, in the portfolio open-to-close 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, and it appears that large stock prices do adjust to information overnight.

<Insert Table 10>

6. Concluding remarks

Daily stock return autocorrelation is one of the most visible stylized facts in empirical finance. While the price discovery literature has clearly demonstrated the existence of PPA in the varying speeds of adjustment across different assets, there is no consensus in the previous literature on the relative contributions of NT, BAB, PPA and TVRP to daily return autocorrelation.

We find compelling evidence that PPA is an important source, and in some cases the main source, of stock return autocorrelation. PPA is an important source of autocorrelation in all of our tests involving small and medium firms, and in some tests involving large firms. Our tests cover both individual stock return autocorrelation and portfolio return autocorrelation.

Our methods for decomposing stock and portfolio return autocorrelation into NT, BAB, PPA, and TVRP are generally applicable. These methods make use of four key ideas: theoretically signing and/or bounding these components; computing returns over disjoint return subperiods separated by a trade to eliminate NT; subdividing our data period into time-horizon subperiods to obtain independent autocorrelation measures; and obtaining a lower bound on the portion of the autocorrelation attributable to PPA by isolating a portion of the autocorrelation that must be PPA.

We use two methods to eliminate NT. The first method computes correlations of open-to-close returns; this method can be applied to eliminate NT 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 NT using any security which, like the SPDRs, is traded nearly continuously.

By dividing our data period into disjoint time-horizon subperiods, we obtain independent tests of the sources of autocorrelation in the different time-horizon subperiods. Aggregating the tests across time-horizon subperiods allows us to increase the statistical power of our tests, and to work around the problem that returns are correlated across stocks. Our methods for aggregating the results of the time-horizon subperiod tests 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 PPA among portfolios in our test using SPDRs, but not our tests involving open-to-close returns. The use of open-to-close returns allows us to measure only a portion of the PPA; 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 PPA. Is there some other way to capture more of the PPA in a single test?
- our tests seem to indicate that PPA among large firms is positive in certain periods and negative in

other periods. We argue that this most plausibly reflects variations in the relative numbers of informed and momentum traders. Is there a way to test this?

- among small and medium firms, all our tests provide strong evidence of the importance of PPA in individual return autocorrelation. Among large firms, we find similar evidence in the test aggregating average autocorrelation tests in the time-horizon subperiods, but not from our analysis of individual firm return autocorrelations. The number of positive autocorrelations of individual large-firm open-to-close returns in the various time-horizon subperiods is usually substantially above the expected value, suggesting that PPA plays a role in return autocorrelation in this setting, but our test involving the first two order statistics is too weak to detect it, at least using five two-year or ten one-year time-horizon subperiods. Would a longer overall data period, and/or a different statistical test, allow one to demonstrate the role of PPA in stock return autocorrelation in this setting?
- risk management measures such as Value-at-Risk (VaR) are computed from the estimated unconditional correlation among assets. The presence of autocorrelation in stock returns suggests that these risk measures should be adjusted, conditional on the price movements of the previous day.

Appendix: derivation of equation (1)

As we can see in Fig. 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 NT 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 Fig. 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 NT induces spurious correlation. The contribution of the NT effect to the correlation between the SPDRs return and the stock return comes solely from the interval B in Fig. 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 Fig. 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 NT. BAB is too small to generate statistically significant autocorrelation. In the absence of PPA, the correlation must be zero, so Equation (1) holds.

References

Adams, G., McQueen, G., Wood, R., 2004. The effects of inflation news on high frequency stock returns. Journal of Business 77, 547-574.

Ahn, D.-H., Boudoukh, J., Richardson, M.P., Whitelaw, R.F., 2002. Partial adjustment or stale prices? implications from stock index and futures return autocorrelations. Review of Financial Studies 15, 655-689.

Anderson, R.M., 2006. Time-varying risk premia and stock return autocorrelation. Working paper, University of California at Berkeley, http://emlab.berkeley.edu/users/anderson/Time-VaryingRiskPremia.pdf.

Andrews, D.W.K., 1991. Heteroskedasticity and autoregression consistent covariance estimation. Econometrica 59, 817-858.

Atchison, M.D., Butler, K.C., Simonds, R.R., 1987. Nonsynchronous security trading and market index autocorrelation. Journal of Finance 42, 533-553.

Badrinath, S.G., Kale, J.R., Noe, T.H., 1995. Of shepherds, sheep, and the cross-autocorrelations in equity returns. Review of Financial Studies 8, 410-430.

Bernhardt, D., Davies, R.J., 2005. Portfolio cross-autocorrelation puzzles. Working paper, University of Illinois.

Bessembinder, H., 1997. The degree of price resolution and equity trading costs. Journal of Financial Economics 45, 9-34.

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

Brennan, M.J., Jegadeesh, N., Swaminathan, B., 1993. Investment analysis and the adjustment of stock prices to common information. Review of Financial Studies 6, 799-824.

Brown, S.J., Goetzmann, W.N., Ross, S.A., 1995. Survival. Journal of Finance 50, 853-873.

Busse, J.A., Green. T.C., 2002. Market efficiency in real time. Journal of Financial Economics 65, 415-437.

Campbell, J.Y., Grossman, S., Wang, J., 1993. Trading volume and serial correlation in stock returns. Quarterly Journal of Economics 108, 905-939.

Campbell, J.Y., Lo, A.W., MacKinlay, A.C., 1997. The Econometrics of Financial Markets. Princeton University Press, Princeton, NJ.

Chan, K., 1993. Imperfect information and cross-autocorrelation among stock prices. Journal of Finance, 1211-1230.

Choi, J.Y., Salandro, D., Shastri, K., 1988. On the estimation of bid-ask spreads: theory and evidence. Journal of Financial and Quantitative Analysis 23, 219-230.

Chordia, T., Sarkar, A., Subrahmanyam, A., 2007. The microstructure of cross-autocorrelation. Federal Reserve Bank of New York Staff Reports 303.

Chordia, T., Swaminathan, B., 2000. Trading volume and cross-autocorrelations in stock returns. Journal of Finance 55, 913-935.

Chordia, T., Roll, R., Subrahmanyam, A., 2005. Evidence on the speed of convergence to market efficiency. Journal of Financial Economics 76, 271-292.

Connolly, R., Stivers, C., 2003. Momentum and reversals in equity-index returns during periods of abnormal turnover and return dispersion. Journal of Finance 58, 1521-1555.

Ederington, L.H., Lee, J.H., 1995. The short-run dynamics of the price adjustment to new inflation. Journal of Financial and Quantitative Analysis 30, 117-134.

Fisher, L., 1966. Some new stock market indexes. Journal of Business 39, 191-225.

Garcia Blandon, J., 2001. New findings regarding return autocorrelation anomalies and the importance of non-trading periods. Working paper, Universitat Pompeu Fabra.

Hasbrouck, J., 1996. Order characteristics and stock price evolution: an application to program trading. Journal of Financial Economics 41, 129-149.

Hasbrouck, J., 2003. Intraday price formation in U.S. equity index markets. Journal of Finance 58, 2375-2399.

Hasbrouck, J., 2007. Empirical Market Microstructure: The Institutions, Economics, and Econometrics of Securities Trading. Oxford University Press, New York, NY.

Jegadeesh, N., Titman, S., 1993. Returns on buying winners and selling losers: implications for market efficiency. Journal of Finance 48, 64-91.

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

Kim, S.T., Lin, J.-C., Slovin, M.B., 1997. Market structure, informed trading and analysts' recommendations. Journal of Financial and Quantitative Analysis 32, 507-524.

Lee, C.M.C., Mucklow, B., Ready, M.J., 1993. Spreads, depths, and the impact of earnings information: an intraday analysis. Review of Financial Studies 6, 345-374.

Llorente, G., Michaely, R., Saar, G., Wang. J., 2002. Dynamic volume-return relation of individual stocks. Review of Financial Studies 15, 1005-1047.

Lo, A.W., MacKinlay, A.C., 1990. An econometric analysis of nonsynchronous trading. Journal of Econometrics 45, 181-212.

Lyons, R.K., 2001. The Microstructure Approach to Exchange Rates. MIT Press, Cambridge, MA.

Madhavan, A., Smidt, S., 1993. An analysis of changes in specialist inventories and quotations. Journal of Finance 48, 1595-1628.

McQueen, G., Pinegar, M., Thorley. S., 1996. Delayed reaction to good news and cross-autocorrelation of portfolio returns. Journal of Finance 51, 889-920.

Mech, T.S., 1993. Portfolio return autocorrelation. Journal of Financial Economics 34, 307-344.

Roll, R., 1984. A simple implicit measure of the effective bid-ask spread in an efficient market. Journal of Finance 39, 1127-1140.

Ross, S.A., 1987. Regression to the max. Working paper, Yale University.

Säfvenblad, P., 2000. Trading volume and autocorrelation: empirical evidence from the Stockholm Stock Exchange. Journal of Banking and Finance 24, 1275-1287.

Scholes, M.S., Williams, J., 1977. Estimating betas from nonsynchronous data. Journal of Financial Economics 5, 309-328.

Sheskin, D.J., 1997. Handbook of Parametric and Nonparametric Statistical Procedures. CRC Press, Boca Raton, FL.

Footnotes

[1] See pages 44-55, 65-78, 84-107 and 128-136 of Campbell, Lo, and MacKinlay (1997).

[2] The momentum effect has been cited as an explanation of medium-term (3 to 12 months) autocorrelation (see Jegadeesh and Titman (1993)). The momentum effect is properly viewed as a form of PPA. We make no attempt in this paper to model PPA, and thus need not be concerned with the various forms of trader behavior that can give rise to it. Rather, we present methods to decompose return autocorrelation into the various components. In addition, the medium-term momentum effect is of little relevance to daily return autocorrelation, which is the focus of the empirical work reported here.

[3] On pages 84-85, Campbell, Lo, and MacKinlay (1997) write "For example, suppose that the returns to stocks A and B are temporally independent but A trades less frequently than B. … Of course, A will respond to this information eventually, but the fact that it responds with a lag induces spurious cross-autocorrelation between the daily returns of A and B when calculated with *closing* prices. This lagged response will also induce spurious own-autocorrelation in the daily returns of A" [emphasis in original]. On page 100, they write "Moreover, as random buys and sells arrive at the market, prices can bounce back and forth between the ask and bid prices, creating spurious volatility and serial correlation in returns, even if the economic value of the security is unchanged."

[4] We do not view the term "spurious" as pejorative in any sense.

[5] 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 NT, PPA, and TVRP.

[6] Inventory costs should lead specialists to make transitory adjustments in prices in order to bring their inventories back to the desired level (see Lyons (2001), pages 130-133). Because the mean reversion of specialists' inventories has quite a long half-life (Madhavan and Smidt (1993)), it seems unlikely that inventory costs are a significant source of *daily* return autocorrelation. Although the autocorrelation resulting from inventory costs is often described as a microstructure *bias*, we see it as a form of PPA. The specialist deliberately adjusts the stock price above (below) the price that equates traders' buy and sell orders in order to increase (reduce) his inventory to the desired level, then gradually lowers (raises) the price, even though the inventory imbalance conveyed no information. The autocorrelation results from the slow decay of the adjustment.

[7] However, as noted in Anderson (2006), the application of the methods of this paper to monthly, quarterly, or annual return periods, or to time horizons substantially longer than two years, may require adjustment for TVRP.

[8] We do not claim that returns are independent across disjoint time periods, only that if the theoretical autocorrelation is zero in each time period, the sample autocorrelations are independent across disjoint time periods. We use three autocorrelation tests, each of which depends on certain assumptions. The Kendall tau test is nonparametric and is completely unaffected by observed forms of return dependence, such as volatility clustering. The Andrews test depends on semiparametric assumptions; if our returns data are generated by a GARCH process, for example, then the assumptions underlying the Andrews test are satisfied. The Pearson test depends on strong parametric assumptions; if our returns data are generated by the returns of a geometric Brownian motion with drift, for example, then the assumptions underlying the Pearson test are satisfied. If the assumptions underlying a given autocorrelation test are satisfied, and the theoretical autocorrelation is zero in each time horizon subperiod, then the

sample autocorrelations will in fact be independent across the disjoint subperiods. Because eliminating NT is a main focus of our work, we need to use data sampled at random times determined by trades, reflecting strategic decisions by traders; we are not aware of any econometric time series model which covers this case, but we can see no plausible reason why independence across the five subperiods would be violated.

[9] Garcia Blandon (2001) examines the return autocorrelation of the IBEX-35, an index composed of the 35 most liquid Spanish companies. He computes returns on an open-to-close basis. It appears he takes the opening price to be the index value when the market opens, rather than the average of the opening prices of the stocks comprising the index; since some of the stocks in the index will not trade at the market opening, the opening price of the index will involve some stale prices, so NT will not be completely eliminated. Garcia Blandon (2001) finds that the autocorrelation disappears when the index returns are computed on an open-to-close basis; this is analogous to our finding for large firms, but contrasts with our finding for small and medium firms.

[10] For example, Ederington and Lee (1995), Busse and Green (2002), and Adams, McQueen, and Wood (2004) established by direct tests that the incorporation of publicly-released information into securities' prices is not instantaneous. Ederington and Lee (1995) found lags of not more than 40 seconds in interest rate and foreign exchange futures in response to scheduled macroeconomic news releases; since trade occurs over this 40 second interval, they found PPA. Kim, Lin, and Slovin (1997) studied price adjustment when favored clients are given access to an analyst's initial buy recommendation prior to the opening of the market: "For NYSE/AMEX stocks, almost all of the private information contained in analysts' recommendations is reflected in the opening trade. There are only minor gains to informed traders subsequent to the initial trade, gains that are typically less than transaction costs. … However, [on the NYSE and AMEX] considerable time elapses, on average 10 minutes, before the specialist sets a market clearing price. Thus, in the specialist market, there is a substantial period during which no market clearing price is found, so there is a loss in liquidity." Note that Kim, Lin, and Slovin (1997) found a delay in price adjustment, but did *not* find PPA on the NYSE and AMEX, since no trade occurs in the period before the specialist sets the price; they did find PPA on the Nasdaq. Busse and Green (2002) found that the prices of NYSE, AMEX, and Nasdaq stocks positively mentioned on the CNBC TV Morning Call and Midday Call segments adjust within one minute; traders who execute within 15 seconds make small but significant profits, indicating that PPA is occurring but for a very brief period. Adams, McQueen, and Wood (2004) studied the response of portfolios of stocks (sorted by size) to inflation surprises. The return of each portfolio was calculated on horizons of one, five, and fifteen minutes, one hour, and one day; if a stock does not trade within a given horizon, it is excluded from the calculation of return on its portfolio. They found substantial responses within one minute and,

indeed, the coefficients at the one-minute horizon exceed those at the one-day horizon in five of their ten regressions (five portfolios times two measures of inflation). The coefficients at the fifteen-minute horizon are even higher, which Adams, McQueen, and Wood (2004) interpreted as implying that the full response takes somewhat more than one minute, but less than fifteen minutes. To us, this indicates that some overshooting is occurring at the fifteen-minute horizon. We acknowledge, as Adams, McQueen, and Wood (2004) pointed out, that the statistical significance is generally lower at the one-day horizon, because additional time allows for additional volatility that raises the standard errors. However, the coefficients at the one-day horizon are still unbiased estimators of the permanent effect of the inflation surprise, so that if the one-minute coefficients roughly agree with the one-day coefficients, while the fifteen-minute coefficients are noticeably higher than the one-day coefficients, this says to us that prices reach the correct level within one minute, then proceed to overshoot over the next fourteen minutes; overshooting should result in negative autocorrelation. Adams, McQueen, and Wood (2004) estimated that prices of large firms complete their adjustment (which we believe involves some overshooting) within six trades, and within a smaller number of trades for small firms. Adams, McQueen, and Wood (2004) conclude that "differential response times to inflation news cannot be driving daily, weekly and monthly cross-autocorrelation." Thus, we are not aware of any previous evidence that the slow incorporation of publicly-released information is a factor in daily return autocorrelation. When information is released publicly, there is no opportunity for informed agents to benefit by trading slowly so as to conceal their information, so we expect price adjustment to be rapid. By contrast, when private information is possessed by some agents and *not* released publicly, market microstructure models predict that informed agents may strategically choose to exercise their informational advantage slowly, over several days, and this slow price adjustment has the potential to generate daily return autocorrelation. Because the private information of traders is generally not observable, one cannot usually apply the methods of Ederington and Lee (1995); Kim, Lin, and Slovin (1997); Busse and Green (2002); and Adams, McQueen, and Wood (2004) to the incorporation of private information into prices. In this paper, we devise direct tests of the relative roles of NT, BAB, PPA, and TVRP in stock return autocorrelation.

[11] 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.

[12] 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 (2003) differs from that of Chordia and Swaminathan (2000), the results from Connolly and Stivers (2003) are similar to those of Chordia and Swaminathan (2000); and Llorente, Michaely, Saar, and Wang (2002); supporting the partial price adjustment hypothesis. For other literature on the PPA hypothesis, see for example Brennan, Jegadeesh, and Swaminathan (1993), Mech (1993), Badrinath, Kale, and Noe (1995), McQueen, Pingar, and Thorly (1996).

[13] From the context, it is clear they are talking about portfolio return autocorrelation, and they note that "the S&P 500 index is virtually a random walk over a horizon of one day." However, Chordia, Sarkar, and Subrahmanyam (2007) "examines the mechanism through which the incorporation of information into prices leads to cross-autocorrelation in stock returns," finding that price discovery in large firms leads price discovery in small firms.

[14] The CRSP dataset reports the average of the final bid and ask quotes as the "closing price" so that returns calculated from CRSP data will generally not be zero on no-trade days. However, because our primary focus is separating PPA from NT, we need to use intraday transaction data, and thus we use the NYSE TAQ dataset, rather than CRSP data.

[15] More precisely, the absence of PPA and TVRP imply the stated conclusion. See Anderson (2006) for a detailed analysis of the magnitude of the potential bias resulting from TVRP.

[16] Since there is no trade in the stock after time t_i, the open-to-close return also equals $r_{s_i ti}$, where t is the closing time.

[17] 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 BAB. Positive (negative) serial correlation in the coin tosses of a given stock induces negative (positive) autocorrelation of open-to-close 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 BAB to induce slight negative autocorrelation of individual stock open-to-close 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 BAB 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 open-to-close returns.

[18] An alternative method for reducing BAB would be to use the midpoint of the closing buy and sell quotes as the closing price. As has been noted in the literature, computing returns from the midpoint of the bid and ask quotes reduces BAB, but need not completely eliminate it. For example, Hasbrouck (2007, page 91 notes "Finally, trade prices are subject to bid-ask bounce, a source of transitory volatility. Although *quote midpoints are not immune from transitory components*, their short-run volatility is lower than that of trades" [emphasis added]. The most important issue for us is separating nonsynchronous trading (NT) from partial price adjustment (PPA). PPA is characterized by *trades* occurring at prices that do not fully reflect the information available. If we were to compute returns using the midpoint of the bid and ask quotes, rather than actual trades, then our tests would not establish the role of PPA in daily return autocorrelation. At best, they would establish a role for a different notion—that the bid and ask quotes do not fully reflect the information available—in daily return autocorrelation. In summary, there are two ways to control for BAB. Neither is perfect. Using the midpoint of the bid and ask quotes is probably better for the specific purpose of controlling for BAB, but is not helpful for our main goal, separating NT from PPA. We chose to use a method which clearly separates NT from PPA, and which is helpful for reducing BAB.

[19] While there may be positive serial correlation in the indicator in successive trades on a single day, we are not aware of any paper that finds positive serial correlation between the indicator on the last trade one day and the first trade the next day. We do not find statistical significance in our negative autocorrelation tests for individual stocks in Table 5; this supports our belief that any residual BAB is too small to matter in our individual stock open-to-close return autocorrelation tests.

[20] The specific tests for autocorrelation will be described in section 3.3.

[21] The reader might have expected us to set the open-to-close return of that stock to be zero for that day. Doing so could introduce an NT 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.

[22] 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.

[23] Note that this does not imply that the stock will trade within two or three minutes of a move of the SPDRs; rather, it indicates that any trade occurring at least three minutes after the SPDR move should be at the correct price and should not exhibit PPA.

[24] See Sheskin (1997), page 633.

[25] The analysis began in the fall of 2004. SPDRs were not introduced until 1993, so we could not go back before 1993. We did not have access to complete data for 2004 at that time, so we could not use six two-year subperiods 1993-2004. Thus, we chose to use five two-year subperiods 1993-2002. At the time, we paid no attention to the presence of the bubble in our dataset.

[26] 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.

[27] 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.

[28] As noted above, this differs from the CRSP dataset, which reports the midpoint of the closing bid and ask quotes as the closing price.

[29] 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.

[30] Small (medium) firms exhibit five (four) rejections out of five at the 0.5% (one-sided) level; the probability of five out of five rejections is $(.005)^5<10^{-11}$, while the probability of four out of five rejections is $5(.005)^4(.995)<10^{-8}$.

[31] The exception is the Kendall tau test of small firms, which barely misses significance at the 5% level (*p*=.0715).

Table 1
Descriptive statistics of data

We use transaction data from TAQ database over the sample period from January 4, 1993 to December 31, 2002, drawing separate samples of firms in each two-year subperiod. 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. ¶ denotes the statistics of portfolios, not the average of individual firms of each group. † SPDRs were introduced on the AMEX on January 29, so SPDRs have 19 fewer trading days compared than other stocks.

	Small firm	Medium firm	Large firm	SPDRs
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 *portfolio* 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 2
Average *individual* daily return autocorrelations: conventional daily returns
(Null Hypothesis I)

Numbers in parenthesis are standard errors. Because the average sample return autocorrelation is a downward-biased estimator of the autocorrelation arising from PPA, we use a one-sided test of the averages, rejecting if they are sufficiently positive. In each time-horizon subperiod, ** and * denote positive significance at the 0.5% and 2.5% (one-sided) levels, respectively. Null Hypothesis I is rejected at the 1% level if the autocorrelation is statistically positive in two or more of the five subperiods; thus, Null Hypothesis I is rejected among all three firm-size groups.

Portfolio	Number of firms	Average *individual* daily returns autocorrelations: conventional daily returns					
		93-94	95-96	97-98	99-00	01-02	1993-2002
Small firm	100	0.0084	0.0034	0.0664**	0.0412**	0.0330**	0.0335**
		(0.0102)	(0.0112)	(0.0093)	(0.0085)	(0.0098)	(0.0043)
Med. firm	100	0.0173*	-0.0058	0.0208*	0.0119	-0.0354	-0.0009
		(0.0082)	(0.0076)	(0.0097)	(0.0080)	(0.0075)	(0.0036)
Large firm	100	0.0115*	-0.0011	-0.0082	0.0217**	-0.0076	0.0026
		(0.0058)	(0.0058)	(0.0057)	(0.0061)	(0.0055)	(0.0026)

Table 3
Autocorrelation of daily *individual* stock returns: conventional daily returns
(Null Hypothesis II)

Because the sample return autocorrelation is a downward-biased estimator of the return autocorrelation arising from PPA, we do a one-sided test on each stock, rejecting if the sample return autocorrelation is sufficiently positive. + denotes the numbers of stocks with statistically significant positive autocorrelation, at the one-sided 2.5% level. X_1 denotes the first order statistic (minimum) of the observations for the five two time-horizon 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 + 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 *individual* 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***
Med. firm: +	19	12	22	16	4	4	0.0954	12	0.0078	*0.0157**
Large firm: +	7	4	6	15	1	1	1.0000	4	0.3815	*0.7629*
Panel B: modified Pearson correlation test										
Small firm: +	18	16	35	23	23	16	0.0001	18	0.0017	*0.0002***
Med. firm: +	19	8	19	13	3	3	0.4019	8	0.0358	*0.0072***
Large firm: +	7	5	5	11	1	1	1.0000	5	0.1875	*0.3750*
Panel C: Kendall tau test										
Small firm: +	9	16	36	26	26	9	0.0017	16	0.0026	*0.0033***
Med. firm: +	20	10	23	15	3	3	0.4019	10	0.0156	*0.0313**
Large firm: +	5	4	5	6	5	4	0.0954	5	0.1875	*0.1907*

Table 4
Average *individual* daily return autocorrelations: open-to-close returns
(Null Hypothesis III)

Because the open-to-close return sample autocorrelation is an unbiased estimator of the open-to-close autocorrelation induced by PPA, we do two-sided tests. Numbers in parenthesis are standard errors. Within each of the five two-year time-horizon subperiods, ** and * denote positive significance at the 1% and 5% (two-sided) levels, respectively; ++ and + denote negative significance. Null Hypotheses III is rejected at the 1% level (p = 0.00125) if the autocorrelation is statistically nonzero in three or more of the five two-year subperiods; thus, Null Hypothesis III is overwhelmingly rejected for all three firm size groups. PPA generates positive return autocorrelation in all five subperiods among small firms, and in four of five subperiods among medium firms.

Portfolio	Number of firms	Average *individual* daily returns autocorrelations: open-to-close returns					
		93-94	95-96	97-98	99-00	01-02	1993-2002
Small firm	100	0.0389**	0.0394**	0.0729**	0.0441**	0.0240**	0.0435**
		(0.0074)	(0.0084)	(0.0090)	(0.0083)	(0.0092)	(0.0038)
Med. firm	100	0.0543**	0.0321**	0.0326**	0.0265**	-0.0376++	0.0218**
		(0.0071)	(0.0065)	(0.0096)	(0.0067)	(0.0073)	(0.0032)
Large firm	100	0.0064	-0.0133+	-0.0324++	0.0126*	-0.0103	-0.0064++
		(0.0062)	(0.0063)	(0.0066)	(0.0058)	(0.0054)	(0.0027)

Table 5
Autocorrelation of daily *individual* stock returns: open-to-close returns
(Null Hypothesis IV)

Because the open-to-close return sample autocorrelation is an unbiased estimator of the open-to-close autocorrelation induced by PPA, we do two-sided tests, but also report positive and negative rejections separately. + and - denote the numbers of stocks with statistically significant positive and negative autocorrelation, at the one-sided 2.5% level; +- denotes the numbers of stocks with statistically significant autocorrelation at the two-sided 5% level. X_1 denotes the first order statistic (minimum) of the observations for the five two time-horizon 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, while those in +- are nonnegative random variables with expectation 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 *individual* stock returns: open-to-close 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 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 6

First-order autocorrelation of daily *portfolio* returns: conventional daily returns

Since NT is positive, the sample conventional portfolio return autocorrelation is an upward-biased estimator of the contribution of PPA to portfolio return autocorrelation, so we use a one-sided test, rejecting if the sample autocorrelation is sufficiently negative. There are no significant negative results. Many of the positive results would be significant in a two-sided test, indicating that the combined contribution of PPA and NT to the autocorrelation is positive.

Portfolio	Number of firms	First-order autocorrelation of daily *portfolio* 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 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 7
First-order autocorrelation of daily *portfolio* returns: open-to-close returns
(Null Hypothesis V)

Because the open-to-close return sample autocorrelation is an unbiased estimator of the open-to-close autocorrelation induced by PPA, we do two-sided tests. Numbers in parenthesis are standard errors. ** and * denote positive significance at the 1% and 5% level, respectively; $^+$ denotes negative significance at the 5% level. Null Hypotheses V is rejected at the 1% level (p = 0.00125) if the autocorrelation is statistically nonzero in three or more of the five two-year subperiods; thus, Null Hypothesis V is overwhelmingly rejected for small and medium firms, but not for large firms. PPA generates positive portfolio return autocorrelation in four of the five subperiods among small and medium firms.

Portfolio	Number of firms	First-order autocorrelation of daily *portfolio* returns: open-to-close 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 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 8
Proportion of PPA in the autocorrelation of the portfolio returns

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
open-to-close 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
open-to-close 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
open-to-close 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
open-to-close 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
open-to-close 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
open-to-close 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
open-to-close return (%)	0.0319	0.0267	0.0892	0.0605	0.0476
residual (%)	0.0363	0.0208	0.0604	0.0318	0.0587
open-to-close return autocovariance as percentage of open-to-close 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
open-to-close return (%)	0.0282	0.0320	0.1389	0.1262	-0.0013
residual (%)	0.0225	0.0084	0.0685	0.0127	0.0252
open-to-close return autocovariance as percentage of open-to-close 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
open-to-close return (%)	-0.0144	-0.0100	-0.0748	0.0798	0.0075
residual (%)	0.0277	0.0445	0.0798	0.0499	0.0318
open-to-close return autocovariance as percentage of conventional daily return autocovariance	34.16	18.41	48.38	64.02	19.14

Table 9
ETFs
(Null Hypothesis VI)

Because the open-to-close return sample autocorrelation is an unbiased estimator of the open-to-close autocorrelation induced by PPA, we do two-sided tests, but report positive and negative rejections separately. + 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 time-horizon 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, while the values in +- are nonnegative random variables with expectation 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 *individual* 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	*0.0251**
	-	2	1	23	25	6	1	1.0000	2	1.0000	*1.0000*
	+-	13	25	35	31	15	13	0.0084	15	0.0453	*0.0168**
Panel C: Kendall tau test											
Small firm:	+	16	20	41	9	16	9	0.0017	16	0.0026	*0.0033***
	-	0	0	1	0	14	0	1.0000	0	1.0000	*1.0000*
	+-	16	20	42	9	30	9	0.0529	16	0.0358	*0.0715*
Med. firm:	+	33	24	49	14	6	6	0.0126	14	0.0044	*0.0087***
	-	0	1	5	7	17	0	1.0000	1	1.0000	*1.0000*
	+-	33	25	54	21	23	21	0.0008	23	0.0092	*0.0015***
Large firm:	+	9	19	26	7	10	7	0.0058	9	0.0232	*0.0116**
	-	4	1	15	25	10	1	1.0000	4	0.3815	*0.7629*
	+-	13	20	41	32	20	13	0.0084	20	0.0156	*0.0168**

Table 10
ETFs: large firms only, lag at least three minutes
(Null Hypothesis VI)

This Table repeats the analysis of Table 9 for large firms, but requires the SPDR trade to be at least 3 minutes prior to the closing trade of the stock. Because the open-to-close return sample autocorrelation is an unbiased estimator of the open-to-close autocorrelation induced by PPA, we do two-sided tests, but report positive and negative rejections separately. + 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 time-horizon 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, while the values in +- are nonnegative random variables with expectation 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 *individual* 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											
Large firm:	+	9	22	12	5	8	5	0.0313	8	0.0358	*0.0625*
	-	3	1	22	26	6	1	1.0000	2	1.0000	*1.0000*
	+-	12	23	34	31	14	12	0.0126	14	0.0581	*0.0251**
Panel B: modified Pearson correlation test											
Large firm:	+	9	22	13	6	9	6	0.0126	9	0.0232	*0.0251**
	-	3	1	23	25	6	1	1.0000	3	0.8038	*1.0000*
	+-	12	23	36	31	15	12	0.0126	15	0.0453	*0.0251**
Panel C: Kendall tau test											
Large firm:	+	8	18	26	7	10	7	0.0058	8	0.0358	*0.0116**
	-	5	1	15	25	10	1	1.0000	5	0.1875	*0.3750*
	+-	13	19	41	32	20	13	0.0084	19	0.0189	*0.0168**

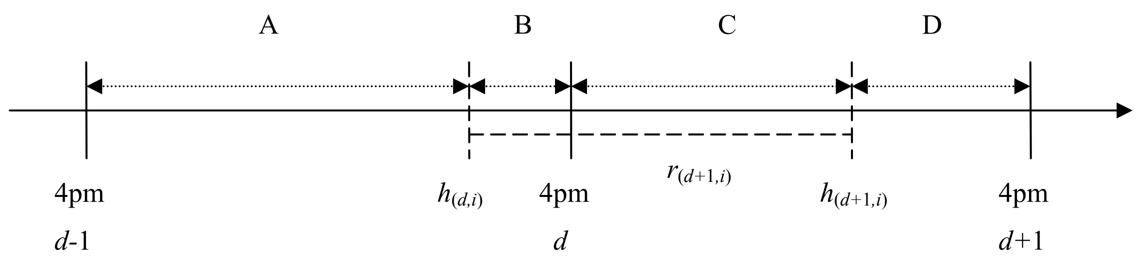

Fig. 1. Time diagram for Null Hypothesis VI

Our Null Hypothesis VI is that the correlation of each of the individual stock returns and the return of the SPDRs is zero. $r_{(d+1,i)}$ is the daily return of each individual stock on day d+1, computed in the conventional way. 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.