A Tale of Three Schools: Insights on Autocorrelations of Short-Horizon Stock Returns

Jacob Boudoukh; Matthew P. Richardson; Robert F. Whitelaw

A Tale of Three Schools: Insights on Autocorrelations of Short-Horizon Stock Returns

Jacob Boudoukh
New York University

Matthew P. Richardson
University of Pennsylvania

Robert F. Whitelaw
New York University

This article reexamines the autocorrelation patterns of short-horizon stock returns. We document empirical results which imply that these autocorrelations have been overstated in the existing literature. Based on several new insights, we provide support for a market efficiency-based explanation of the evidence. Our analysis suggests that institutional factors are the most likely source of the autocorrelation patterns.

There is strong support for the view that short-horizon stock returns are predictable.[1] This view has grown

[1] See, for example, Fama (1965), Gibbons and Ferson (1985), Lo and MacKinlay (1988, 1990a), Conrad and Kaul (1988, 1989), and Conrad, Kaul, and Nimalendran (1991).

in stature with two seemingly important pieces of evidence: (*a*) short-horizon portfolio returns are significantly autocorrelated [e.g., .36 for small stocks (see Table 1)], and (*b*) short-horizon returns are highly cross-serially correlated [e.g., .28 between small and lagged large-firm portfolios (see Table 1)]. That these autocorrelations and cross-serial correlations are large and statistically significant is irrefutable. However, the economic meaning of these short-horizon correlations is less clear and has been debated in the recent finance literature.

We can broadly divide the prevailing views on the meaning of these correlations into three schools of thought. The first school, the *loyalists*, believes that markets rationally process information. Their view is that large autocorrelations at short horizons are not due to fundamentals; instead, they argue that the correlations arise from market frictions. Specifically, both the pattern and the magnitude of the correlations are consistent with measurement error in the data (e.g., nonsynchronous trading, price discreteness, or bid-ask spreads), institutional structures (e.g., trading mechanisms such as different market structures or trading/nontrading periods), or microstructure effects (e.g., systematic changes in either inventory holdings or the flow of information).[2]

Similar to the loyalists, the second school of thought, the *revisionists*, believes that markets are efficient. However, they believe that, even in frictionless markets, short-horizon stock returns can be autocorrelated. Specifically, their view is that the correlation patterns are consistent with time-varying economic risk premiums. Changing risk premiums, they argue, can be explained by intertemporal asset pricing models, such as conditional versions of the arbitrage pricing theory or the consumption-based asset pricing model. That is, variation in risk factors, such as past market returns, past size returns, or interest rate spreads, can induce variation in short-horizon risk premiums.[3]

The third school of thought, the *heretics*, takes a different approach. They believe that markets are not rational, that profitable trading strategies do exist (even on a risk-adjusted basis), and that psycho-

---

[2] As examples of this literature: see Cohen et al. (1986) and Lo and MacKinlay (1990b) for a direct analysis of nonsynchronous trading; see Conrad, Gultekin, and Kaul (1992) for the potential effect of bid-ask spreads; see Keim (1989) for an investigation (albeit in a different setting) of the relation between transaction costs and market seasonals; for a look at returns around trading and nontrading periods, see Bessembinder and Hertzel (1993); for information-related microstructure effects, see Hasbrouck (1991), or, in a related way, for the potential effect of strategic trading, see Admati and Pfleiderer (1989); and for an analysis of transactions costs and portfolio autocorrelations, see Mech (1993).

[3] As examples of this literature: see Keim and Stambaugh (1986) and Lo and MacKinlay (1992) for monthly returns using linear factor models; see Conrad and Kaul (1989) for weekly returns in univariate settings; and for an analysis of weekly returns in a multivariate framework, see Connolly and Conrad (1991); Conrad, Gultekin, and Kaul (1991); Conrad, Kaul, and Nimalendran (1991); and Hameed (1992).

logical factors are important for pricing securities. For example, heretics argue that time series patterns in returns occur because investors either overreact or only partially adjust to information arriving to the market. Thus, for "astute" investors, excess profits can exist even if financial markets are well functioning.[4]

To understand the importance of the short-horizon autocorrelation evidence and its implications for expected returns, consider the expected return on the small-firm portfolio, conditional on the previous week's returns on the small- and large-firm portfolios both being positive versus both being negative. The annualized conditional expected returns on this portfolio are 60 percent and −44 percent, respectively, in these two states, a 104 percent difference in ex ante premiums! Although transaction costs may explain why this premium cannot be arbitraged away, they do not explain its existence. The aforementioned schools of thought all have opinions on the source of this premium, and this article's purpose is to provide some insights on these sources.

With respect to the debate between the loyalists, revisionists, and heretics, it is important to distinguish between correlation and economic causality. Therefore, we first provide an interesting perspective on the existing literature for short-horizon return autocorrelations. Specifically, Conrad, Gultekin, and Kaul (1991) reported rather striking evidence that seems to be at odds with existing interpretations of cross-serial correlation patterns documented by Hawawini (1980), Lo and MacKinlay (1990a), and Mech (1990). We argue that these cross-serial correlations can be explained by the portfolios' own autocorrelation patterns coupled with high contemporaneous correlations across portfolios [see also Hameed (1992)]. Given this point and the fact that small firms exhibit by far the most autocorrelation, understanding the dynamics of short-horizon returns is very much related to explaining the magnitude of small-firm portfolio return autocorrelations.

We provide an analysis of one potential explanation for the magnitude of this autocorrelation based on the nontrading and risk characteristics of small firms. As an alternative to Lo and MacKinlay's (1990b) homogeneous model of nonsynchronous trading, we develop different specifications of nontrading, which can produce potentially important effects. Using stylized facts on nontrading documented by

---

[4] At short horizons, several papers document *excess* profits from various trading rules based on positions in either individual securities or portfolios. As examples of this literature, see Jegadeesh (1990) and Lehmann (1990) for an analysis of contrarian strategies; see Jegadeesh and Titman (1992) and Lo and MacKinlay (1990a) for an analysis of profits due to cross-serial correlation across stocks; and see Brock, Lakonishok, and LeBaron (1992) for a detailed examination of technical trading strategies based on past movements in the Dow Jones 30. Some of these authors admit the possibility that these profits are compensation for risk.

Foerster and Keim (1993), we show that the effect of nonsynchronous trading has most probably been understated in the literature. For example, we consider portfolios with autocorrelations as low as 7 percent under standard assumptions. We show that these can be as high as 20 percent when the researcher takes account of the degree of heterogeneity within the portfolio in both the nontrading and the betas of individual stocks. This suggests that researchers should be especially cautious in ruling out nonsynchronous trading as an important determinant of the magnitude of autocorrelations in portfolio returns.

With the goal of providing an ex ante test of the three schools' implications, we study the relation between the autocorrelations of futures returns and the returns on the underlying spot index of two small-firms-weighted portfolios. We find that, although returns on small-firms-based indices display significant autocorrelation, returns on the corresponding futures contracts display almost none. Coupled with some diagnostic tests, this result points toward a loyalist explanation of the autocorrelation evidence.

The article is organized as follows. Section 1 explains, in the context of the existing evidence, why the cross-serial correlation pattern in portfolio returns is simply a restatement of existing autocorrelation patterns, coupled with high correlations across asset returns. In Section 2, we provide a closer look at the autocorrelation evidence. In Section 3, we focus on the issue of nonsynchronous trading and show that it can be more important than previously thought. Section 4 focuses on differentiating the three schools' implications for the time-series data by studying the relation between the spot index and corresponding futures markets. Section 5 makes some concluding remarks. Throughout the article, various data sources and corresponding specific procedures for forming portfolios are used. These are described in detail in the Appendix.

1. Another Look at the Cross-Serial Correlation Evidence: Why Only Autocorrelations Matter

Consider weekly returns on five size-sorted portfolios over the sample period 1962–1990. In a widely cited article, Lo and MacKinlay (1990a) found that the returns on large stocks led those on smaller stocks, but not vice versa. For example, over the 1962–1990 period, Table 1B provides the own- and cross-serial correlations between the returns on the five size portfolios.[5] As evident from this table, there is an

---

[5] Our autocorrelation calculations differ from those of Lo and MacKinlay (1990a) and other research in this area in two significant ways. First, we estimate the weekly autocorrelation day by day using overlapping observations, whereas they estimate the weekly autocorrelation on a particular day

asymmetric lead-lag relation in the cross-serial correlations across firms. Although not explicitly related to their analysis, Conrad, Gultekin, and Kaul (1991) and Conrad, Kaul, and Nimalendran (1991) reported an interesting stylized fact that bears on this evidence. They found that, in multiple regressions of small-firm portfolio returns on lagged returns, lagged large-firm returns have no predictive power beyond lagged small-firm portfolio returns [Conrad, Gultekin, and Kaul (1991, p. 609)].[6] This is in stark contrast to the popular lead-lag explanation of the cross-serial correlation evidence of Lo and MacKinlay (1990a) [i.e., that there is a delayed (possibly irrational) stock price reaction of smaller firms to information arriving to the market]. Absent the existing literature on lead-lag relations, what type of cross-serial correlation pattern should we expect in portfolio returns?

In the context of this evidence, Hameed (1992) showed that a time-varying factor model can explain the asymmetric cross-serial correlations in portfolio returns. His argument is that the asymmetry arises from differences in the level of time variation in expected returns across portfolios. So, for example, in comparing large firms to small firms, the key feature is how autocorrelated the expected return component is for small firms versus large firms. A more general way of looking at Hameed's (1992) point is to recognize that lagged returns on large firms are simply proxying for the small-firm portfolio's own lagged return, given the high degree of contemporaneous correlation across portfolios.

To see how asymmetric cross-serial correlation across assets can easily arise, consider a simple AR(1) model of the return-generating process for each size portfolio. For example, from the loyalist perspective, this process is implied by Lo and MacKinlay's (1990b) model of nonsynchronous trading. That is,

$$R_{it} = \alpha_i + \rho_i R_{it-1} + \epsilon_{it} \qquad \forall i, \tag{1}$$

where R_{it} is the return on size portfolio i from $t-1$ to t and ϵ_{it} is the unexpected shock to portfolio i over this period. Note that these shocks are contemporaneously correlated across the size portfolios. For example, Table 1A shows that the correlation between weekly

---

using nonoverlapping data. A priori, our procedure is more efficient [see Hansen and Hodrick (1980) and Richardson and Smith (1991)]. For example, heteroskedasticity aside, results in Richardson and Smith (1991) imply that the variances of the autocorrelation estimators using nonoverlapping data will be (asymptotically) 47 percent higher, which can be interpreted as the percentage number of additional observations needed to achieve the same power as the overlapping method. Second, although most researchers use a fixed portfolio throughout the sample, we resample quarterly [as in Lo and MacKinlay (1988)]. We refer to the differences between these approaches at various times in the article (also see the Appendix).

[6] We have performed some additional empirical tests of this result in nonlinear frameworks. The result seems quite robust to more general specifications.

Table 1
Lead-lag theoretical and implied correlations

A: Empirical contemporaneous correlations

	R_{st}	R_{2t}	R_{3t}	R_{4t}	R_{lt}	R_{ewt}
R_{st}	1.000	0.939	0.886	0.830	0.720	0.920
R_{2t}	0.939	1.000	0.972	0.937	0.849	0.983
R_{3t}	0.886	0.972	1.000	0.977	0.913	0.990 ·
R_{4t}	0.830	0.937	0.977	1.000	0.960	0.978
R_{lt}	0.720	0.849	0.913	0.960	1.000	0.920
R_{ewt}	0.920	0.983	0.990	0.978	0.920	1.000

B: Empirical cross-serial correlations

	R_{st-1}	R_{2t-1}	R_{3t-1}	R_{4t-1}	R_{lt-1}	R_{ewt-1}
R_{st}	0.362	0.357	0.346	0.328	0.284	0.352
R_{2t}	0.253	0.271	0.278	0.273	0.248	0.276
R_{4t}	0.189	0.209	0.219	0.223	0.210	0.219
R_{4t}	0.134	0.155	0.170	0.176	0.171	0.167
R_{lt}	0.033	0.049	0.063	0.069	0.070	0.058
R_{ewt}	0.210	0.225	0.231	0.230	0.210	0.231

C: Implied cross-serial correlations

	R_{st-1}	R_{2t-1}	R_{3t-1}	R_{4t-1}	R_{lt-1}	R_{ewt-1}
R_{st}	0.362	0.340	0.321	0.300	0.260	0.333
R_{2t}	0.254	0.271	0.263	0.253	0.230	0.266
R_{4t}	0.194	0.213	0.219	0.214	0.200	0.217
R_{4t}	0.146	0.165	0.172	0.176	0.169	0.173
R_{lt}	0.050	0.059	0.064	0.067	0.070	0.064
R_{ewt}	0.212	0.227	0.229	0.226	0.212	0.231

Tables 1A–1C provide estimated and implied autocorrelations of weekly stock returns of five size-sorted portfolios and the equal-weighted index over the 1962–1990 sample period. The estimates are calculated using overlapping weekly return data. Table 1A documents the contemporaneous cross-correlation across portfolio returns, Table 1B documents the cross-serial correlation across portfolio returns, and Table 1C documents the implied cross-serial correlations, using the contemporaneous cross-correlations and each portfolio's own autocorrelation of returns. Note that stocks are sorted from smallest to largest with corresponding returns R_{st}, R_{2t}, R_{3t}, R_{4t}, and R_{lt}. R_{ewt} is the return on the equal-weighted index.

returns on the size portfolios varies from .72 to .98 over the 1962–1990 subperiod.

This model therefore allows for contemporaneously correlated shocks to returns. However, consistent with Conrad, Gultekin, and Kaul (1991), this model implies that own lagged returns on a portfolio *completely* describe conditional expected returns on that portfolio. From Equation (1), each portfolio return has an infinite order moving average representation in terms of the disturbance terms:

$$R_{it} = \frac{\alpha_i}{1 - \rho_i} + \sum_{k=0}^{\infty} \rho_i^k \epsilon_{it-k} \qquad \forall i. \tag{2}$$

Using Equation (2), it is possible to calculate the first-order cross-serial correlation between return R_{it} and lagged return R_{jt-1} in terms

of R_{it}'s first-order autocorrelation and the contemporaneous correlation between R_{it} and R_{jt}. Specifically,[7]

$$\text{corr}(R_{it}, R_{jt-1}) = \text{corr}(R_{it}, R_{jt}) \times \text{corr}(R_{it}, R_{it-1}). \tag{3}$$

Table 1C provides Equation (3)'s implied values of the cross-serial correlations among the five size portfolios. Note that all that is required for this calculation are the relevant estimates of the autocorrelations and contemporaneous correlations across assets. Tables 1B and 1C show that the implied autocorrelations and the actual estimates appear quite close in value. For example, consider a representative portfolio, size quintile 3. Its implied cross-serial correlation values with respect to quintiles 1 through 5 and the market are .194, .213, .219, .214, and .200, as compared to the estimated correlations of .189, .209, .219, .223, and .210, respectively. This illustrates not only that the values are of equal magnitude, but also that the ordering in magnitudes across portfolios is similar.[8]

Thus, even in a world in which large-firm returns have no information beyond that contained in small firms, there can be large amounts of lagged cross-predictability. One way to interpret these results is that the lead-lag relation is a "red herring" with respect to the dynamics of short-horizon returns. That is, the lead-lag relation is simply another (and less efficient) way of describing the individual autocorrelation patterns of short-horizon portfolio returns. Since Table 1 implies that the small-firms portfolio's autocorrelation is especially important, it seems worthwhile trying to better understand the sources of this autocorrelation. With this in mind, we examine the autocorrelation more closely in the next section.

2. The Autocorrelation of Weekly Returns: Another Look

Our procedure for calculating the autocorrelation of weekly returns is based on the asymptotic arguments of Hansen and Hodrick (1980) and the analytical calculations of Richardson and Smith (1991), who showed that using all available data (i.e., in the presence of overlapping observations) is in general more efficient than its nonoverlapping counterpart. Previous researchers have chosen a particular day, say Wednesday, to estimate the weekly autocorrelations. What type of effects might this have in small samples?

A growing literature in finance documents patterns in daily auto-

---

[7] Hameed (1992) derived a related, but different, result in the context of a factor model for time-varying expected returns.

[8] Across all five portfolios, there is a slight tendency for our model's values to be lower than the actual estimates. However, the difference in magnitudes is economically quite small and is consistent with sampling error in the estimates.

Table 2
Weekly autocorrelations and the effect of outliers and subperiods

	Over-lapping	Monday	Tuesday	Wednes-day	Thursday	Friday	Wald Test for Seasonality, χ_i^2
Overall period	.362	.313	.306	.373	.408	.429	21.2
(quarterly resampling)	(.041)	(.045)	(.054)	(.044)	(.037)	(.032)	(.000)
w/o outliers		.311	.362	.409	.400	.429	
Number of outliers		77	68	76	80	83	
Overall period	.272	.209	.221	.306	.319	.324	19.2
(no resampling)	(.041)	(.050)	(.055)	(.039)	(.035)	(.034)	(.001)
1962–1969	.346	.273	.271	.359	.437	.411	10.5
	(.066)	(.070)	(.089)	(.084)	(.052)	(.059)	(.033)
1970–1977	.386	.320	.304	.397	.470	.476	17.2
	(.069)	(.083)	(.101)	(.071)	(.058)	(.057)	(.002)
1977–1984	.365	.308	.362	.423	.360	.395	11.2
	(.068)	(.064)	(.074)	(.061)	(.072)	(.045)	(.024)
1984–1990	.295	.325	.273	.268	.281	.352	2.8
	(.106)	(.085)	(.055)	(.104)	(.086)	(.065)	(.586)

Table 2 provides autocorrelations of returns on the small-firm portfolio over the period 1962–1990 and over four subperiods of equal length. The autocorrelations are estimated for weekly returns for weeks ending on different days of the week. Wald tests for equality between the weekly autocorrelations across the different days of the week are also given. Standard errors (of the autocorrelations) and p-values (of the statistics) are given in parentheses. Note that all of the estimates and test statistics have been adjusted for possible heteroskedasticity and serial correlation using the method of Newey and West (1987). Also provided is the autocorrelation of weekly returns over the whole period without outliers included. In particular, all observations outside 2 standard deviations were dropped for this case.

correlations across days of the week [e.g., Keim and Stambaugh (1984), and Bessembinder and Hertzel (1993)]. Here, we reexamine the influence of these and other patterns on the autocorrelation of weekly returns. Because the small-firm portfolio provides the most extreme example of time variation in short-horizon expected returns, in Table 2 we calculate weekly return autocorrelations for this portfolio ending on different days of the week. Because of the overlap these autocorrelations are highly correlated and should be very similar given the 1486 observations. The results in Table 2 suggest that this is not the case.

Specifically, irrespective of the sampling procedure of the portfolios [i.e., quarterly resampling as in Lo and MacKinlay (1988) versus no resampling as in Lo and MacKinlay (1990a)], the autocorrelations using weeks ending on Monday and Tuesday are substantially different from those using the other days of the week. For example, for the full sample period, the weekly autocorrelation pattern for the returns on the small-firm portfolio Monday through Friday are .313, .306, .373, .408, and .429 resampled quarterly and .209, .221, .306, .319, and .324 sampled just once. This can have important implications for interpreting the results. For example, consider the Lo and

MacKinlay (1990a) sampling procedure for portfolios described in the Appendix. The autocorrelation of weekly returns is .272 in the presence of overlapping observations, whereas it is .306 for days of the week ending on Wednesday. Since Wednesday has been the preferred choice of days among researchers, they have inadvertently overstated the autocorrelation magnitude in existing work.

Table 2 also provides a Wald test of the equality of the autocorrelations across the days of the week. The test statistic is adjusted for heteroskedasticity and cross-serial correlation (because of the overlapping errors across days of the week) using the method of Newey and West (1987). From a statistical point of view, the magnitude of the difference in the autocorrelations is large. For quarterly resampling, the Wald statistic is 21.2. Since the underlying distribution of this statistic is a χ^2_4, the *p*-value is .000.

We explore several possible explanations for the seasonal pattern in the estimated autocorrelations. First, holidays or other market closures are more likely to occur on some days than others, which may lead to a bias in the estimate. This is because our method for estimating the weekly autocorrelation follows Lo and MacKinlay (1988, 1990a), who treat holidays (and other "missing observations") as a zero return. During this period, Monday and Friday have the most missing observations (approximately 6.9 percent and 3.7 percent of the observations, respectively). In contrast, Tuesday has the least number of observations missing, approximately 1.8 percent. Given that weekly returns for weeks ending on Monday and Tuesday lead to similar autocorrelations, it seems unlikely that the pattern in autocorrelations can be explained via holidays or other market closures.

Second, there are several extreme observations over the 1962–1990 period (such as the market crash in October 1987). This could potentially lead to poor large-sample approximations even with a sample size of 1486 observations. We therefore check to see whether outliers can explain the seasonal pattern.[9] These results are given in Table 2. With or without extreme observations, both the seasonal pattern and the magnitude of the autocorrelations persist. For example, excluding outliers, the autocorrelations of weekly returns for weeks ending on Tuesday and Wednesday are .362 and .409, respectively, which are similar in magnitude to those reported earlier.

Third, nonstationarities may be present in the data, such that the usual standard errors, whether homoskedasticity- or heteroskedasticity-consistent, are not appropriate. Although it is not clear to us what

---

[9] We choose a standard objective criterion for omitting outliers. Specifically, we ignored observations that are 2.0 standard errors from the mean of small-firm returns. Depending on the day over which weekly returns are calculated, approximately 4.5 to 5.6 percent of the observations lie outside 2 standard errors.

Figure 1
Daily autocorrelations
The figure plots the daily autocorrelation of returns on a portfolio of small firms over the period 1962–1990. First- through fifth-order autocorrelations are provided for each day of the week. Thus, the point (Thursday, 2), for example, represents the autocorrelation between Thursday's return and the prior Tuesday's return (i.e., a two-day lag).

type of nonstationarity in the data can produce these seasonal differences at such short horizons, evidence of nonstationarity might show up in a subperiod analysis. Table 2 provides autocorrelations of weekly returns (for weeks ending on each day) over four equal subperiods: 1962–1969, 1970–1977, 1977–1984, and 1984–1990. As an example, consider the autocorrelation of weekly returns for weeks ending on Tuesday versus Wednesday. The Tuesday-week weekly autocorrelations are .271, .304, .362, and .273, respectively, over the four subperiods. In contrast, the Wednesday-week weekly autocorrelations are higher in the first three of four subperiods—.359, .397, .423—and slightly lower in the last—.268.

Therefore, there is some evidence that the seasonal may not be stationary. In particular, the difference in the weekly autocorrelations for weeks ending on Wednesdays versus Tuesdays drops from .088, .093, and .061 in the first three subperiods to −.005 in the last subperiod. Table 2 provides a joint Wald test of the no-seasonal hypothesis across the five days of the week for the four subperiods. The first three subperiods imply a seasonal at the 5 percent level. The Wald statistic for this final period is only 2.8, which represents a .586 p-value. Though this suggests that the seasonal in weekly returns may be less

important than implied by the overall 1962–1990 period, it should be pointed out that, in every subperiod, there is some seasonal (albeit small) and that this seasonal tends to be a particular direction.

Another way to look at the stationarity issue is to consider the magnitude of the autocorrelations across the subperiods. In a joint test of the null hypothesis that the autocorrelations are equal across subperiods for each day of the week, we obtain a statistic of 24.1. With 15 restrictions (3 for each of the five days of the week), the statistic is distributed asymptotically as a $\chi^2(15)$. The corresponding *p*-value is .064—a borderline rejection. This result, in conjunction with the large variability of the autocorrelations (e.g., .423 versus .268 for the latter two subperiods on Wednesdays), suggests caution in interpreting the standard errors of these estimates.

Given the overlap and the obvious high contemporaneous correlation across weekly returns on different days, what accounts for the seasonal patterns in weekly autocorrelations? Note that the weekly autocorrelations (if continuously compounded) are just different weighted combinations of correlations between different days. For example, the weekly autocorrelation for weeks ending on Wednesday is made up of daily autocorrelations from first to ninth order, including the first-order autocorrelation between Thursday's return and Wednesday's return, the ninth-order autocorrelation between Wednesday's return and the prior week's Thursday's return, and so on. Therefore, differences across days must account for the seasonal patterns.

Figure 1 provides estimates of the correlation between a particular day's return and each daily return during the previous five-day period. Figure 1 shows that extreme differences persist across both days and lags.[10] Given 1486 observations for each day, these daily autocorrelations impose sharp restrictions on viable explanations for short-horizon dynamics in stock returns. For example, as evident from the figure, there is a major difference in the first-order autocorrelations between Monday's return and the prior Friday's return (i.e., .52) compared to the correlation between Wednesday's return and the prior Tuesday's return (i.e., .18).[11] One possible explanation of this difference may be the importance of weekend returns versus nonweekend returns as documented by Keim and Stambaugh (1984) and Bessem-

---

[10] Note that Keim and Stambaugh (1984) explored, among other things, daily correlations in Dow Jones 30 firms over different days and also found a seasonal pattern (especially around weekends). Bessembinder and Hertzel (1993) extended the analysis to a longer time series, to additional periods of regular nontrading (such as holidays), and to a wide array of assets.

[11] Note that Monday's return is defined as the return from the previous Friday close to the Monday close. Thus, it is the return over both the weekend and Monday day.

[12] As a check to see if this pattern is unique to a particular market, such as the NYSE/AMEX, we repeated the tests using NASDAQ data over the 1972–1990 period. Essentially, the same pattern also occurs in this market.

binder and Hertzel (1993).[12] This view is supported by the fact that the test statistic for seasonals drops dramatically when we ignore the restriction that autocorrelations of weekly returns on Friday should equal those of Monday.[13] The only day not in common for these weekly returns is the weekend return (Friday close to Monday close), and, not surprisingly, this contributes to the rejection of the no-seasonality null in Table 2.

3. Nonsynchronous Trading: New Results

Since Fisher (1966) and Scholes and Williams (1977) first pointed out that nonsynchronous trading can induce positive autocorrelation in stock returns, researchers have investigated nontrading as the potential source of portfolio serial correlation. The effect of nonsynchronous trading on the autocorrelation of portfolio returns relies on two facts. First, daily closing stock prices on the Center for Research in Security Prices (CRSP) tapes reflect either the last trade of the day, which may have occurred hours before the official close of the exchange, or the average of bid and ask quotes, which may not have been updated for hours or days. Second, the vast majority of stock prices respond in the same direction to aggregate economic news. As an illustration, consider two stocks, one that trades at the close and one that only trades at the open but at no other time during the day. If significant economic news comes out during a particular trading day, the price reported on the CRSP tapes for the first stock on that day will reflect the news because the stock trades at the end of the day. In contrast, the reported price for the second stock on that day will not reflect the news because the last trade in this stock occurred before the news was released. The price of the second stock, however, will respond to the news when it trades the following day at the open of trade. As a result, the return on the first stock will seem to lead or predict the return on the second stock, even though this phenomenon is purely an artifact of the trading frequency of the two stocks. Spurious predictability across stocks leads to positive serial correlation in portfolio returns if these stocks are grouped into portfolios.

From an empirical standpoint, Atchison, Butler, and Simonds (1987),

---

[13] In particular, the statistic drops from 21.2 for the test

$$\rho_w = \rho_r = \rho_j = \rho_m = \rho_n$$

to 2.5 for the test

$$\rho_w = \rho_t = \rho_j, \qquad \rho_m = \rho_n$$

where, for example, ρ_m is the Monday week return autocorrelation.

Cohen et al. (1986), and Perry (1985) all found that the autocorre-lation is more prevalent for portfolios with more thinly traded stocks (such as small firms). Nevertheless, their analysis suggested that non-synchronous trading is not the only cause of autocorrelation in returns. More recently, Berglund and Liljeblom (1988), Conrad and Kaul (1988), Lo and MacKinlay (1988, 1990b), McInish and Wood (1991), Mech (1993), and Muthuswamy (1988) have investigated the auto-correlations of portfolio returns more closely. The overall conclusion, although not unanimous, is even stronger than in previous work. In particular, only a small part of the autocorrelation patterns in portfolio returns can be explained by nontrading effects. Our analysis, however, suggests that these nontrading effects may play a more important role than previously thought.

With respect to the prevailing wisdom, one of the most widely cited studies is Lo and MacKinlay (1990b). They developed a model of nontrading that yields closed-form solutions for both autocorrelation and cross-serial correlation patterns in portfolio returns. Using this model, they concluded that the existing evidence is consistent with nonsynchronous trading only under unreasonable assumptions about the probability of nontrading. The majority of the empirical results in Lo and MacKinlay (1990b) are derived in a setting in which the probability of nontrading in any fixed time interval is constant and in which stocks within a portfolio are homogeneous in terms of non-trading probabilities and covariation with the aggregate stock mar-ket.[14] For example, in their framework, if the probability of nontrading in any given hour is 20 percent, then there is an 80 percent probability that the last trade of the day for a particular stock will occur in the last hour. The probability of the last trade of the day occurring in the second-to-last hour is .2 × .8, or 16 percent. Similarly, the probability of the last trade occurring in the third-to-last hour is .2 × .2 × .8, or 3.2 percent. A priori there is no reason to believe that the assumption of time independence in nontrading is reasonable. This is especially the case given that most researchers believe that information flow (and thus trading) is clustered both within and across days. In addi-tion, it is likely that there will be heterogeneity in both nontrading probabilities and covariations with the market within portfolios, especially within portfolios of small stocks. In this section, we con-sider the effect of time dependence in nontrading and heterogeneity on portfolio return autocorrelations in the context of numerical exam-ples using data from the empirical literature.

---

[14] The theoretical model in Lo and MacKinlay (1990b) accommodates heterogeneous stocks, and they discussed extensions of their model to serial dependence in the nontrading process in a concluding paragraph. However, they argued that the induced autocorrelations are still too small to explain the existing evidence.

3.1 Time dependence in nontrading

The Markov structure in Lo and MacKinlay (1990b) places severe restrictions on the intradaily and intraweekly patterns of nontrading. Of particular importance, the geometrically declining probabilities of the last trade occurring in a fixed time interval as we move backward from the close of trade generate less severe nontrading than is apparent from the available data. The exact nature of the time dependence that creates these non-Markovian nontrading patterns is unclear, but the resultant increase in nontrading relative to the time-independent model will induce additional spurious autocorrelation in portfolio returns. Unfortunately, reliable intraday data does not exist over an extensive sample period. However, some studies do provide evidence on nontrading patterns. For example, Foerster and Keim (1993) and Keim (1989) provided estimates of daily nontrading for firms during the week. Examining small firms at the turn of the year over the period 1972–1987, Keim (1989) estimated that 73 percent of stocks last trade on the final day of the week, 12 percent last trade on the second-to-last day, 6 percent on the third-to-last day, 3 percent on the fourth-to-last day, and the remaining 6 percent do not trade during the final four days of the week, but no further breakdown is given.[15] This distribution of nontrading can be contrasted with the distribution given by Lo and MacKinlay's Markov model of nontrading. Specifically, if we match the 27 percent nontrading probability for the final day of the week, then the corresponding probabilities of the last trade occurring on a given day as we move backward from the end of the week are 73, 19.7, 5.3, 1.4, and 0.4 percent.

What are the implications of these two different nontrading distributions for portfolio autocorrelations? The model of Scholes and Williams (1977) provides a convenient framework in which to answer this question. This model has the advantage of being able to accommodate any distribution of nontrading with the restriction that all stocks must trade within a fixed time interval. In our case, this means that all stocks must trade once during the week. We impose this restriction on the time-independent distribution by adding the small remaining fraction to the first day of the week, yielding the nontrading distribution (73, 19.7, 5.3, 1.4, 0.5 percent). The Keim (1989) numbers are adjusted by distributing the 6 percent that did not trade in the final four days of the week to the first and second days of the week, yielding the nontrading distribution (73, 12, 6, 4.5, 4.5 percent).

In the Scholes and Williams (1977) setting, it is straightforward to

---

[15] The probability of trade on the last day of the week is consistent with data for other times of the year provided in Foerster and Keim (1993). Moreover, evidence in that article suggests that trading activity is high during the turn of the year, so the estimates are most probably conservative [for corroborating evidence, also see McInish and Wood (1991)].

analyze portfolio autocorrelations. Specifically, let s_{it} be the fraction of the trading period before the close during which security i does not trade, that is, s_{it} is the time between the close of trade and the last observed trade of security i. For example, using weekly data (or approximately 30 hours of trading), s_{it} might equal $\frac{1}{6}$ (or five hours). Following Scholes and Williams (1977), assume that s_{it} is independent through time.[16] In addition, it is well known that the autocorrelation of a well-diversified portfolio is approximately equal to the average cross-serial covariance between the stocks in the portfolio divided by the average contemporaneous covariance between these stocks [see, for example, Atchison, Butler, and Simonds (1987)]:

$$\operatorname{corr}(R_{pt}, R_{pt-1}) \approx \frac{\sum_{i=1}^{n} \sum_{j=1, \neq i}^{n} \operatorname{cov}(R_{it}, R_{jt-1})}{\sum_{i=1}^{n} \sum_{j=1, \neq i}^{n} \operatorname{cov}(R_{it}, R_{jt})}. \tag{4}$$

The intuition behind Equation (4) is analogous to the intuition underlying the concept of portfolio diversification. In particular, the cross-serial covariance and cross-covariance terms simply "outnumber" the autocorrelation and variance terms in the numerator and denominator of Equation (4), respectively.

Using results in Scholes and Williams (1977, p. 313) and making assumptions about either (a) the correlation of s_{it} and s_{jt} or (b) the coefficient of variation in returns, it is possible to show that, for an equally weighted portfolio of n securities,[17]

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

$$\tag{5}$$

The expression in Equation (5) depends solely on how the distribution of nontrading across securities is allocated within a given

---

[16] Even in a weekly time interval, it is possible that s_{it} is correlated through time. However, extending the model to allow this assumption complicates the derivations. It suffices to say that independent s_{it} tend to lower the induced autocorrelations; thus, our analysis understates the effect of nonsynchronous trading with respect to this assumption.

[17] If s_{it} and s_{jt} are positively correlated, this will reduce the amount of autocorrelation present in portfolio returns that is induced by nonsynchronous trading. However, two features of weekly returns diminish the effect from positive correlation of nontrading across securities. The first is that this correlation is most likely less for weekly data, since idiosyncratic reasons for nontrading most probably dominate in the longer run. Second, and most important, the coefficient of variation is generally high for smaller firms (which dominate the portfolio looked at here), thus weakening the influence of cross-correlation of the s_{it} and s_{jt}.

portfolio. Of particular interest, note that the distribution of the s_{it} itself matters, since we are concerned with the distribution of its maximum and minimum values. Thus, assuming equal probability of trading in any period [as in much of the empirical work in Lo and MacKinlay (1990b)] may lead to quite different effects than some alternative specification.

If we had access to reliable intraday data on small-firm returns over this sample period, then it would have been possible to estimate the implied autocorrelation in Equation (5) explicitly via estimation of the distribution of each firm's nontrading period, namely s_{it}. Instead, we use the formulation in Equation (5) and the nontrading estimates discussed above. For the moment, assume that the s_{it} are homogeneous, so that each firm's fraction of nontrading has the same distribution. Since each s_{it} has the same distribution, the average expectations in Equation (5) can be represented by the expectation for any pair of stocks:

$$\text{corr}(R_{pt}, R_{pt-1}) \approx \frac{E[\max(s_{it} - s_{jt}, 0)]}{1 - E[\max(s_{it}, s_{jt})] + E[\min(s_{it}, s_{jt})]}. \quad (6)$$

Finally, we need to make an assumption about the distribution of nontrading within the day. The Lo and MacKinlay (1990b) model with daily nontrading probabilities implicitly assumes that, if a stock trades during a specific day, then it trades at the close. In other words, stocks that trade on the final day of the week implicitly exhibit no nontrading at all. This assumption corresponds to setting s_{it} equal to 0 for stocks that last trade on the last day, s_{it} equal to .2 (one-fifth of a week) for stocks that last trade on the second-to-last day, and so forth. Initially, we will maintain this assumption, although it potentially understates the severity of nontrading.

Because closed-form solutions are nontrivial, we simulate the two distributions described above using a random number generator. We conduct 10,000 replications of each simulation and then numerically calculate the implied autocorrelation in Equation (6) using the distribution of the s_{it}. Using the slightly modified time-independent distribution (73, 19.7, 5.3, 1.4, 0.5 percent), nonsynchronous trading induces 6.3 percent autocorrelation. For comparison purposes, the closed form in Lo and MacKinlay (1990b) gives an autocorrelation of 6.6 percent for a daily nontrading probability of 27 percent. The discrepancy between these numbers can be attributed to the slight differences in nontrading distributions and the random fluctuations across simulations. To make all succeeding simulations as directly comparable as possible, the same set of random numbers is used in every case. Using the time-dependent nontrading distribution (73, 12, 6, 4.5, 4.5 percent), nonsynchronous trading induces 10.5 percent

autocorrelation. By relaxing the Markov assumption, we obtain auto-correlations that are approximately two-thirds higher than those implied by the time-independent probability structure. What is crucial is the amount of extreme nontrading (even if for only a few stocks) that takes place. Recall that the time-independent Lo and MacKinlay (1990b) model cannot precisely match the Keim (1989) estimates— it puts more weight on the last trade occurring on the second-to-last day of the week and very little weight on the first and second days of the week. The results here show that it is nontrading over these days, however slight, that induces higher autocorrelations.

We now proceed to relax the assumption that all daily trading occurs at the close of trade. No reliable data are available on the precise distribution of nontrading within the day, so we take a relatively conservative stance. Specifically, we assume that, of all the stocks that last trade on a given day, a fraction $\frac{6}{21}$ last trade at the end of the final hour (i.e., at the close), a fraction $\frac{5}{21}$ last trade at the end of the second-to-last hour, a fraction $\frac{4}{21}$ last trade an hour before that, and so on, until the final $\frac{1}{21}$ trade at the end of the first hour of trade (assuming a six-hour trading day). Using this nontrading distribution within each day and the daily nontrading distribution based on the Keim (1989) numbers, nonsynchronous trading generates 12.8 percent autocor-relation. This level of autocorrelation is approximately 20 percent higher than that generated with all trading occurring at the close of trade and more than twice as high as the autocorrelation generated under the original time-independent distribution.[18]

The results above have an important implication for interpreting the potential effect of nontrading. In particular, the distribution of nontrading within the day and across days has a large influence on weekly return autocorrelations. The problem with the Lo and MacKinlay (1990b) Markov model is that it places little weight on the tails of s_{it}. Thus, given the Foerster and Keim (1993) and Keim (1989) estimates of nontrading, it is the small probability of trading *very infrequently* that spuriously increases the autocorrelation. In addition, neglecting the within-day distribution of nontrading can also lead to an underestimation of the autocorrelation induced by nonsynchronous trading.

3.2 Heterogeneity in nontrading and betas
Another issue that has received little attention in the literature is the effect of the heterogeneity of stocks within portfolios. For example,

---

[18] Note that the within-day nontrading distribution corresponds closely to the nontrading distribution generated by applying the Markov model to hourly nontrading. In other words, setting the hourly nontrading probability to 80.4 percent, which generates daily nontrading of 27 percent, induces increases in autocorrelations similar in magnitude to those of switching from all trade at the end of the day to the distribution described above.

stocks within a portfolio may have both different nontrading probabilities and different covariations with the stock market as a whole. Thus, the nontrading distribution reported by Keim (1989) is most likely not representative of each firm, but more likely illustrates the fact that some stocks (for whatever reason) are always thinly traded relative to other stocks. Heterogeneity in nontrading across securities is important because the expression in Equation (5) concerns the joint distribution of the maximum and minimum values. So, for example, if some securities within the portfolio trade frequently and others trade infrequently, then the autocorrelation magnitude will be accentuated. As mentioned above, our analysis thus far assumes each stock in the portfolio has the same distribution of nontrading throughout the week. These estimates, however, are averages of nontrading across small firms. Keim (1989) also provided evidence of nontrading separately for NYSE versus AMEX firms. Within the small decile, there is clear heterogeneity, with AMEX firms having much larger nontrading periods than NYSE firms [see Keim (1989), Table 1].

Foerster and Keim (1993) provided some detailed evidence regarding the prevalence of heterogeneity of nontrading over the period 1973–1990. They calculated not only the average frequency of daily nontrading, but also various distribution percentiles for the size deciles. We use their numbers for the two smallest deciles to calibrate the degree of heterogeneity in nontrading. Note that our sample of stocks corresponds closely to that of Foerster and Keim (1993); therefore, nontrading in their two smallest deciles will correspond closely to nontrading in our smallest quintile.

For the examples that follow we use the Markov model in Lo and MacKinlay (1990b) to calculate portfolio autocorrelations. Specifically, consider a group of securities with unobservable returns R_{it}^* generated by the processes

$$R_{it}^* = \mu_i + \beta_i \Lambda_t + \epsilon_{it}, \tag{7}$$

where Λ_t is a common factor with variance σ_Λ^2 and ϵ_{it} are cross-sectionally and temporally independent idiosyncratic errors. In each period, each security trades with a constant probability p_i, which is time independent. If a security does not trade, its observed return is 0. If a security does trade, its observed return is the sum of the unobserved returns for that period and all past consecutive periods for which it did not trade. In other words, economic news (i.e., movements in the common factor) are reflected in stock prices only when a security trades. Under these assumptions, consider observed security returns aggregated over q periods, where the aggregated periods are indexed by τ, that is,

$$R_{i\tau}(q) = \sum_{t=(\tau-1)q+1}^{q\tau} R_{it}. \tag{8}$$

The covariance between the observed returns on two securities is given by

$$\mathrm{cov}[R_{i\tau}(q), R_{j\tau+n}(q)]$$

$$= \begin{cases} \left[q - \dfrac{p_i(1 - p_i^q)(1 - p_j)^2 + p_j(1 - p_j^q)(1 - p_i)^2}{(1 - p_i)(1 - p_j)(1 - p_ip_j)} \right] \beta_i\beta_j\sigma_\lambda^2 \\ \qquad \text{for } n = 0 \\[2mm] \dfrac{(1 - p_i)(1 - p_j)}{1 - p_ip_j} \left(\dfrac{1 - p_j^q}{1 - p_j} \right)^2 p_j^{(n-1)q+1}\beta_i\beta_j\sigma_\lambda^2 \\ \qquad \text{for } n > 0. \end{cases} \tag{9}$$

Equation (9) can easily be used to calculate portfolio autocorrelations using the result in Equation (4).[19]

Although stocks will be described in terms of their daily nontrading probabilities, the actual implementation uses the corresponding hourly nontrading probabilities (for a trading day of six hours) to avoid the understatement in autocorrelations associated with assuming that all stocks trade at the close. It is important to note, however, that no adjustment is made for potential time dependence in nontrading, and consequently the autocorrelations may be understated. In each case we use six classes of stocks, which comprise the fractions 5, 20, 25, 25, 20, and 5 percent of the stocks in the portfolio. Throughout, the average daily nontrading probability in the portfolio is kept at 27 percent to make for easier comparisons with the earlier results.

We consider four different distributions of nontrading, which are motivated by the results in Foerster and Keim (1993, Table 3). The first distribution is homogeneous, with nontrading probabilities (27, 27, 27, 27, 27, 27 percent), and the other three are heterogeneous with nontrading probabilities—(0, 11, 21, 32, 44, 55 percent), (0, 0, 11, 32, 60, 85 percent), and (0, 0, 0, 43, 60, 85 percent). The distributions are ordered in terms of increasing heterogeneity and reflect various degrees of conservatism with respect to heterogeneous probabilities provided by Foerster and Keim (1993).[20] We also consider three different distributions of market betas across the six classes of

---

[19] Note that Equation (9) here differs somewhat from Equations (3.6) and (3.12) in Lo and MacKinlay (1990b), which contain some errors.

[20] For completeness we provide the relevant numbers. For the smallest decile, the nontrading probabilities for the 99th, 95th, 90th, 75th, 50th, 25th, and 10th percentiles are 75.9, 60.2, 51.0, 36.7, 20.0, 7.4, and 0.8 percent, with a mean nontrading probability of 31 percent. For the second smallest decile, the nontrading probabilities are 62.7, 47.1, 37.1, 22.5, 9.1, 2.0, and 0.0 percent, with a corresponding mean of 19 percent.

Table 3
**The effects of heterogeneity on portfolio autocorrelations**

	Nontrading probabilities, %	Betas (1, 1, 1, 1, 1, 1)	(0.8, 0.9, 1, 1, 1.1, 1.2)	(0.8, 1.2, 1.6, 2, 2.4, 2.8)
1	(27, 27, 27, 27, 27, 27)	8.90	8.90	8.90
2	(0, 11, 21, 32, 44, 55)	10.33	10.72	11.60
3	(0, 0, 11, 32, 60, 85)	13.18	14.20	16.43
4	(0, 0, 0, 43, 60, 85)	13.97	15.08	17.82

Table 3 provides autocorrelations of weekly returns for portfolios of stocks with heterogeneous nontrading probabilities and betas. Each portfolio consists of six classes of stocks in proportions (5, 20, 25, 25, 20, 5 percent). Nontrading probabilities are expressed on a daily basis, although the actual computations use the corresponding hourly nontrading probability (six hours per day). Autocorrelations are calculated using the closed form solutions to the Markov model in Lo and MacKinlay (1990b).

stocks. Again, the first distribution is homogeneous with betas (1, 1, 1, 1, 1, 1), and the other two are heterogeneous with betas (0.8, 0.9, 1, 1, 1.1, 1.2) and (0.8, 1.2, 1.6, 2.0, 2.4, 2.8).

What are the implications of the heterogeneity described above for portfolio autocorrelations? Table 3 reports the autocorrelations for the 12 combinations of nontrading probabilities and betas described above. The clear implication of the results is that heterogeneity across stocks within a portfolio can cause dramatic increases in the induced spurious autocorrelation. In particular, the autocorrelation for the portfolio with the most heterogeneity in both nontrading and betas is 17.82 percent. This autocorrelation is more than twice the magnitude of the autocorrelation for the homogeneous portfolio (8.90 percent). Table 3 illustrates that large spurious autocorrelations are driven principally by severe nontrading in some stocks coupled with frequent trading in others (i.e., by the tails of the nontrading distribution). Although the last two distributions differ substantially in the "center," they have identical tail distributions, and their autocorrelations are within approximately 1 percent of each other. Table 3 also illustrates the interaction effect of heterogeneity in betas and nontrading. In particular, for homogeneous nontrading, the distribution of betas has no effect on portfolio autocorrelations, but variation in betas can amplify the effects of heterogeneity in nontrading. These results are in stark contrast to conclusions reached by Lo and MacKinlay (1990b).

To understand better the relation among nontrading, betas, and autocorrelations, consider a simple portfolio that consists of two types of stocks. Three-quarters of the stocks trade all the time and have a beta of 1. The daily nontrading probability of the remaining quarter is varied from 20 to 95 percent, and their betas take on the values 1,

Figure 2
The effects of heterogeneity in nontrading and betas on portfolio autocorrelations
This figure shows the spurious weekly autocorrelation induced by nonsynchronous trading for a portfolio composed of two types of stocks. Of the stocks 75 percent trade all the time and have a beta of 1. We vary the daily nontrading probability of the other 25 percent of the stocks from 20 to 90 percent (*x*-axis) and plot autocorrelations for betas of 1, 1.5, 2, and 2.5.

1.5, 2, and 2.5. Figure 2 plots the resulting portfolio autocorrelations, where the *x*-axis measures the daily nontrading probability and the four lines represent the four different betas. It is immediately clear that variation in betas is relatively unimportant for small degrees of relatively homogeneous nontrading. However, as the probability of nontrading in some of the stocks increases, the effect of variation in betas increases. In particular, for a nontrading probability of 80 percent, an increase in beta from 1 to 2.5 increases the portfolio autocorrelation from 12 to 25 percent. Seen from a slightly different perspective, an increase in nontrading from 30 to 80 percent increases portfolio autocorrelations from 6 to 25 percent for beta equal to 2.5, but only from 3 to 12 percent for beta equal to 1. Finally, increasing nontrading beyond a certain point decreases rather than increases portfolio autocorrelations. It is important to note that, for the purposes of this analysis, it is the ratio of the betas of the two classes of stocks rather than their absolute levels that is important. For example, the graph would look identical if 75 percent of the stocks had a beta of 0.5 and the other 25 percent had betas of 0.5, 0.75, 1, and 1.25.

Note that the caution with which researchers should interpret portfolio autocorrelations is not restricted to the equity market. Return autocorrelations for portfolios of corporate bonds, for example, may also be subject to the same nonsynchronous-trading-induced biases discussed above. The problem will be exacerbated by combining infrequently traded, low-grade securities with their less risky counterparts. To see this, first recall that heterogeneity in betas increases the spurious autocorrelation and that it is the ratio of the betas that determines the magnitude of this effect. As an illustration, consider the case in which the 5 percent of securities with highest nontrading probabilities also have a return beta of 1 while all other securities have a beta of 0.01. For nontrading distributions 3 and 4 described above, this beta distribution induces an autocorrelation of approximately 55 percent. The intuition is simple. Recall that the autocorrelation of returns on a portfolio is just the average cross-serial covariance of these securities divided by the average cross-covariance, as in Equation (4). Thus, in the setting of Equation (7), high beta securities will almost completely determine the autocorrelation of the portfolio. Although we are not claiming that this is a relevant set of assumptions in our particular context, the magnitude of the spurious autocorrelation caused by nonsynchronous trading suggests caution in interpreting all autocorrelations at short horizons.

Although heterogeneity in betas and nontrading across stocks have dramatic effects on spurious portfolio autocorrelations, the analysis above is still limited to the case in which distributions for individual stocks are independent and identically distributed over time. Temporal heterogeneity may induce additional biases that are impossible to quantify without a detailed model. For example, it is conceivable that seasonals in nonsynchronous trading during the week could lead to seasonal patterns in the autocorrelations. In this particular instance, however, two characteristics of the data suggest that this is not the case. First, in a detailed study, Foerster and Keim (1993) did not find seasonals in nontrading. Although their results refer to the amount of average nontrading across portfolios and not the distribution of nontrading for each stock, we have no a priori reason to believe that heterogeneity in nontrading across stocks is more severe on one day than any other. Second, as we shall see in the next section, the seasonal, albeit with smaller magnitudes, appears also in the futures market. Clearly, seasonals in nontrading for individual stocks should have no effect in the futures market.

The main conclusion from our analysis is that the amount of autocorrelation due to nonsynchronous trading in stocks has been understated in the literature. Accounting for heterogeneity within a portfolio can lead to spurious autocorrelations 2 to 3 times higher than

previously believed. Nonetheless, nonsynchronous trading falls short of explaining all of the autocorrelation in portfolios of small stocks. The remaining autocorrelation may also be consistent with the loyalist viewpoint, and we explore this issue from a different angle in the next section.

4. A Horse Race between the Three Schools

We introduced this article by describing three schools of thought currently at work in finance. The loyalist view is that the magnitude of short-horizon autocorrelations is due to market frictions. One explanation, which is discussed in detail above, is that nontrading has a far more important role than previously thought. In this section, we proceed one step further by providing an ex ante test of hypotheses implied by the three schools—loyalists, heretics, and revisionists. Here, we claim that comparisons between autocorrelation patterns of returns on the futures on an index and the underlying index itself are a useful way to differentiate theories of short-horizon dynamics in returns. First, consider the revisionist view of the short-horizon autocorrelations of stock returns. They believe the magnitude can be explained via time-varying factor risk premiums. If autocorrelations on index returns are due to changing risk premiums, then (institutional features aside) this evidence should also show up in returns on the futures on that same index. Since the futures are priced off the "true" spot index, futures should display the same magnitude of autocorrelation patterns as spot returns.

To see this, note that there is strong support in the literature for the cost of carry model as applied to futures pricing.[21] Via an arbitrage argument, this model implies that the futures price is simply the current spot price times the compounded rate of interest (adjusted for paid dividends):

$$F_{t,T} = S_t e^{(i-d)(T-t)}$$

where $F_{t,T}$ is the futures price of the index, maturing in $T - t$ periods; S_t is the current level of the index; i is the continuously compounded rate of interest; d is the continuously compounded rate of dividends paid, and T is the maturity date of the futures contract. Thus, under the cost of carry model and appropriately adjusting for time to maturity, we can write the return on the futures as

$$r_F = r_S + \Delta(i - d), \tag{10}$$

---

[21] MacKinlay and Ramaswamy (1988) provided a discussion of this model. As they pointed out, under stochastic interest rates, the model is strictly valid only for forward contracts. Nevertheless, they argued that differences between forwards and futures prices are negligible from a practical point of view.

where r_F and r_S are the continuously compounded returns on the futures and the underlying spot index, respectively, and $\Delta(i - d)$ is the change in the continuously compounded interest rate (adjusted for the dividend rate).[22] Thus, if autocorrelations of spot index returns are due to fundamental movements in factors and not the institutional structure of the spot market, futures returns should have approximately the same autocorrelation pattern as spot returns.[23]

Second, consider the heretic view of the autocorrelation patterns. They argue that market participants only partially adjust to information arriving to financial markets. Thus, positive autocorrelation is induced as information slowly gets incorporated into stock prices. Since these spot prices are real (albeit irrational) transaction prices, one implication of this type of market inefficiency is that partial adjustment should show up in both the spot and futures markets. This is because the "true" spot index and futures are linked via the no-arbitrage condition. Although this requires some degree of rationality in the marketplace, it is well known that index arbitrage continually takes place between these two markets.

The other school of thought, the loyalists, describes autocorrelation of the index returns via market frictions. The idea is that market microstructure biases (such as measurement error) induce non-risk-based or nonexploitable autocorrelation in portfolio returns. For example, consider the aforementioned measurement problem in observed prices due to nonsynchronous trading. It is well known that nonsynchronous trading in individual securities can induce positive autocorrelation at the portfolio (i.e., the spot index) level. Moreover, this autocorrelation can be severe if the portfolio contains securities that trade relatively infrequently (such as small firms). In contrast, the futures return will not pick up this autocorrelation, since it is priced off the true fundamentals. Of course, the futures contract may suffer from nonsynchronous trading itself, which can generate negative autocorrelation in its return series. However, at the weekly level, and given that futures contracts are frequently traded, nonsynchronous trading should not be a problem. Therefore, in contrast to the other schools, for this particular microstructure bias, the loyalists' main implication is that returns on the spot index should display some amount of spurious autocorrelation, whereas the futures index should display none.

---

[22] Note that under the cost of carry model, $\Delta(i - d)$ equals zero. Assuming the cost of carry model still approximately holds under stochastic interest rates, this term can in theory time-vary. We address this issue in more detail below.

[23] Of course, this assumes that the futures market has no unique institutional structure, which can induce large differences in autocorrelations. In addition, variations in $\Delta(i - d)$ are assumed to be small.

Table 4
The autocorrelation of futures versus spot returns

Index	Over-lapping	Monday	Tuesday	Wednes-day	Thurs-day	Friday	Wald for Season-ality, χ_i^2
Small-firms	.319	.330	.296	.299	.309	.365	2.3
	(.088)	(.071)	(.049)	(.089)	(.073)	(.055)	(.685)
Value-Line spot	.174	.151	.120	.180	.204	.220	8.4
	(.086)	(.078)	(.062)	(.083)	(.068)	(.071)	(.078)
Value-Line futures	.046	−.015	−.003	.087	.111	.076	13.5
	(.065)	(.069)	(.050)	(.073)	(.066)	(.085)	(.009)
Wald Test for equality of futures and spot, χ_i^2	18.8	44.7	30.1	11.8	15.3	8.6	
	(.000)	(.000)	(.000)	(.001)	(.000)	(.003)	
NYSE spot	.003	−.026	−.055	.044	.057	.025	8.1
	(.063)	(.061)	(.057)	(.077)	(.068)	(.081)	(.088)
NYSE futures	−.063	−.111	−.104	.026	.014	−.054	6.00
	(.048)	(.047)	(.037)	(.072)	(.069)	(.077)	(.202)
S&P spot	−.024	−.054	−.080	.022	.033	−.008	8.3
	(.060)	(.059)	(.056)	(.075)	(.066)	(.080)	(.081)
S&P futures	−.067	−.109	−.106	.034	.000	−.074	7.1
	(.049)	(.049)	(.038)	(.072)	(.064)	(.071)	(.131)
(4NYSE − 3S&P) spot	.059	.039	.020	.092	.102	.049	4.0
	(.073)	(.072)	(.056)	(.081)	(.075)	(.087)	(.403)
(4NYSE − 3S&P) futures	−.076	−.119	−.119	−.038	.020	−.033	4.9
	(.044)	(.045)	(.040)	(.073)	(.077)	(.093)	(.299)
Wald Test for equality of futures and spot, χ_i^2	7.3	8.3	11.1	5.8	4.2	1.4	
	(.007)	(.004)	(.001)	(.016)	(.042)	(.238)	

Table 4 provides a comparison of the autocorrelation of weekly returns on both spot indices and corresponding futures on these indices. The autocorrelations are estimated for weekly returns for weeks ending on different days of the week, covering the sample period 1982–1991. In particular, we consider three different portfolios that place weight on small firms: (*a*) the small-firm portfolio (described in Table 1) for the 1982–1991 period, (*b*) the Value-Line index (which is an equally weighted index of a broad cross-section of stocks), (*c*) the NYSE composite spot and futures and the S&P index spot and futures, and (*d*) 4 times the NYSE composite minus 3 times the S&P 500 index (which leaves, by construction, a portfolio of smaller firms). Wald tests for equality between the futures and spot return autocorrelations are provided, as well as a Wald test that the weekly autocorrelations are equal across the different days of the week. Standard errors (of the autocorrelations)) and *p*-values (of the statistics) are given in parentheses. Note that all of the estimates and test statistics have been adjusted for possible heteroskedasticity and serial correlation using the method of Newey and West (1987).

There is a growing literature in finance that relates prices of futures to the pricing of the underlying assets. For example, MacKinlay and Ramaswamy (1988), among others, looked at intraday arbitrage trading between the futures and spot for the Standard & Poor's (S&P) 500 index. They reported some differences in autocorrelations between returns on the futures and spot at short intervals but find that most of this goes away at the daily level. Harris (1989), in a different context, reached similar conclusions. Furthermore, in a recent article, Bessembinder and Hertzel (1993) documented interesting patterns in autocorrelations of futures and spot returns around trading and non-

trading periods. Their conclusion, however, is that both the futures and spot returns have similar patterns. These researchers have focused on especially short intervals and, with respect to stock indices, on market value-weighted portfolios like the S&P 500.

In terms of our analysis of weekly returns, note that the S&P 500 displays little autocorrelation. Thus, there is apparently little difference in the behavior of the futures and spot returns. As we have argued, however, much of the action in the autocorrelation patterns of weekly returns comes from the small-firms stock portfolio. In studying this question, therefore, it is of paramount importance to use indices dominated by smaller firms. That is, the larger the proportion of small firms in the index, the larger the potential difference between the autocorrelations of spot and futures returns (i.e., under the loyalist viewpoint). Table 4 reports the autocorrelations of the spot and futures returns of the NYSE composite and the S&P 500. As predicted, the difference between the NYSE spot and futures return autocorrelation is .06 versus a smaller difference in the S&P spot and futures' return autocorrelation of .04.

However, since both of these portfolios are value-weighted, there is clearly little autocorrelation in either market. To coincide with our previous analysis, it is important to use indices with much more weight on small firms (e.g., an equal-weighted market index). One such index traded in futures markets is the Value-Line index on the Kansas City Board of Trade. Using data on this contract from 1982 to 1991, we calculate the weekly autocorrelation of returns on both the spot and futures of the Value-Line index.

We focus our analysis on the overlapping weekly returns; however, to coincide with existing studies, we report autocorrelation patterns for nonoverlapping weekly returns.The results are presented in Table 4. The spot index returns exhibit an autocorrelation of .173 versus only .046 for the autocorrelation of the corresponding futures returns.[24] Moreover, the difference in autocorrelations is statistically significant. In particular, the Wald statistic for equality between the autocorrelations on the spot index and futures equals 18.8, which represents a p-value equal to .000. While the autocorrelation of the return on the Value-Line index is significantly different from zero (at the 5 percent level), the autocorrelation of the corresponding futures return

---

[24] The index was geometric until 1988, when it changed to arithmetic. Some recent evidence in Thomas (1992) suggests some mispricing in the futures market during the early years of the contract. However, this mispricing seems to be based on differences between geometric and arithmetic means and thus should not affect the autocorrelations [see Thomas (1992)]. We also performed our analysis on the two unspliced series separately (i.e., the geometric from 1982 to 1988 and the arithmetic from 1988 to 1991). The results are qualitatively similar to the ones reported here. In particular, the spot's autocorrelation was .166 and .167 for the respective subperiods, while that of the futures was .027 and .065.

is not significant at the 5 percent level. This is consistent with the loyalist view described above.

There are two reasons why the autocorrelation is somewhat lower than the 36 percent for small-firm portfolio returns given in Table 1. First, the Value-Line index, although equally weighted, includes a substantial number of large firms (which do not seem to be serially correlated). For example, Table 1B shows that the equal-weighted index of NYSE and AMEX stocks over the 1962–1990 sample period also has a lower autocorrelation (i.e., 23 percent). Second, the auto-correlation for the small-firm portfolio returns is somewhat lower in the 1980s than the earlier periods. To see this, Table 4 provides the autocorrelation for the small-firm portfolio returns during the 1982–1990 sample period. The overlapping autocorrelation estimate equals 32 percent, which is lower than that of the overall period. Thus, the 17 percent autocorrelation of the Value-Line index partly reflects the general sample period.

An alternative way to obtain information about small-firm portfolio returns in spot and futures markets is to construct a proxy using a judiciously chosen combination of indices that contain both small and large stocks. Unfortunately, this is necessary because futures contracts on small-firms portfolios are not traded over this period. In particular, assume that the stocks in the S&P 500 index are a subset of those in the NYSE composite index. This is not exactly the case, but the majority of S&P 500 firms are in fact contained in the NYSE composite. Thus, the NYSE composite return can be approximately represented as a weighted sum of the S&P 500 index return and the return on the remaining stocks in the NYSE (which are, for the most part, smaller stocks).

To see this mathematically, note that the level of the NYSE composite can be broken up into two components:

$$P_t^{\text{NYSE}} \equiv P_t^{\text{S\&P}} + P_t^{\text{SM}},$$

where P_t^{NYSE}, $P_t^{\text{S\&P}}$, and P_t^{SM} are the market values of the NYSE composite, the S&P 500, and the remaining smaller firms portfolio, respectively. Therefore, the corresponding return on the NYSE index is

$$R_{t,t+1}^{\text{NYSE}} = \frac{P_{t+1}^{\text{S\&P}} + P_{t+1}^{\text{SM}}}{P_t^{\text{S\&P}} + P_t^{\text{SM}}}$$

$$= \left(\frac{P_t^{\text{S\&P}}}{P_t^{\text{S\&P}} + P_t^{\text{SM}}}\right)\left(\frac{P_{t+1}^{\text{S\&P}}}{P_t^{\text{S\&P}}}\right) + \left(\frac{P_t^{\text{SM}}}{P_t^{\text{S\&P}} + P_t^{\text{SM}}}\right)\left(\frac{P_{t+1}^{\text{SM}}}{P_t^{\text{SM}}}\right)$$

$$\equiv w_t^{\text{S\&P}} R_{t,t+1}^{\text{S\&P}} + w_t^{\text{SM}} R_{t,t+1}^{\text{SM}}, \tag{11}$$

where $R_{t,t+1}$ is the return from t to $t+1$ and $w_t^{\text{S\&P}}$ and w_t^{SM} are the weights of the S&P 500 index and the remaining portfolio in the NYSE

composite index. During the period 1982–1991, the S&P 500 index accounts for approximately 75 percent of the value of the NYSE composite index.[25] Given that the weights over the 1982–1991 period are $w_t^{S\&P} = .75$ and $w_t^{SM} = .25$, we can use Equation (11) to construct the return on the smaller firms portfolio. Specifically,

$$R_{t,t+1}^{SM} = 4R_{t,t+1}^{NYSE} - 3R_{t,t+1}^{S\&P}. \qquad (12)$$

We therefore construct two return series using Equation (12): (*a*) spot returns on the smaller firm portfolio and (*b*) its corresponding futures returns. Table 4 provides autocorrelation estimates and corresponding test statistics for these constructed series. For overlapping data, the spot return autocorrelation of this series is .059, while that of the futures is −.076 percent. The difference in autocorrelations between the futures and spot index return is significant at the 5 percent level (the Wald statistic for equality equals 7.3, with a corresponding *p*-value of 0.7 percent). The evidence is consistent with the earlier findings that the spot autocorrelations tend to be higher than those of the futures. Note, however, that the magnitude of the spot return autocorrelation is much lower than that of the equal-weighted small-firm return series and the Value-Line series given in Tables 1 and 4, respectively. This is not surprising, since the constructed series is based on a value-weighted portfolio return series of smaller stocks from the NYSE. The evidence is, however, consistent with the loyalist view, as the futures return series exhibits an autocorrelation indistinguishable from zero.

There are, however, some alternative explanations of the difference in autocorrelation patterns between futures and spot markets. First, note that we have focused our argument on the institutional structure of equity markets. Perhaps it is the futures markets that suffer from microstructure effects. For example, consider the revisionist view of the 17 percent autocorrelation of the Value-Line spot return. They might argue that this is the true autocorrelation due to time-varying factors and that the 4 percent autocorrelation of futures returns is in fact understated because of institutional reasons. There is some credibility to this argument. For example, note that the futures contract is subject to a bid-ask bias, which will lead to negative autocorrelation and thus a downward bias in the autocorrelation estimate. This reason is unlikely to be the complete explanation, however, because the bid-ask bias is fairly small on the various futures contracts and our return horizon is relatively long (i.e., a week).

To see this, consider the bid-ask model of Blume and Stambaugh (1983), as applied to the futures price:

---

[25] The value of the S&P 500 index, as a fraction of that of the NYSE index, is quite stable over the period. Its highest value is 78 percent, and its lowest value is 73 percent.

$$F_t^m = F_t(1 + \delta_t),$$

where F_t^m is the observed futures price at time t and δ_t equals the adjustment due to the trade taking place at the bid or ask price. For illustrative purposes, suppose there is an equal probability of being at a bid versus an ask, such that $\delta_t = s$ if trade takes place at an ask and $\delta_t = -s$ if trade takes place at a bid. Over the 1982–1991 sample period, the futures price on the index ranged from 123 to 316. To be conservative, consider the lowest range for the index (i.e., 123) and various values for the level of the spread (i.e., $2s \times F_t$) equal to 10 cents, 20 cents, . . . , 50 cents. For the extreme case of a 50 cents spread, it is possible to show that the autocorrelation of the futures weekly return adjusted for bid-ask bias increases from .043 to only .050. Furthermore, note that in active futures markets, such as those discussed here, the bid-ask spread is approximately 2 ticks. Since the minimum price movement is 5 cents for the Value-Line futures, a more reasonable estimate of the spread is 10 cents. Thus, empirically, bid-ask bias in the futures price cannot explain the differences in autocorrelations of futures and spot returns on the Value-Line index.

A second potential explanation deals with violations of assumptions implicit in the cost of carry model. In particular, since interest rates are stochastic, it is possible that a positively autocorrelated factor in the index returns could be offset by negative correlation in an "interest rate" factor, leading to a small autocorrelation of returns on the futures. To see this, consider the equation for the futures return in Equation (4). Assuming that the dividend rate is constant, it is possible to express the autocorrelation of the futures return as

$$\rho_{r_F} = \frac{\text{cov}(r_{S_t}, r_{S_{t-1}}) + \text{cov}(\Delta i_t, \Delta i_{t-1}) + \text{cov}(\Delta i_t, r_{S_{t-1}}) + \text{cov}(r_{S_t}, \Delta i_{t-1})}{\text{var}(r_{S_t} + \Delta i_t)}$$

$$(13)$$

As an illustration, suppose that interest rate changes and the spot index returns are uncorrelated. Then the autocorrelation of the futures return is just

$$\rho_{r_F} = \frac{\text{cov}(r_{S_t}, r_{S_{t-1}}) + \text{cov}(\Delta i_t, \Delta i_{t-1})}{\text{var}(r_{S_t}) + \text{var}(\Delta i_t)}.$$

Thus, the futures return can exhibit a small autocorrelation only if (*a*) changes in interest rates are negatively autocorrelated and (*b*) these changes display substantial variation relative to equity returns.

We collected weekly data on three-month interest rates for every Wednesday over the sample period 1982–1991. In terms of the cost of carry model, the interest rates should match the rates that mature on the date of expiration of the futures contract. Although these data

Table 5
Correlations and covariations between returns on the Value-Line spot and futures index and changes in interest rates

	R_{ft-1}	R_{ft}	R_{st-1}	R_{st}	Δi_{t-1}	Δi_t
R_{ft-1}	375.6	.0459	.9417	.1814	.0775	.0286
R_{ft}	17.18	373.5	.0672	.9414	−.0242	.0863
R_{st-1}	302.7	21.54	275.0	.1727	.1041	.0188
R_{st}	58.11	300.7	47.34	273.2	.0083	.1136
Δi_{t-1}	.0478	−.0149	.0549	.0044	.0010	.0101
Δi_t	.0176	.0527	.0099	.0594	.0000	.0010

Table 5 provides estimates of the comovements between the lagged and current weekly return on the Value-Line index, the lagged and current weekly return on futures on the Value-Line index, and lagged and current weekly changes in three-month interest rates. Specifically, overlapping covariances, autocovariances, and correlations are provided for weekly data over the period 1982–1991. The top triangle of Table 5 provides the correlations, while the bottom triangle and diagonal provide covariances and variances. The covariances and variances are given in annualized percentage terms. R_{st} denotes the return on the spot index, R_{ft} the return on the futures on the spot index, and Δi_t the change in interest rates.

are not available, variation in short-term rates is so similar that the use of three-month rates is practically equivalent to the correct procedure [see Thomas (1992)]. Using Equation (13), we estimated the implied futures autocorrelation from the data on the Value-Line spot index returns and changes in interest rates. Table 5 provides estimates of each element of Equation (13), as well as the corresponding correlation matrix of the variables. Substituting these values into Equation (13), the implied autocorrelation of the futures return is not close to the estimated autocorrelation of .046. In fact, the implied autocorrelation is .183, which slightly exceeds that of the underlying spot index, namely .173. Thus, movements in interest rate changes do not seem to provide a plausible explanation for the different autocorrelations of futures and spot returns on small-firm-based indices.

There is some evidence that the seasonal pattern documented in the equity markets may in fact carry through to the futures market. Even though the futures weekly return autocorrelation is economically small and, in fact, statistically insignificant for each day (see Table 4), a Wald statistic of equality across the five days equals 13.5, which is significant at the 1 percent level. This is because even slight deviations can be statistically important (although perhaps of little economic consequence) when the correlation between the weekly autocorrelation estimators is so high. While the futures seasonal is not predicted a priori by any of the schools of thought, some may interpret it as either (*a*) consistent with a microstructure explanation based on information flow [e.g., Admati and Pfleiderer (1988, 1989)] or (*b*) consistent with systematic seasonals in risk factors, with microstructure biases that drive a wedge between futures and spot return autocorrelations.

5. Concluding Remarks

The evidence presented in this article should not be looked at in isolation. Instead, we must also consider the volumes of evidence suggesting that markets react quickly to information, such as announcements of earnings, dividends, and takeovers. These studies cast further doubt on the heretic view that lead-lag relations, such as those discussed in this article, are due to the delayed reaction of small stocks to news.

With respect to the revisionist explanation, which alludes to economic risk premiums as the source of autocorrelations and cross-serial correlations in size portfolio returns, we conclude that these are an unlikely source. For example, the effect of nonsynchronous trading, in a case where returns are not autocorrelated, can conceivably be 18 percent or higher, for the portfolios of interest. It remains an open question whether the remaining amount of autocorrelation can be attributed to time-varying economic risk premiums or, more likely, some other microsctructure effects.

In an attempt to perform an ex ante test of the three schools' views, we examine the autocorrelation properties of small-firm indices and futures contracts written on them. Supportive of the loyalist view, we ·find that the spot index's autocorrelation is significantly higher than that of the futures. In addition, we find that the futures' index auto-correlation is indistinguishable from zero.

Apparently, reports of the death of *market efficiency* have been premature and greatly exaggerated.

Appendix: The Data

In this appendix we describe the various data sources used throughout the article.

The size portfolios

The stock data come from the daily 1990 CRSP NYSE/AMEX tapes, and stocks are sorted into portfolios every 13 weeks based on the market value of their equity. Specifically, starting with price and shares outstanding data for Monday, July 2, 1962, all of the firms that have price data for that day and existed for the 13-week period beginning on the subsequent Thursday are sorted into five portfolios of approximately equal size. Daily return data for these firms are compiled for the 13-week period, with missing observations replaced by zero returns. Daily, equal-weighted, simple returns are then calculated for each portfolio. Weekly portfolio returns are calculated by aggregating the daily data over five business-day periods. Daily returns for the

following 13 weeks are calculated for portfolios formed based on data for the Monday preceding the beginning of the 13-week period. The first daily return observation is on Thursday, July 5, 1962, and the last observation is on Monday, December 31, 1990 (7433 observations). Consequently, for weekly returns there are 1486 observations.

Rebalancing the portfolios less frequently generates qualitatively similar results, but the magnitudes of the autocorrelations decrease because the requirement that a firm exist over the whole period between rebalancing tends to exclude smaller firms. For example, we also formed portfolios using a method similar to the procedure used by Lo and MacKinlay (1990a) (i.e., sorting firms into portfolios only once over the whole sample period). Specifically, the sample of firms is restricted to those firms that traded throughout the period from July 2, 1962 to December 31, 1990. In addition, all firms with more than 50 missing daily observations are excluded. Firms are sorted into size quintiles based on the market value of their equity at the midpoint of the sample period (October 1, 1976). Using this sorting technique, weekly portfolio autocorrelations (Wednesday close to Wednesday close) for the five quintiles are .306, .250, .173, .129, and .043.

Spot and futures data on indices

The spot and futures data on the Value-Line Composite Average, the S&P 500 index, and the NYSE composite index cover the sample period 1982–1991. The Value-Line futures contracts trade on the Kansas City Board of Trade, and the underlying index is a geometric mean of stock prices before 1988 and an arithmetic mean thereafter. The data are constructed according to usual conventions. In particular, a single time series of futures prices are spliced together from individual futures contracts. For liquidity, the nearest contract's prices are used until 10 days to maturity and then the next nearest is used, and so on [see Thomas (1992) for more details].

The S&P 500 and NYSE composite futures contracts trade on the Chicago Mercantile Exchange (CME) and the New York Futures Exchange (NYFE), respectively. Note that both are value-weighted indices. In calculating the implied small-firms portfolio from both the futures contracts and the spot markets, it is necessary to calculate the relative weight of the S&P 500 in the NYSE composite index. To do this, we collected data on the total market value of the stocks in the NYSE composite index from the NYSE Fact Book and the total market value of the stocks in the S&P 500 from the Standard & Poor's Corporation "S&P Information Bulletin" (some recent numbers were obtained over the phone directly from Standard & Poor's). All market values are end-of-year values.

Interest rates

The interest rates used in this article are three-month rates, reported daily over the sample period 1982–1991. The source for the data is the Board of Governors of the Federal Reserve System. Ideally, the maturity of the rates should equal the maturity of the futures contract. Although these data are not available, Thomas (1992) concluded that this approximation works well over the same sample period. Note that the interest rates are converted to continuously compounded rates.

References

Admati, A., and P. Pfleiderer, 1988, "A Theory of Intraday Patterns: Volume and Price Variability," *Review of Financial Studies*, 1, 3–40.

Admati, A., and P. Pfleiderer, 1989, "Divide and Conquer: A Theory of Intraday and Day-of-the-Week Mean Effects," *Review of Financial Studies*, 2, 189–223.

Atchison, M., K. Butler, and R. Simonds, 1987, "Nonsynchronous Security Trading and Market Index Autocorrelation," *Journal of Finance*, 42, 533–553.

Berglund, T., and E. Liljeblom, 1988, "Market Serial Correlation on a Small Security Market: A Note," *Journal of Finance*, 43, 1265–1274.

Bessembinder, H., and M. Hertzel, 1993, "Return Autocorrelations around Nontrading Days," *Review of Financial Studies*, 6, 155–189.

Blume, M., and R. Stambaugh, 1983, "Biases in Computed Returns: An Application to the Size Effect," *Journal of Financial Economics*, 12, 387–404.

Brock, W., J. Lakonishok, and B. LeBaron, 1992, "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns," *Journal of Finance*, 47, 1731–1764.

Cohen, K., et al., 1986, *The Microstructure of Security Markets*, Prentice-Hall, Englewood Cliffs, N.J.

Connolly, R., and J. Conrad, 1991, "Cointegration, Error-Correction, and the Lagged Crosscorrelation Structure of Security Prices," working paper, University of North Carolina.

Conrad, J., M. Gultekin, and G. Kaul, 1991, "Asymmetric Predictability of Conditional Variances," *Review of Financial Studies*, 4, 597–622.

Conrad, J., M. Gultekin, and G. Kaul, 1992, "Profitability of Short-Term Contrarian Profit Strategies: Implications for Market Efficiency," working paper, University of Michigan and University of North Carolina.

Conrad, J., and G. Kaul, 1988, "Time Varying Expected Returns," *Journal of Business*, 61, 409–425.

Conrad, J., and G. Kaul, 1989, "Mean Reversion in Short-Horizon Expected Returns," *Review of Financial Studies*, 2, 225–240.

Conrad, J., G. Kaul, and M. Nimalendran, 1991, "Components of Short-Horizon Individual Security Return," *Journal of Financial Economics*, 29, 365–384.

Fama, E., 1965, "The Behavior of Stock Market Prices," *Journal of Business*, 30, 34–105.

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

Foerster, S., and D. Keim, 1993, "Non-trading, Bid-Ask Spreads, and Daily Return Autocorrelations," working paper, University of Pennsylvania.

Gibbons, M., and W. Ferson, 1985, "Tests of Asset Pricing Models with Changing Expectations and an Unobservable Market Portfolio," *Journal of Financial Economics,* 14, 217–236.

Hameed, A., 1992, "Time Varying Factors and Cross-Autocorrelation in Short-Horizon Stock Returns," working paper, University of North Carolina.

Hansen, L. P., and R. J. Hodrick, 1980, "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis," *Journal of Political Economy,* 88, 829–853.

Harris, L., 1989, "The October 1987 Crash S&P 500 Stock-Futures Basis," *Journal of Finance,* 44, 77–99.

Hasbrouck, J., 1991, "The Summary Informativeness of Stock Trades: An Econometric Analysis," *Review of Financial Studies,* 4, 571–595.

Hawawini, G., 1980, "Intertemporal Cross Dependence in Securities Daily Returns and the Short-Term Intervailing Effect on Systematic Risk," *Journal of Financial and Quantitative Analysis,* 15, 139–149.

Jegadeesh, N., 1990, "Evidence of Predictable Behavior of Security Returns," *Journal of Finance,* 45, 881–898.

Jegadeesh, N., and S. Titman, 1992, "Overreaction, Delayed Reaction, and Contrarian Profits," working paper, UCLA.

Keim, D., 1989, "Trading Patterns, Bid-Ask Spreads, and Estimated Security Returns: The Case of Common Stocks at Calendar Turning Points," *Journal of Financial Economics,* 25, 75–98.

Keim, D., and R. Stambaugh, 1984, "A Further Investigation of the Weekend Effect in Stock Returns," *Journal of Finance,* 39, 819–835.

Keim, D., and R. Stambaugh, 1986, "Predicting Returns in the Stock and Bond Markets," *Journal of Financial Economics,* 17, 357–390.

Lehmann, B., 1990, "Fads, Martingales, and Market Efficiency," *Quarterly Journal of Economics,* 105, 1–28.

Lo, A., and C. MacKinlay, 1988, "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test," *Review of Financial Studies,* 1, 41–66.

Lo, A., and C. MacKinlay, 1990a, "When are Contrarian Profits Due to Stock Market Overreaction?" *Review of Financial Studies,* 3, 175–205.

Lo, A., and C. MacKinlay, 1990b, "An Econometric Analysis of Nonsynchronous Trading," *Journal of Econometrics,* 45, 181–211.

Lo, A., and C. MacKinlay, 1992, "Maximizing Predictability in the Stock and Bond Markets," working paper, MIT.

MacKinlay, C., and K. Ramaswamy, 1988, "Index-Futures Arbitrage and the Behavior of Stock Index Futures Prices," *Review of Financial Studies,* 1, 137–158.

McInish, T., and R. Wood, 1991, "Autocorrelation of Daily Index Returns: Intraday-to-Intraday versus Close-to-Close Intervals," *Journal of Banking and Finance,* 15, 193–206.

Mech, T., 1990, "The Economics of Lagged Price Adjustment," working paper, Boston College.

Mech, T., 1993, "Portfolio Return Autocorrelation," *Journal of Financial Economics,* 34, 307–344.

Muthuswamy, J., 1988, "Asynchronous Closing Prices and Spurious Correlation in Portfolio Returns," working paper, University of Chicago.

Newey, W., and K. West, 1987, "A Simple, Positive Semi-Definite, Heteroscedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica,* 55, 703–708.

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

Richardson, M., and T. Smith, 1991, "Tests of Financial Models in the Presence of Overlapping Observations," *Review of Financial Studies,* 4, 227–254.

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

Thomas, S., 1992, "The Behavior of Prices in the Value Line Stock Index Futures Market under Both Versions of the Spot Index," Ph.D. dissertation, Wharton School, University of Pennsylvania.