
NEW TRADING PRACTICES AND SHORT-RUN MARKET EFFICIENCY *

KENNETH A. FROOT
ANDRÉ F. PEROLD

INTRODUCTION

Many observers claim that new institutional trading practices have fundamentally altered the behavior of stock prices. (Examples of such practices include index arbitrage and other portfolio trading strategies—strategies which involve simultaneous trades in many securities.) Some argue that these practices have made market prices better, by helping them more rapidly reflect market-wide information. Others hold that the new practices have made the markets worse by increasing volatility, particularly at short horizons. A third group—consisting primarily of academic researchers—argues that there is little evidence of an important shift in the statistical properties of stock returns.¹

*This article is a much expanded version of a paper prepared for the NYSE's Market Volatility and investor Confidence Panel, coauthored with Jim Gammill. We thank Larry Harris, Jim Shapiro, Eric Sirri, and Jeremy Stein for helpful discussions, Phil Hamilton for efficient research assistance, and the Division of Research at Harvard Business School and the Q-Group for generous research support.

¹Most attention to date has focused on the volatility of stock returns. But there is little evidence of a recent upward shift in volatility. Harris (1989), for example, studies the recent short-run volatility of the S&P 500, and finds a statistically significant, but economically trivial, rise in the conditional volatility of returns. However, because there is no accepted equilibrium model of volatility, it is hard to know whether an increase—even if it could be convincingly established—is for better or for worse. See Schwert (1989) for an investigation of the statistical properties of return volatilities over longer sample periods.

-
- *Kenneth A. Froot is a Professor of Business Administration at Harvard University.*
 - *André F. Perold is a Professor of Business Administration at Harvard University.*

This article examines one particular property of stock-index returns—short-run autocorrelation—and finds a dramatic decline in recent years. For example, over the 1983–1989 period, 15-minute returns on the S&P 500 went from being highly positively serially correlated (with an autocorrelation coefficient of about 0.4) to practically uncorrelated.² Over the past 20 years, daily and weekly autocorrelations on the Dow-Jones, S&P 500, and NYSE value-weighted indices have fallen markedly also. (The autocorrelation of equally weighted NYSE returns has fallen too, but not as dramatically.) The data show that positive index autocorrelation found in earlier studies [e.g., Lo and MacKinlay (1988), Poterba and Summers (1988)] was a result of high autocorrelation during the 1960s and 1970s, and that it had vanished completely by the late 1980s.

Several explanations for these declines in autocorrelation are considered:

- i. That the dissemination of market-wide information has improved;
- ii. That the tendency for stock prices to overreact has increased;
- iii. That the bid–ask bounce component of the return index has increased; and
- iv. That staleness in prices due to nontrading has decreased.

To motivate the first hypothesis, a simple model is provided to show that slow dissemination of market-wide information results in index returns which are sluggish, in that they exhibit positive autocorrelation and relatively low variance. This occurs even when information about individual stocks is processed efficiently. It is shown that the index's theoretical autocorrelation falls and that its volatility (over short return horizons) rises with an increase in the speed at which market-wide information is disseminated. Such increases in the speed of dissemination are shown to alter cross-stock moments, but need have no effect on the autocorrelation or variance of individual stock returns.

Hypothesis *ii* has different implications for own- versus cross-stock autocorrelations. If overreaction has increased, one might expect to find a decline in own- as well as cross-autocorrelation. The data suggest that this is, in fact, not the case: while cross-autocorrelations have

²Consistent with the decline in autocorrelation, it is found also that from 1983 to 1989 the variance of short-run index returns rose steadily by almost 50% relative to the variance of longer-horizon returns.

declined, own-autocorrelations have actually *risen* during the 1980s. This finding seems at odds with the argument that the new trading practices have brought with them an increased tendency toward short-run overreaction.³ The overreaction hypothesis is also inconsistent with this finding that higher-order index autocorrelations, which were formerly statistically negative, have *risen* to become indistinguishable from zero.

The decline in index autocorrelation could, in principle, be explained by measurement problems associated with increases in bid-ask bounce or decreases in nontrading effects (hypotheses *iii* and *iv*, respectively). Transactions data on individual NYSE stocks are used to estimate the importance of these alternatives. Index returns, which are based on last-trade prices, are decomposed into bid-ask bounce, nontrading, and current midquote components. The data show that the first two of these components explain very little of the decline in autocorrelation. Moreover, increases in trading volume (i.e., decreases in nontrading staleness) appear to have *increased* measured index autocorrelation.

This study concludes that the autocorrelation found in early short-term index returns appears to have been due to inefficient processing of market-wide information, and that recent technological and institutional improvements in the processing of this information has removed much of the autocorrelation.⁴

The rest of the article is organized as follows. The second section provides a simple model which relates the speed of dissemination of market-wide information to the autocorrelation and variance of returns. The third section explores the decline in autocorrelation in 15-minute returns on the S&P 500. The fourth section is devoted to interpreting the findings by performing the decomposition mentioned above. The fifth section then looks at the historical behavior of daily- and weekly-return autocorrelations and reports evidence of a similar secular decline in autocorrelation.

A SIMPLE MODEL

A simple model suffices to demonstrate how reductions in transaction costs and improvements in information technology can affect the be-

³See Lo and MacKinlay (1990) for evidence of the importance of cross-stock effects in generating predictable index returns.

⁴Cutler, Poterba, and Summers (1990), among others, have claimed that positive short-run autocorrelations might result from time-varying required returns. The results of this study suggest a very different interpretation.

havior of index returns. Imagine that the market consists of N stocks, each of which is managed by a risk-neutral specialist. Suppose that the true value of the i th stock at time t is given by V_t^i , which is defined as the sum of a market-wide "factor," V_t , plus an idiosyncratic value term, ξ_t^i : $V_t^i \equiv V_t + \xi_t^i$. For simplicity, assume that the components of V_t^i follow independent random walks, and the mean-zero innovations $\Delta V_t = u_t$, and $\Delta \xi_t^i = e_t^i$ are iid normal, with variance σ_u^2 and σ_e^2 , respectively.

To capture the notion that trading costs and technological delays hamper the dissemination of information, assume that the specialist cannot observe V_t^i instantly, but must wait until time $t + 1$ to observe V_t^i (and its components). In the spirit of Kyle (1985), assume that the specialist also observes at time t an order flow which is comprised of an informed traders' component, here given simply by the change in true value, $u_t + e_t^i$, plus a random component from "liquidity" traders, ν_t^i :

$$F_t^i = u_t + e_t^i + \nu_t^i \quad (1)$$

with ν_t^i iid normal (both across time and over stocks) and with zero mean and variance σ_ν^2 .⁵ Thus, at time t , the i th specialist observes his own private order flow, F_t^i , plus the components of true value of the i th stock at time $t - 1$, V_{t-1}^i . Those informed traders who observe u_t (and therefore V_t) contemporaneously are referred to as index traders to distinguish them from traders who observe stock-specific information, e_t^i .

If specialists set time- t prices optimally, according to their current conditional expectation of V_t^i , it is easy to show that the price of the i th stock at time t is just:

$$P_t^i = \lambda F_t^i + V_{t-1}^i \quad (2)$$

where $\lambda = (\sigma_u^2 + \sigma_e^2)/(\sigma_u^2 + \sigma_e^2 + \sigma_\nu^2)$.⁶ The change in price between times t and $t - 1$ is then:

$$\Delta P_t^i = \lambda(u_t + e_t^i + \Delta \nu_t^i) + (1 - \lambda)(u_{t-1} + e_{t-1}^i) \quad (3)$$

⁵The results below would continue to hold if one were to derive informed traders' optimal order flow, rather than positing it exogenously. All that matters here is that current information about value is not fully incorporated into prices, which is a general feature of equilibrium models of informed trading.

⁶Given knowledge of V_{t-1}^i and the fact that the random variables are normally and independently distributed, the best unbiased predictor of current price has the linear form in eq. (2), with λ set to the OLS estimator in a regression of $u_t + e_t^i$ (the portion of current value unknown to the specialist) on F_t^i .

Straightforward algebra yields that own-stock price changes are serially uncorrelated, i.e., $\text{cov}(\Delta P_t^i, \Delta P_{t-1}^i) = 0$, a result that follows directly from specialists' optimal choice of the market-depth parameter, λ .

Even though individual price changes are not predictable based on their own past behavior, an index of stock prices is positively autocorrelated. This is because market-wide information is not simultaneously incorporated into the prices of all stocks. To see this, define the change in price of an equally weighted stock index from $t - 1$ to t as $\Delta P_t = N^{-1} \sum_{i=1}^N \Delta P_t^i$. The autocovariance in the index can be written as the sum of the own-stock plus the cross-stock autocovariances:

$$\begin{aligned} \text{cov}(\Delta P_t, \Delta P_{t-1}) = N^{-2} & \left(\sum_{i=1}^N \text{cov}(\Delta P_t^i, \Delta P_{t-1}^i) \right. \\ & \left. + \sum_{i=1}^N \sum_{j \neq i}^N \text{cov}(\Delta P_t^i, \Delta P_{t-1}^j) \right) \end{aligned} \quad (4)$$

As mentioned above, the own-covariance on the right-hand side of eq. (4) is zero, and from eq. (3), the i, j th cross-covariance on the right-hand side of eq. (4) is given by $\lambda(1 - \lambda)\sigma_u^2$. Index autocovariance is therefore:

$$\begin{aligned} \text{cov}(\Delta P_t, \Delta P_{t-1}) &= \frac{(N - 1)}{N} (1 - \lambda)\lambda\sigma_u^2 \\ &= \frac{(N - 1)}{N} \frac{(\sigma_e^2 + \sigma_u^2)\sigma_v^2\sigma_u^2}{(\sigma_u^2 + \sigma_e^2 + \sigma_v^2)^2} > 0 \end{aligned} \quad (5)$$

Thus, even though specialists use all the information available to them to set prices and individual stock returns are serially uncorrelated, an index of returns exhibits positive autocorrelation.

What happens if market-wide information is disseminated more rapidly, so that the lag in observing V_t is reduced? To see this in the model above, imagine that there is a change in market technology such that V_t is observable at time t to specialists. This would occur if it becomes costless to trade V_t , whereupon index traders would earn positive net profits *unless* innovations in V_t are fully incorporated in current prices. Alternatively, a futures market for the index might open and serve as a billboard, making the current value of the index publicly observable.

For either of these reasons, once innovations in V_t are fully incorporated into current prices, the price of the i th stock is given by $P_t^i = \lambda'(F_t^i - u_t) + V_{t-1}^i + u_t$, and the new level of market-depth

by $\lambda' = \sigma_e^2 / (\sigma_e^2 + \sigma_v^2)$.⁷ By substitution, the change in the i th stock's price becomes:

$$\Delta P_t^i = \lambda'(e_t^i + \Delta v_t^i) + (1 - \lambda')e_{t-1}^i + u_t \quad (6)$$

As before, λ' is set such that the own-stock autocovariance is zero. Using eqs. (4) and (6), index autocovariance is now given by:

$$\text{cov}(\Delta P_t, \Delta P_{t-1}) = N^{-1}(N - 1) \text{cov}(u_t, u_{t-1}) = 0 \quad (7)$$

i.e., the cross-stock autocovariance disappears and, hence, the index is serially uncorrelated. Although this model is simple, it demonstrates a very general point: more rapid dissemination of market-wide information lowers the autocorrelation of index returns.⁸

Consider next how faster dissemination of market-wide information affects the variance of index. Simple algebra yields that when market-wide information is observed with a one-period lag, the variance of the index is:

$$\text{var}(\Delta P_t) = N^{-1}(\sigma_u^2 + \sigma_e^2 + \alpha(N - 1)\sigma_u^2) \quad (8)$$

with $\alpha = (1 - 2\lambda(1 - \lambda)) < 1$. Alternatively, if information is instantly disseminated, variance increases to:

$$\text{var}(\Delta P_t) = N^{-1}(\sigma_u^2 + \sigma_e^2 + (N - 1)\sigma_u^2) = \sigma_u^2 + N^{-1}\sigma_e^2 \quad (9)$$

These formulas demonstrate that own-stock variances remain constant at $\sigma_u^2 + \sigma_e^2$ under both regimes, while contemporaneous cross-stock covariances rise from $\alpha\sigma_u^2$ to σ_u^2 when information is disseminated more quickly. This makes intuitive sense, since the decline in autocovariance is exclusively a cross-stock effect. Thus, index variance rises to reflect the compression of market-wide movements.

AUTOCORRELATIONS IN HIGH-FREQUENCY S&P 500 RETURNS

The actual behavior of the variance and serial correlation of short-term returns are explored next. In this and the following sections, the behavior of very short-run returns—15 minutes returns on the S&P 500 cash index from February 1983 to December 1989, are examined.

⁷The specialist subtracts u_t , which is now directly observable, out of the order flow to obtain the best unbiased predictor of e_t^i (the component of current value which he cannot observe).

⁸None of the results depend on the symmetrical nature of the model. Similar conclusions would be reached if, for example, one were to assume that some stock prices react more rapidly to aggregate information than others.

TABLE I
Measures of Variance of 15-Minute S&P 500 Returns ($\times 10^6$)

<i>Year</i>	<i>Intraday Variance</i>	<i>Interday Variance</i>
1983	1.194	1.216
1984	1.469	1.502
1985	0.907	0.918
1986	2.211	2.216
1987 (precrash)	3.344	3.264
1987 (full year)	7.714	7.808
1988	2.776	2.780
1989	1.751	1.727

Notes: Intraday variance measures the average 15-minute return during the trading day, excluding close-open returns. Interday variance includes the close-open return, treating it as though it were another 15-minute interval.

Table I shows the average 15-minute variances for each year from 1983 to 1989.⁹ Each measure of variance is calculated in two ways: the first column reports the average variance during each trading day (beginning at the open and ending at the close), averaged across all trading days in the period; the computation in the second column does the same, but also includes the overnight return between the close and open, treating the overnight as though it were just another 15-minute interval.¹⁰

The main result from Table I is that the level of variance moves around so much year by year that it is difficult to discern an upward trend over this seven-year span. While the variances for 1989 are about 50% above those for 1983, they remain about 25% below the average variances during 1986. Similar results emerge when longer horizon returns are used to compute measures of volatility. This is not strong evidence of a sustained upward trend.¹¹

Tables IIa and IIb show average variances by year and by time of day for both the cash and futures indexes.¹² Table IIa indicates that as much as 25% of an average day's cash-market volatility occurs during the first half-hour of trading. Variances in each year are greatest in

⁹To clarify the effect of the October 1987 crash, two measures are calculated for 1987. The first includes only trading days up until the crash and the second includes the entire year.

¹⁰Note that the latter column is only slightly higher than the former. Indeed, the overnight variance is not much larger than the variance for an average 15-minute interval during the day. If hourly variances remained constant around the clock, the second column would be about three times as large as the first. French and Roll (1986) document that variance per unit time is much lower when the market is closed than when it is open. These data may even exaggerate this effect because, due to nontrading, overnight price changes may get incorporated only slowly into the opening index.

¹¹Harris (1989) studies conditional as well as unconditional variances of S&P returns, and finds that there is an economically small (but statistically significant) increase in recent variance.

¹²Variances are of log returns over each full year for each 15-minute interval of the trading day.

TABLE IIa
Variances of 15-Minute S&P 500 Returns by Year and Time of Day ($\times 10^6$)

Time of Day	Year							
	1983	1984	1985	1986	1987:1	1987	1988	1989
9:45				13.29	23.18	51.59	22.05	10.89
10:00				4.222	6.838	27.92	3.730	2.677
10:15	8.761	9.856	5.667	1.990	3.897	10.46	3.137	1.733
10:30	2.699	2.602	1.422	1.885	2.867	5.160	1.761	4.351
10:45	1.048	1.625	0.947	1.149	2.183	6.069	1.937	1.767
11:00	0.955	0.966	0.598	1.440	2.451	8.949	1.327	1.078
11:15	0.684	0.705	0.571	1.823	1.773	4.748	1.374	1.156
11:30	0.537	1.00	0.463	1.620	1.527	6.771	1.200	1.041
11:45	0.524	0.606	0.415	1.789	1.209	3.807	1.041	0.752
12:00	0.651	0.777	0.437	1.398	1.556	3.157	0.986	0.991
12:15	0.577	0.603	0.365	1.473	1.445	3.352	1.027	0.748
12:30	0.507	0.609	0.536	0.933	2.045	3.306	1.223	0.686
12:45	0.402	0.552	0.413	1.126	1.155	4.847	1.041	0.535
1:00	0.402	0.614	0.238	1.049	1.425	6.063	1.143	0.526
1:15	0.482	0.579	0.559	1.066	1.358	2.443	1.341	0.607
1:30	0.520	0.947	0.450	1.262	1.353	4.052	1.661	0.732
1:45	0.592	0.858	0.488	1.161	1.330	3.780	1.152	0.832
2:00	0.786	0.760	0.480	1.185	2.879	5.883	1.212	0.789
2:15	0.934	1.166	0.892	1.417	1.805	2.858	1.446	0.812
2:30	1.063	1.045	0.658	1.102	1.931	4.799	2.375	1.203
2:45	1.028	1.167	0.855	1.212	3.212	5.736	2.240	1.482
3:00	1.111	1.578	0.901	2.040	2.208	5.766	2.439	1.860
3:15	1.107	2.128	1.109	2.682	4.128	5.444	3.860	2.069
3:30	1.434	1.699	1.457	3.367	4.061	5.753	2.904	2.378
3:45	1.423	2.091	1.246	3.967	5.482	7.636	5.836	2.255
4:00	1.579	1.780	1.174	2.896	3.638	9.581	3.923	1.967
Overnight	0.262	0.987	0.402	1.018	0.708	0.945	1.442	0.341

Notes: Opening times during 1985–1985 were at 10:00 AM, so the first recorded 15-minute return for each day is at 10:15. Opening times were 30 minutes earlier for the rest of the sample. The column entitled 1987:1 includes only trading days before the October crash; the column entitled 1987 includes trading days from the entire calendar year.

the early morning and near the end of the day, remaining uniformly low in between. What is responsible for such large price movements in the morning after the open? One possibility is that staleness in the index results in information that has accumulated overnight to seep only slowly into prices. In such a case, one would expect to see a very different picture in the futures market, where sluggish trading at the open should not be a problem. Thus, one would expect the overnight return variance in the futures index to be greater and the early morning variances smaller than in the cash market.

Table IIb shows that this is indeed the case. The table compares the volatility of the cash and future S&P 500 indexes for 1988 and 1989 (the only years for which such high-frequency futures data were avail-

TABLE IIb
Variances of 15-Minute Returns by Year and Time
of Day for Cash and Futures S&P 500 Indices ($\times 10^6$)

<i>Time of Day</i>	<i>Cash</i>		<i>Futures</i>	
	1988	1989	1988	1989
9:45	15.32	11.35	3.806	4.068
10:00	2.239	2.761	3.539	2.805
10:15	2.658	1.778	3.686	4.093
10:30	1.192	4.662	2.623	4.337
10:45	1.738	1.828	2.896	2.777
11:00	0.995	1.007	1.241	1.486
11:15	0.824	1.206	1.435	2.393
11:30	0.930	1.059	1.640	1.142
11:45	0.964	0.721	2.107	1.182
12:00	1.023	1.034	1.606	1.336
12:15	1.009	0.700	1.970	1.207
12:30	0.961	0.641	1.318	1.367
12:45	0.642	0.518	1.343	0.9345
1:00	0.679	0.535	1.294	1.075
1:15	0.737	0.629	1.844	1.372
1:30	0.939	0.706	1.793	1.109
1:45	0.592	0.765	1.456	1.416
2:00	0.768	0.726	1.544	1.542
2:15	1.031	0.832	2.364	1.605
2:30	1.648	1.269	2.875	1.513
2:45	1.726	1.473	2.694	1.824
3:00	1.689	1.762	3.735	2.616
3:15	1.882	2.113	3.257	2.706
3:30	1.647	2.480	3.027	3.450
3:45	2.757	2.375	4.808	14.24
4:00	3.105	2.087	3.736	1.921
4:15	0.289	0.196	2.375	1.225
Overnight	0.255	0.506	16.72	10.85

Notes: The futures data cover only the period from April 1988 to November 1989. For comparability, the cash index variances above are computed for the same sets of trading days.

able).¹³ As expected, the high early-morning variances evident in the spot market index are conspicuously absent in the futures data, while the overnight futures variances are about ten times as large as those during the day. Interestingly, the variance of the futures index in the middle of the day is consistently *greater* than that of the mid-day cash index.

The overall result evident from Tables I and II is not very exciting: levels of index volatility are too variable to isolate with much confidence any recent increase. However, if it were possible to scale volatility properly, so that noise is eliminated, perhaps these observations would

¹³The data used to construct Table IIb run only from April 1988 until November 1989. As a consequence, the cash-market estimates in Tables IIa and IIb are not identical.

yield more information. One approach would be to scale 15-minute volatility by volatility at some long horizon. To do so, the ratio of 15-minute to weekly volatility is computed and is plotted in Figure 1. The figure shows a definite increase in short-horizon volatility relative to that at longer horizons. This finding is consistent with both the overreaction hypothesis as well as the model of information dissemination in the preceding section.

Has there been a decline in low-order autocorrelation alongside of this increase in variance? To answer this, the intraday variance ratios and first-order autocorrelations are examined further. Table IIIa reports variance ratios comparing the variance of 15-minute returns with the variance of returns at 30, 60, 120, and 180 minutes. If the index follows a random walk so that returns are completely random, each of the variance ratios would be close to 1.0. Numbers higher than 1.0 indicate positive serial correlation in returns.^{14,15}

Also, Table IIIb reports estimates of average first-order autocorrelation coefficients from 15-minute returns. These should be (and indeed are) similar to the 30- to 15-minute ratios in Table IIIa, which are approximately equal to $1 + \rho$, where ρ is the first-order autocorrelation coefficient. Differences between the two measures are due to the different weighting of first and last returns on each day, and are small, though detectable, for these data.¹⁶

¹⁴Table IIIa reports variance ratios measured in several different ways. In the top panel, the ratios are computed for each trading day, and are then averaged over the year. In the second panel, overnight price changes are once again included just as though they were 15-minute returns, and the average across days is reported. The numbers in the top panel are generally lower than those in the second panel, in part because of the behavior of prices at the beginning of each day. Average daily variance ratios, such as those computed in the top panel, will generally be biased downward when a disproportionate share of the day's variance occurs at the beginning of the day. The third and fourth panels are computed analogously to the first and second, except that they omit the opening-return effect by leaving out the first 30 minutes of each day's trading. All of the estimates are corrected for small sample biases. For details on this procedure, see, for example, Cochrane (1988).

¹⁵Standard errors from Monte Carlo simulations under the null hypothesis that returns are independently and identically distributed are reported in each panel. Interestingly, conditional heteroskedasticity does not appear to be a problem in returns over such short intervals. White tests for conditional heteroskedasticity on the actual data are performed. These tests are unable to reject the null hypothesis of no heteroskedasticity. This is in striking contrast to returns for daily intervals, where there is strong evidence of heteroskedasticity conditional on the prior day's returns.

¹⁶In calculating these coefficients and the variance ratios which precede them, expected returns are allowed to vary freely across trading days. While this method imposes no restrictions on expected returns, it does lead to some implausible results (for example, expected returns on some trading days are calculated to be negative). An alternative, but equally extreme method would be to force expected returns to be constant over the entire year. When this is done, the autocovariances are higher by about 0.03, but the change from 1983 to 1988 remains essentially unaffected. Because the return horizon is so short, when expected returns are fixed over the year, just which value is chosen for expected returns makes little difference to the results.




FIGURE 1
Ratio of annualized 15-minute to weekly volatility S&P 500, 1983–1989.

In spite of differences in computational technique, identical conclusions come consistently out of both tables: *there has been a dramatic decline in the high-frequency positive serial correlation present in the index in the 1980s.* Indeed the majority of the initial positive correlation

TABLE IIIa
Variance Ratios Based on 15-Minute S&P 500 Returns

Year	Minutes of Longer Horizon				1 Day
	30	60	120	180	
Panel 1: Averages of intraday ratios					
1983	1.375	1.490	1.574	1.726	
1984	1.220	1.291	1.341	1.416	
1985	1.150	1.216	1.362	1.470	
1986	1.027	1.011	1.033	1.119	
1987 (precrash)	0.956	0.908	0.938	1.046	
1987 (full year)	0.957	0.908	0.946	1.061	
1988	0.967	0.953	0.968	1.077	
1989	0.970	0.966	0.986	1.090	
Simulated standard errors	0.015	0.028	0.044	0.048	

continued

TABLE IIIa (continued)

Year	Minutes of Longer Horizon				1 Day
	30	60	120	180	
Panel 2: Interday ratios					
1983	1.677	1.817	1.954	2.249	2.249
1984	1.465	1.638	1.782	2.081	2.081
1985	1.293	1.429	1.519	1.773	1.773
1986	1.127	1.228	1.295	1.527	1.527
1987 (precrash)	1.103	1.182	1.235	1.338	1.338
1987 (full year)	1.413	1.389	1.459	1.817	1.817
1988	1.150	1.229	1.267	1.384	1.384
1989	1.130	1.260	1.285	1.344	1.344
Simulated standard errors	0.014	0.025	0.039	0.049	
Panel 3: Averages of intraday ratios (excluding first 30 minutes)					
1983	1.463	1.768	1.904	1.979	
1984	1.362	1.584	1.690	1.818	
1985	1.193	1.295	1.382	1.374	
1986	1.115	1.167	1.213	1.283	
1987 (precrash)	1.081	1.114	1.158	1.282	
1987 (full year)	1.085	1.120	1.172	1.285	
1988	1.085	1.134	1.173	1.322	
1989	1.102	1.152	1.166	1.280	
Simulated standard errors	0.015	0.028	0.044	0.048	
Panel 4: Interday ratios (excluding first 30 minutes)					
1983	1.237	1.375	1.476	1.588	1.787
1984	1.219	1.373	1.549	1.670	1.911
1985	1.116	1.228	1.360	1.451	1.565
1986	1.070	1.178	1.285	1.360	1.589
1987 (precrash)	1.059	1.134	1.234	1.283	1.379
1987 (full year)	1.161	1.248	1.284	1.354	1.639
1988	1.094	1.184	1.262	1.302	1.420
1989	1.077	1.156	1.274	1.293	1.352
Simulated standard errors	0.015	0.028	0.044	0.048	

Notes: Panel 1 is the average across variance ratios computed for each day (and ignoring overnight returns). Panel 2 treats the overnight return as though it were another 15-minute return. Standard errors are from Monte Carlo experiments, using the null hypothesis that returns are iid. Monte Carlo simulation are run also using several models of conditional heteroskedasticity. None of these result in standard errors importantly different than those reported above. Panels 3 and 4 are comparable to Panels 1 and 2, except that the returns from the first 30 minutes of each day are omitted. Standard errors are from Monte Carlo experiments, using the null hypothesis that returns are iid. Monte Carlo simulations are run also using several models of conditional heteroskedasticity estimated from the actual data. None of these result in standard errors importantly different than those reported above.

has since disappeared. The estimated standard errors—which are less than 0.02—indicate that these changes are highly statistically significant. The largest declines appear to occur in 1985 and 1986, although (with the exception of the crash of 1987) the point estimates have

TABLE IIIb
First-Order Autocorrelation Coefficients based on 15-Minute S&P 500 Returns

<i>Year</i>	<i>ρ</i>
Panel 1: Averages of intraday coefficients	
1983	0.423
1984	0.264
1985	0.197
1986	0.073
1987 (precrash)	0.020
1987 (full year)	0.034
1988	0.038
1989	0.023
Simulated standard errors	0.015
Panel 2: Averages of intraday coefficients, excluding first 30 minutes	
1983	0.446
1984	0.322
1985	0.154
1986	0.102
1987 (precrash)	0.068
1987 (full year)	0.077
1988	0.086
1989	0.088
Simulated standard errors	0.016

Notes: Panel 1 is the average across variance ratios computed for each day (and ignoring overnight returns). Panel 2 is comparable, except that the first 30 minutes of each trading day are omitted. ρ denotes the first-order autocorrelation coefficient of the index returns. Standard errors are from Monte Carlo experiments, using the null hypothesis that returns are conditionally heteroskedastic following a White (1980) model of heteroskedasticity.

continued falling since then. In some cases, there remains currently no statistically significant autocorrelation in the index.¹⁷ Figure 2 graphs the autocorrelations from the top panel of Table IIIb.

Tables IIIa and IIIb may hide a great deal of information by averaging autocorrelations over the day. To look beneath these numbers, Table IV presents evidence on the predictability of consecutive 15-minute returns, showing first-order autocorrelation coefficients by year and time of day. The first two columns of Table IV show that the predictability of upcoming 15-minute returns is very large during 1983 and 1984. The average correlation coefficient in those years is 0.44 and 0.33, respectively. In addition, the degree of predictability is basically constant throughout the day, and not importantly different at the beginning of each day when volatility is greatest.

¹⁷ MacKinlay and Ramaswamy (1988) also report a recent decline in the first-order autocorrelation of index returns.




FIGURE 2

Average daily first-order autocorrelation in 15-minute returns on the S&P 500.

Table IV suggests that the reduction in the predictability of returns is not restricted to some portion of daily trading: essentially all of the daily correlation coefficients fall from their high levels at the beginning of the sample. The steady reduction in autocorrelations appears to be a fairly general feature of the market, and does not appear concentrated in a portion of the trading day.¹⁸

So far, the predictions of subsequent 15-minute returns by current 15-minute returns have been described. How well do current returns forecast price changes which are further into the future? Tables Va and Vb address this issue by reporting higher-order autocorrelation coefficients. To read Table Va, note for example that -0.0122 in the fourth line, first column, represents the autocorrelation between a current 15-minute return and the 15-minute return one hour later.¹⁹ The most readily obvious feature of Table Va is that none of the higher-order autocorrelations are anywhere near as large as the first-order

¹⁸To save space, standard errors for the correlation coefficients in Table IV are not presented. Generally, almost all of the coefficients in excess of 0.2 are statistically different from zero at the 1% level.

¹⁹These autocorrelations treat the overnight return like any other 15-minute interval, and ignore any time-of-day heterogeneity in autocorrelation coefficients.

TABLE IV
Correlation Coefficients of Adjacent 15-Minute
S&P 500 Returns by Year and Time of Day

Daily Interval	Year								Average
	1983	1984	1985	1986	1987:1	1987	1988	1989	
1	0.7317	0.4158	0.4229	0.0570	-0.0552	0.3053	-0.1875	-0.1590	0.1914
2	0.1091	0.2346	0.0455	-0.0483	0.0137	0.5010	0.1059	0.0972	0.1323
3	0.4141	0.5112	0.2815	0.2928	0.0108	0.3009	0.0490	0.0255	0.2357
4	0.4720	0.4024	0.1907	0.2785	0.2374	0.2247	0.2447	-0.1082	0.2428
5	0.6252	0.4019	0.2897	0.1424	0.2670	0.5361	0.3734	0.1703	0.3507
6	0.4938	0.3988	0.3437	0.3769	0.1845	0.1129	0.1987	0.2365	0.2932
7	0.5107	0.4357	0.4180	0.2640	0.2350	0.5632	0.0413	0.0935	0.3202
8	0.6118	0.4552	0.3022	0.2884	0.2299	0.4601	0.2138	0.2200	0.3477
9	0.6509	0.5485	0.4256	-0.1159	0.1384	0.3017	0.2018	0.1475	0.2873
10	0.5037	0.3999	0.3226	-0.0047	0.0265	0.2911	0.0946	0.0132	0.2059
11	0.4939	0.5020	0.2741	0.0827	0.1344	0.3687	0.1344	0.0729	0.2579
12	0.5757	0.3054	0.3457	0.1319	0.1427	0.0501	0.1143	0.1050	0.2214
13	0.5920	0.4144	0.1842	0.0287	0.0559	0.6281	0.2748	0.1684	0.2933
14	0.5812	0.3434	0.1806	0.1220	0.1208	0.3052	0.1305	0.1792	0.2454
15	0.4599	0.2786	0.0729	0.1103	0.1102	-0.0645	0.3998	0.1701	0.1922
16	0.5624	0.3524	0.2358	0.2055	0.0943	0.3272	-0.0931	0.1543	0.2298
17	0.5817	0.3561	0.0741	0.0456	0.0181	0.3173	0.0387	0.1068	0.1923
18	0.5247	0.3653	0.2594	0.1650	-0.1188	0.0467	0.1380	0.0223	0.1753
19	0.5145	0.5271	0.2456	-0.0590	-0.0634	-0.0678	0.2542	0.1576	0.1886
20	0.3593	0.2342	0.1351	0.1333	0.1750	0.2198	0.1020	-0.0449	0.1642
21	0.5498	0.3434	0.0923	0.1913	0.2503	0.1774	0.0603	-0.1530	0.1890
22	0.3684	0.3133	0.2992	0.0701	-0.0683	0.2549	0.1666	0.3822	0.2233
23	0.5905	0.4243	0.2609	0.0391	-0.0194	0.2542	0.0653	0.1555	0.2213
24				0.0101	0.1347	0.2543	0.1746	-0.0238	0.1247
25				0.0891	0.2137	0.4167	0.2736	0.3958	0.1339
26	0.3591	0.1327	0.3151	0.2773	0.0981	0.3290	0.0599	0.1366	0.1126
27	0.1296	-0.0667	-0.2514	-0.2572	-0.0910	-0.0632	0.2193	0.3241	0.0165
Average	0.4399	0.3344	0.2136	0.1080	0.0917	0.2723	0.1425	0.1128	0.2144

Notes: Line numbers 1–25 indicate the daily time interval of the regressor. For example, line 1 is the correlation coefficient between the second and first (or opening) return on each trading day. In years 1983–1985 the market opened 30 minutes later than in subsequent years; hence, there are two fewer correlation coefficients for 1983–1985. Line 26 is the correlation between the overnight return (close to open) and the return in the last 15 minutes of trading. Line 27 is the correlation between the overnight return and the return in the first 15 minutes of the next day's trading. The column entitled 1987:1 includes only trading days before the October crash; the column entitled 1987 includes trading days from the entire calendar year.

autocorrelations. In 1983, for example, the second-order coefficient is 0.038, an order of magnitude lower than the first-order estimate (but still statistically positive). The second-order autocorrelations have also fallen over time. Indeed, all the estimates after 1984 (with the exception of the subsample which includes the 1987 crash) are statistically indistinguishable from zero.

TABLE Va
 Serial Correlation Coefficients of 15-Minute S&P 500 Returns at Longer Lags by Year

Lag Number	Year							
	1983	1984	1985	1986	1987:1	1987	1988	1989
1	0.4327	0.2942	0.1884	0.0741	0.0558	0.2617	0.0942	0.0667
2	0.0381	0.0213	-0.0087	0.0006	-0.0026	0.0328	0.0010	0.0076
3	-0.0221	0.0025	0.0408	0.0370	0.0526	-0.0240	0.0281	0.0450
4	-0.0122	0.0247	0.0430	0.0326	0.0126	-0.0514	0.0342	0.0413
5	0.0178	0.0423	0.0129	0.0005	0.0128	-0.0498	0.0077	0.0251
6	0.0295	0.0434	0.0184	0.0327	0.0025	-0.0148	-0.0040	0.0267
7	0.0437	0.0266	0.0125	0.0172	0.0164	0.0267	0.0142	0.0096
8	0.0385	0.0385	0.0186	-0.0020	0.0242	0.0450	0.0040	-0.0417
9	0.0302	0.0326	0.0135	0.0117	-0.0157	0.0388	-0.0036	-0.0237
10	0.0251	0.0161	0.0156	0.0078	0.0126	0.0554	0.0145	0.0080
11	0.0278	0.0172	0.0216	0.0337	0.0102	0.0554	-0.0033	0.0001
12	0.0310	0.0390	0.0059	0.0158	-0.0167	0.0273	-0.0072	0.0046
13	0.0198	0.0586	0.0289	0.0215	-0.0062	0.0118	-0.0062	-0.0008
14	0.0165	0.0244	0.0253	0.0216	0.0000	0.0304	0.0079	-0.0029
15	0.0138	0.0199	0.0266	-0.0026	0.0105	0.0346	0.0219	0.0231
16	0.0165	0.0094	0.0088	0.0129	0.0121	0.0310	0.0152	0.0327
17	0.0017	-0.0048	0.0171	0.0032	0.0047	0.0052	0.0173	0.0320
18	-0.0291	-0.0034	0.0029	0.0219	0.0097	-0.0195	0.0083	-0.0173
19	-0.0290	-0.0302	-0.0013	0.0214	0.0154	0.0263	0.0403	-0.0212
20	0.0198	-0.0097	0.0048	0.0187	0.0021	0.1070	0.0042	-0.0362
21	0.0571	-0.0004	0.0024	0.0078	0.0401	0.0803	0.0079	0.0261
22	0.0484	-0.0010	0.0058	0.0268	-0.0033	-0.0014	0.0114	-0.0148
23	0.0446	-0.0207	0.0148	-0.0066	-0.0046	-0.0621	0.0043	0.0011
24	0.0246	-0.0169	0.0132	-0.0106	0.0083	-0.0480	-0.0038	-0.0287
25	-0.0289	-0.0152	-0.0007	0.0145	-0.0122	-0.0503	0.0121	0.0279
26	-0.0360	0.0141	-0.0072	-0.0051	-0.0053	-0.0186	-0.0102	-0.0105
27	-0.0338	0.0040	-0.0027	0.0015	-0.0028	-0.0086	-0.0200	0.0106
28	-0.0436	-0.0244	-0.0251	-0.0065	-0.0200	-0.0062	-0.0277	-0.0100
29	-0.0486	-0.0214	-0.0285	-0.0060	0.0011	0.0110	0.0023	-0.0025
30	-0.0207	-0.0563	-0.0089	-0.0094	-0.0292	0.0022	-0.0396	-0.0242
31	0.0049	-0.0294	0.0085	-0.0228	-0.0028	0.0293	-0.0022	-0.0271
32	-0.0153	-0.0174	-0.0038	-0.0046	-0.0223	0.0132	-0.0112	0.0137
33	-0.0367	-0.0224	-0.0130	-0.0119	-0.0196	-0.0070	-0.0081	0.0047
34	-0.0248	-0.0271	0.0269	-0.0158	-0.0222	-0.0304	0.0081	0.0063
35	-0.0206	-0.0128	-0.0197	0.0020	0.0142	-0.0152	-0.0128	-0.0031
36	-0.0101	-0.0006	-0.0248	-0.0145	-0.0060	-0.0203	-0.0206	0.0062
37				-0.0206	-0.0227	-0.0297	-0.0007	0.0116
38				-0.0306	0.0099	-0.0149	0.0123	0.0137
39				0.0100	-0.0023	-0.0196	-0.0068	0.0135

Notes: Standard errors of these coefficients are approximately 0.012. The column entitled 1987:1 includes only trading days before the October crash; the column entitled 1987 includes trading days from the entire calendar year.

To help digest the information in Table Va, a summary of the coefficients is presented in Table Vb. To do this, the coefficients are averaged over half-day intervals. Thus, for example, the second line in

TABLE Vb

Summary of S&P 500 Serial Correlation Coefficients at Longer Lags by Year

Lag Numbers	Year							
	1983	1984	1985	1986	1987:1	1987	1988	1989
1	0.4327	0.2942	0.1884	0.0741	0.0558	0.2617	0.0942	0.0667
First half-day	0.0225	0.0277	0.0177	0.0174	0.0085	0.0128	0.0066	0.0085
Second half-day	0.0171	0.0021	0.0124	0.0095	0.0060	0.0088	0.0105	0.0009
Third half-day	-0.0262	-0.0174	-0.0083	-0.0099	-0.0096	-0.0074	-0.0098	0.0010

Notes: Standard errors of the last three rows are approximately 0.004. Half-days are equivalent to 12 15-minute return intervals during 1983–1985 and to 13 15-minute return intervals during 1986–1989. The column entitled 1987:1 includes only trading days before the October crash; the column entitled 1987 includes trading days from the entire calendar year.

Table Vb represents the average of the second- to 12th-order correlation coefficients from Table Va. Because the first-order coefficient is distinctly large, it is not included in this average.²⁰

Focus first on the pattern of the estimates for 1983. The first-half-day coefficient is smaller than the first-order coefficient, and the average second-half-day coefficient is smaller still (although it is still statistically significant). The third-half-day coefficient drops further and is actually statistically negative. Next, notice that this pattern disappears slowly over time. By 1988 and 1989, the average higher-order coefficients show no real downward trend and none remain statistically different from zero. Thus, the decline toward zero in first-order autocorrelations seems to occur in higher-order autocorrelations as well. These results are inconsistent with the view that new trading practices have led to short-term overreactions.

INTERPRETING CHANGES IN THE PREDICTABILITY OF THE CASH INDEX

So far, this study has concentrated on the decline in autocorrelation in the reported index which is based on last-trade prices. However, these prices include measurement errors due to bid-ask bounce and nontrading effects, which could in principle account for the decline in autocorrelation. In this section, an attempt is made to measure the contribution of these two sources of measurement problems, and to isolate the portion of the decline that is generated by more efficient

²⁰For the years 1986–1989, there are 26 15-minute intervals in each trading day (as compared with 24 intervals during 1983–1985), so that the half-day averages include an extra coefficient during this period.

processing of market-wide information. To do this, transactions data for individual stocks are used which allow identification of own- and cross-stock components of the decline and insight into how market-wide information is actually disseminated. This is important for distinguishing the overreaction hypothesis from the faster-dissemination-of-information hypothesis.

Bid-Ask Bounce

The first source of measurement error is bid-ask bounce. When there are discrete differences in the prices at which buys and sells are executed, random buys and sells may lead to the appearance of upward and downward movements in prices, even when quoted prices are constant over time. This component of price changes will exhibit negative serial correlation: when the index is at the ask, all else equal, it tends, on average, to move down toward the bid.²¹ If bid-ask bounce is present in the last-trade index, then the *level* of autocorrelation coefficients is lower as a result. What is important for this analysis, however, is whether the bounce can explain the *change* in autocorrelation through time.

There are at least two ways that the importance of bid-ask bounce has increased in the 1980s. First, all else equal, bid-ask bounce is an increasing function of the size of the bid-ask spread, so an increase in the spread could produce a corresponding increase in bounce. Second, if investors tend to trade more frequently in portfolios of stocks rather than in individual stocks, then buys and sells will have greater synchronousness and, all else equal, bid-ask bounce will increase. Consider, as an example, the case in which buys and sells across stocks are random, so that at any given time 50% of the stocks are at the bid and 50% at the ask. In such a case, the index would contain only a negligible bounce component, even though bounce may be important for individual stocks. Compare this with the case where buys and sells are perfectly synchronized as a result of portfolio trading, i.e., stocks are simultaneously all at the bid or all at the ask. In this latter case, the synchronousness of buys and sells would create bid-ask bounce and reduce the serial correlation in the last-trade index.

The evidence from Table V suggests that these explanations are unlikely to explain most of the decline in autocorrelation. If bid-ask bounce is responsible for the change of -0.36 (0.07 in 1989 minus

²¹ Roll (1984) presents a simple model of such bid-ask bounce, and shows that bounce induces negative covariation between current and future returns.

0.43 in 1983) in the first-order autocorrelation, one would expect an equal-size reduction in *all* correlation coefficients, and not just that of the first-order (see Roll, 1984). It is clear from Table V that the change in the first-order coefficient is more than an order of magnitude greater than changes in higher-order coefficients. However, there is some evidence that bid-ask bounce has increased slightly. From Table Va, it is evident that the second-order autocorrelation coefficient falls by -0.030 , from 0.038 in 1983 to 0.008 in 1989. Similarly, from Table Vb, the average autocorrelation coefficient over the first 24 hours falls by about -0.016 over the same period. Both of these changes are statistically significant. But even if one supposes that they are due entirely to increases in bid-ask bounce, they are clearly too small to explain the overall change in autocorrelation of the index.

Transactions Data

A second, more direct, piece of evidence on bid-ask bounce can be obtained by attempting to isolate the bid-ask component of the last-trade index. To do this, data for all NYSE transactions for the years 1983 and 1988 are examined. Working with as large a subset of the S&P 500 as possible, an approximate S&P 500 last-trade index and corresponding indexes of bid and ask prices prevailing just before the last trade are constructed.²²

Let L_t be the index of last-trade prices and M_t be the index of extant midquotes for each last-trade price. The difference between the two, $L_t - M_t$, measures the distance between the last trade index and the center of the then-prevailing spread. Now let $l_t = \ln(L_t/L_{t-1})$ be the return index of last-trade prices, and $m_t = \ln(M_t/M_{t-1})$ be the last-trade-midquote return index. In principle, the last-trade-midquote index is not contaminated by a Roll-type bid-ask spread. Therefore, the importance of bounce can be learned by exploring the changes in serial correlation of l_t and m_t .

²²The intraday transactions database is from the Center for the Study of Security Prices. Only those stocks in the S&P 500 whose primary market is the NYSE are included. Only those transactions and specialist quotes which were reported on the NYSE are considered. This is done to minimize complications arising from quotation and trade reporting standards that vary between markets. Stocks are excluded on days when there are apparent data errors, or on days when quote and price data are available only after the first 30 minutes. Certain whole days are excluded due to gaps in the data. Indices are computed for 236 days in 1983 and for 252 days in 1988. The number of stocks varies between 269 and 430 in 1983, and between 374 and 455 in 1988. By leaving out the first hour of each day's trading, the number of includable stocks increases. However, doing so has no material effect on any of the estimates.

TABLE VI
Decomposition of Last-Trade Index Returns

Variable	1983	1988	Change
Averages of intraday ratios, excluding first 30 minutes			
1. $\text{cov}(l_t, l_{t-1})/\text{var}(l_t)$	0.422	0.090	-0.332
$\sum_{i=1}^N \omega^i \omega^j \text{cov}(l_t^i, l_{t-1}^j)/\text{var}(l_t)$	-0.026	-0.011	0.015
$\sum_{i=1}^N \sum_{j \neq i} \omega^i \omega^j \text{cov}(l_t^i, l_{t-1}^j)/\text{var}(l_t)$	0.448	0.101	-0.347
2. $\text{cov}(m_t, m_{t-1})/\text{var}(l_t)$	0.389	0.168	-0.221
$\sum_{i=1}^N \omega^i \omega^j \text{cov}(m_t^i, m_{t-1}^j)/\text{var}(l_t)$	-0.006	-0.003	0.003
$\sum_{i=1}^N \sum_{j \neq i} \omega^i \omega^j \text{cov}(m_t^i, m_{t-1}^j)/\text{var}(l_t)$	0.395	0.171	-0.224
3. $(\text{cov}(l_t, l_{t-1}) - \text{cov}(m_t, m_{t-1}))/\text{var}(l_t)$	0.033	-0.078	-0.111
$\sum_{i=1}^N \omega^i \omega^j (\text{cov}(l_t^i, l_{t-1}^j) - \text{cov}(m_t^i, m_{t-1}^j))/\text{var}(l_t)$	-0.020	-0.009	0.011
$\sum_{i=1}^N \sum_{j \neq i} \omega^i \omega^j (\text{cov}(l_t^i, l_{t-1}^j) - \text{cov}(m_t^i, m_{t-1}^j))/\text{var}(l_t)$	0.053	-0.069	-0.122
4. $\text{cov}(\epsilon_t, \epsilon_{t-1})/\text{var}(l_t)$	-0.035	-0.086	-0.051
$\sum_{i=1}^N \omega^i \omega^j \text{cov}(\epsilon_t^i, \epsilon_{t-1}^j)/\text{var}(l_t)$	-0.040	-0.016	0.024
$\sum_{i=1}^N \sum_{j \neq i} \omega^i \omega^j \text{cov}(\epsilon_t^i, \epsilon_{t-1}^j)/\text{var}(l_t)$	0.005	-0.070	-0.075
5. $\text{cov}(m_t, \epsilon_{t-1})/\text{var}(l_t)$	0.146	0.126	-0.020
$\sum_{i=1}^N \omega^i \omega^j \text{cov}(m_t^i, \epsilon_{t-1}^j)/\text{var}(l_t)$	0.017	0.006	-0.011
$\sum_{i=1}^N \sum_{j \neq i} \omega^i \omega^j \text{cov}(m_t^i, \epsilon_{t-1}^j)/\text{var}(l_t)$	0.129	0.120	-0.009
6. $\text{cov}(\epsilon_t, m_{t-1})/\text{var}(l_t)$	-0.077	-0.118	-0.041
$\sum_{i=1}^N \omega^i \omega^j \text{cov}(\epsilon_t^i, m_{t-1}^j)/\text{var}(l_t)$	0.003	0.002	-0.001
$\sum_{i=1}^N \sum_{j \neq i} \omega^i \omega^j \text{cov}(\epsilon_t^i, m_{t-1}^j)/\text{var}(l_t)$	-0.080	-0.120	-0.040

Notes: Indexes are constructed to approximate the S&P 500, using NYSE stocks only. See footnote 21 in the text for more details. l_t represents the last-trade return index, m_t the last-trade-midquote return index, m_t^i the current midquote return index, and $\epsilon_t = l_t - m_t$ is measurement error introduced by the bid-ask spread.

Table VI reports estimates of the first-order autocorrelations of the indexes that are constructed from the transactions data: l_t and m_t . Table VI provides average intraday autocorrelations excluding the first 30 minutes of each day's trading. Estimates of the own-autocorrelation of the total index are reported below each number. Since total covariance is the sum of own-covariance plus cross-covariance, the own-covariances allow assessment of the amount of the change in autocovariance that is attributable to cross-stock effects. For example, the first line in Table VI shows the change in the autocorrelation of the l_t index from 1983 to 1988 of -0.332.²³ Of this, the number beneath line 1 says that 0.015

²³Note that this number is close (but not precisely equal) to the autocorrelations of the S&P 500 reported in Table IIIb. Discrepancies are due to the differences in the way the indexes are calculated. See footnote 20 above.

is attributable to a decline in own-autocovariance (i.e., an increase in $\sum_{i=1}^N (w_i)^2 \text{cov}(l_t^i, l_{t-1}^i) / \text{var}(l_t)$, where w^i is the weight of the i th stock in the index) and that the remaining -0.347 is attributable to a decline in cross-covariance ($\sum_{j=1}^N \sum_{i \neq j} w^j w^i \text{cov}(l_t^j, l_{t-1}^i) / \text{var}(l_t)$). As discussed earlier, if the decline in index autocorrelation is due predominantly to better processing of market-wide information, one would expect changes in cross-covariance to explain most of the decline.

The second line of Table VI reports the first-order autocovariance of m_t divided by the variance of the last-trade index, $(\text{cov}(m_t, m_{t-1})) / (\text{var}(l_t))$.²⁴ This index explains a change of -0.221 , or two thirds of the decline in the autocorrelation of l_t . On the third line is the difference between the first two lines, which is an explicit measure of bid-ask effects. It is clear from line 3 that the bid-ask component has risen by more than estimated from Table V: the change from 1983 to 1988 is -0.111 . Nevertheless, this is still about one third of the change in the autocorrelation of the last trade index.

It would be wrong, however, to interpret differences between the autocovariances of the last-trade and last-trade-midquote indices on line three as pure measures of bounce. To see why, define the bid-ask error as the difference between the last-trade and last-trade-midquote return indices: $\epsilon_t \equiv l_t - m_t$, which is approximately the change in $L_t - M_t$ (expressed as a percent of M_t).²⁵ If $\epsilon_t > 0$, then, loosely speaking, there is a greater fraction of buys at time t than at time $t - 1$. Then rewrite the third line of Table VI as:

$$\begin{aligned} \frac{\text{cov}(l_t, l_{t-1})}{\text{var}(l_t)} - \frac{\text{cov}(m_t, m_{t-1})}{\text{var}(l_t)} &= \frac{\text{cov}(\epsilon_t, \epsilon_{t-1})}{\text{var}(l_t)} + \frac{\text{cov}(m_t, \epsilon_{t-1})}{\text{var}(l_t)} \\ &+ \frac{\text{cov}(\epsilon_t, m_{t-1})}{\text{var}(l_t)} \end{aligned} \quad (10)$$

Equation (10) indicates that the difference between the autocorrelations can be subdivided into three parts. Each of these turns out to have a distinct interpretation.

²⁴Note that this is just the first-order autocorrelation of m_t multiplied by $(\text{var}(m_t)) / (\text{var}(l_t))$. This expression is used rather than the simple autocorrelation of m_t because it has the same denominator as the autocorrelation of l_t , and is, therefore, amenable to additive decomposition. These ratios are all calculated daily, then averaged over the year.

²⁵Using the notation from above, the relationship between ϵ_t and deviations from the bid-ask midpoint are:

$$\epsilon_t = \ln(L_t / L_{t-1}) - \ln(M_t / M_{t-1}) = \ln\left(\frac{1 + \xi_t}{1 + \xi_{t-1}}\right) \approx \xi_t - \xi_{t-1}$$

where $\xi_t = (L_t - M_t) / M_t$.

The first term on the right-hand side of eq. (10) is a direct measure of the serial correlation induced by synchronized buy and sell orders. That is, it is a pure measure of Roll-type bounce. The fourth line in Table VI reports that the change from 1983 to 1988 in $(\text{cov}(\epsilon_t, \epsilon_{t-1})/(\text{var}(l_t)))$ over time is negative, but relatively small at -0.051 . Notice, however, that the move in the own-autocovariance in line 4 is actually *positive*, with an increase of about 0.024 . This indicates that bid-ask bounce in individual stocks has become less important (rising over time toward zero), which suggests that the average stock's bid-ask spread has narrowed. This implication is checked by computing the average bid-ask spread, and it is found that the average spread indeed fell from 1983 to 1988 from 19.5 basis points to 17.1 basis points.

These facts also imply that cross-correlation in bid-ask bounce has become (more) negative (falling by $-0.051 - 0.024 = -0.075$). Such a decline would follow from an increase in the synchronousness of buy and sell transactions across stocks. Note that this is exactly what one would expect if portfolio trading has increased over time. In any case, the own- and cross-components of bid-ask bounce are small in size, and, moreover, tend to cancel. Overall, bid-ask bounce is, therefore, responsible for only a tiny part of the change in the correlation of the last-trade index.

The second term on the right-hand side of eq. (10) is slightly more complex. It measures the correlation between past increases in buys (sells) and current increases (decreases) in the midquote index. Table VI reports this term in line 5. The correlation is clearly large and positive. In addition, these estimates move toward zero over time, and account for 20% of the unexplained change reported in the last column of line 3. There are two possible explanations for such behavior. One is termed the eating-through-the-order-book (ETOB) and the other is called the sluggish-response-to-order-flow (SRFI) hypothesis.

The ETOB hypothesis describes a market where order flow is positively autocorrelated and limit orders are sticky. Suppose, for example, that the specialist has a limit order at the ask price, with more limit orders at prices above that. Suppose also that as buy orders come in and the specialist executes and (eventually) exhausts the current limit order, he simply moves the ask price up to the next limit order. When order flow is positively autocorrelated (which might occur if big trades are broken up and executed sequentially), an increase in buy orders tends to forecast an increase in the ask price, and, therefore, an increase in the future midquote index. The ETOB

hypothesis would, therefore, predict that the covariance between ϵ_{t-1} and m_t is positive.²⁶

The other possibility to explain line 5 is the SRFI hypothesis. This posits that order-flow information for a given stock is incorporated into quotes for other stocks slowly over time. For example, suppose that at time $t - 1$, the GM specialist executes a buy order at the ask (which increases ϵ_{t-1}). This buy order might provide incremental information about the value of Ford, and, therefore, might be associated with an increase in the Ford specialist's quotes. The SRFI hypothesis says that the full increase happens not instantaneously, but slowly through time. Thus, the positive covariance between ϵ_{t-1} and m_t .

The important difference between these two hypotheses is that SRFI is a measure of how rapidly market-wide information is disseminated, and is, therefore, central to the point of this study. How can one distinguish between these two hypotheses? One way is to observe that SRFI is clearly a statement about correlation of ϵ_{t-1} and m_t across stocks, while ETOB is an *own*-stock effect. Using the own-covariance numbers in Table VI, one can separate out the cross-stock component in the last columns of line 5. The estimates imply that of the change in line 5 of -0.020 , about -0.009 ($-0.020 + 0.011$) is attributable to cross-stock effects.

This result suggests that, by trying to cleanse the last-trade price index, l_t , of bid-ask bounce, one loses some evidence that the processing of market-wide information has improved. Since the last-trade-midquote index, m_t does not use transactions prices, it ignores the fact that *deviations* from midquotes may be a form of market-wide information, and that this form of improved information dissemination helps explain the reduction in autocorrelation of l_t . This reasoning implies that if SRFI is correct, as it appears to be, one should not attribute the decline in line 5 to an increase in bid-ask bounce, but to improved processing of market-wide information.

Finally, consider the last term on the right-hand side of eq. (10), $(\text{cov}(\epsilon_t, m_{t-1})) / (\text{var}(l_t))$. This term measures the covariance between past increases (decreases) in the midquote index and current buys (sells). Table VI presents the estimates of this term on line 6. The covariance is negative and decreasing over time, accounting for a fall in the autocovariance of the last-trade index of about -0.041 . As with the previous

²⁶See also Glosten and Milgrom (1985), who present a model with the same prediction. In their model, bid and ask rates readjust upward after buys and downward after sells. This creates positive correlation between ϵ_{t-1} and m_t for individual stocks.

term, there are two potential explanations: the “see-’em-coming” (SEC) and “slow-response-to-price-information” (SRPI) hypotheses.

Under the SEC hypothesis, the specialist appears able to anticipate the upcoming order flow, tending to raise (reduce) prices just as buy (sell) orders arrive. This would lead one to expect that bid–ask prices rise as buy orders (locally) peak, and, therefore, that the covariance between m_{t-1} and ϵ_t is negative. Clearly, specialist anticipation of future order flows is not in itself bad for other investors. If specialists are responding to the same information that generates trading in the first place, then SEC may result in better information being incorporated into current prices.²⁷

The alternative—the SRPI hypothesis—holds that some stocks’ quoted prices respond slowly to information, making it attractive to buy or sell them when the index changes. To see this more clearly, suppose that the index is comprised of two stocks: GM, whose quoted prices respond immediately to information, and Ford, whose quoted prices are sticky.²⁸ When positive *market-wide* information is released, GM trades immediately at higher quoted prices, while Ford’s quoted prices remain the same. If there are a few smart traders observing this, they will profit if they buy Ford as the price of GM rises. The buying of Ford subsequently subsides as its price slowly rises. Thus, a current increase in the index of GM and Ford quotes predicts that the index of GM and Ford buys is currently high (and falling).

The SEC hypothesis is an own-stock effect, while the SRPI hypothesis is a cross-stock effect. To see this in the example above, note that the price increase in GM is not associated with current buys of GM, and that the current buying of Ford is assumed not to drive up current Ford quotes. Once again, the estimates show that the own-stock change is essentially zero (-0.005 from the last column of line 6). Thus, the cross-stock effect accounts for most of the change of -0.041 from 1983 to 1988 in line 6. The SRPI hypothesis seems to be the right explanation for the decline in the covariance of m_{t-1} and ϵ_t .

There are several interesting implications of this last set of findings. First, they suggest that, conditional on some prices changing in response to news, trades do not cause price changes, but that the lack of price changes does cause trades. To see this, note that if in these circumstances, trades cause price changes, then there would be buying of GM

²⁷For evidence that specialists are able to anticipate order flows, see Sirri (1990).

²⁸Quoted prices would be sticky if the specialist could adjust them, but does not; or, if investors place on the specialist’s book limit orders which are not immediately revised when information is released.

when its price rises. This, however, should lead to a negative own-stock correlation of m_{t-1} and ϵ_t , which is not found. To generate negative cross-stock effects and zero own-stock effects, it must be that when some prices rise and others do not, outside investors predominantly buy the laggards.

Second, this cross-stock effect is once again closely related to the processing of market-wide information. It is shown above that the covariance between m_{t-1} and ϵ_t is falling because of more aggressive trading of stocks whose quotes are slow to respond to market-wide news. Such trading is clearly helpful in eliminating positive correlation in last-trade quotes. Of course, if the processing of market-wide information were *completely* efficient, all stock prices would respond instantly, and this would tend to choke off such cross-stock trading in the first place. But, in such a world one would observe zero autocorrelation in an index of current midquotes, which, as shown in the following subsection, is not yet the case. In sum, the negative cross-stock covariance of m_{t-1} and ϵ_t suggests that trading pressures are working toward enhancing the efficiency of the market index. Therefore, it may be desirable to include line 6 of Table VI in the portion of the decline in autocorrelation due to improved market efficiency.

This subsection has shown that, after purging the last-trade index of bid-ask effects, the decline in the first-order autocorrelation from 1983 to 1988 is about -0.221 , or about two thirds of the -0.332 decline in the autocorrelation of the last-trade index. Of the remaining -0.111 , -0.045 might be attributed to the slow-response-to-information hypotheses (-0.009 to SRFI and -0.036 to SRPI), which is felt to reflect better processing of market-wide information. Thus, only -0.051 of the -0.332 decline in the autocorrelation of l_t can be attributed to measurement error induced by classic bid-ask bounce.

Nontrading Effects

The more frequently mentioned—and potentially more serious—form of measurement error comes from nontrading. Because the m_t index is computed from last-trade quotes, some fraction of individual stock quotes will always be stale. As trades occur in these stocks, any apparent staleness will disappear, creating the impression that information seeps slowly into the index. Thus, the last-trade index appears *positively* correlated, even if the prices at which these stocks would trade (if they *were* to trade) might respond instantaneously to information.

The size of this nontrading correlation depends on two factors: the frequency with which stocks trade, and the degree to which trades are synchronized across stocks. Clearly, greater trading volume works to reduce nontrading and, hence, to reduce the serial correlation in returns. Alternatively, greater synchronization in trades across stocks can affect index correlation, even holding fixed the volume of trade. Common models of nontrading, such as that of Scholes and Williams (1977), are not easily able to capture the importance of the latter effect. Rather than try to test a particular model of nontrading, the transactions data are used in an attempt to purge the index of the effects of nontrading.²⁹

Transactions Data

The 1983 and 1988 transaction data discussed above are used to calculate explicitly a measure of staleness. Following Harris, Sofianos, and Shapiro (1990), the last-trade–midquote index, m_t , can be thought of as equal to the *current* midquote index plus a staleness term—the difference between the last-trade midquote and the current midquote:

$$m_t = cm_t + s_t = cm_t + (m_t - cm_t) \quad (11)$$

To understand eq. (11), think of the true underlying index as equaling the average of *current* bid and ask prices. The return on this current midquote index, which is free of staleness and bid–ask bounce, is given by cm_t . Then the error in measuring returns using information available at the times when stocks last traded (as opposed to using current information) is given by $s_t = m_t - cm_t$. By examining s_t 's role in the decline in the autocorrelation of m_t , more direct evidence on the importance of nontrading is gained. Note that the autocorrelation of cm_t has real economic implications. Positive autocorrelation in cm_t would, for example, say that it is better not to sell after an up-tic in the market, but to wait until after a down-tic.

Table VII begins the decomposition by comparing the first-order autocovariances of m_t and cm_t in the first two lines. For purposes of comparability with the previous table, these are scaled by the variance of l_t .³⁰ In the first line of Table VII is the autocovariance of m_t from Table VI.

²⁹ Atchison, Butler, and Simonds (1987) use actual transaction arrival rates to estimate the Scholes and Williams (1977) model for daily returns on the NYSE. They find that the model can explain only 10–15% of the observed correlation in this index.

³⁰ The variances of these variables are broadly similar. For example, in 1983 the average daily return variances ($\times 10^6$) were: $l_t = 0.8310$, $m_t = 0.6622$, and $cm_t = 0.7535$.

TABLE VII
Decomposition of Last-Trade-Midquote Index Returns

Variable	1983	1988	Change
Averages of intraday ratios, excluding first 30 minutes			
1. $\text{cov}(m_t, m_{t-1})/\text{var}(l_t)$	0.389	0.168	-0.221
$\sum_{j=1}^N \omega^j \omega^j \text{cov}(m_t^j, m_{t-1}^j)/\text{var}(l_t)$	-0.006	-0.003	0.003
$\sum_{j=1}^N \sum_{j \neq i} \omega^i \omega^j \text{cov}(m_t^i, m_{t-1}^j)/\text{var}(l_t)$	0.395	0.171	-0.224
2. $\text{cov}(cm_t, cm_{t-1})/\text{var}(l_t)$	0.412	0.142	-0.270
$\sum_{j=1}^N \omega^j \omega^j \text{cov}(cm_t^j, cm_{t-1}^j)/\text{var}(l_t)$	-0.007	-0.003	0.004
$\sum_{j=1}^N \sum_{j \neq i} \omega^i \omega^j \text{cov}(cm_t^i, cm_{t-1}^j)/\text{var}(l_t)$	0.419	0.145	-0.274
3. $(\text{cov}(m_t, m_{t-1}) - \text{cov}(cm_t, cm_{t-1}))/\text{var}(l_t)$	-0.023	0.026	0.049
$\sum_{j=1}^N \omega^j \omega^j (\text{cov}(m_t^j, m_{t-1}^j) - \text{cov}(cm_t^j, cm_{t-1}^j))/\text{var}(l_t)$	0.001	-0.000	-0.001
$\sum_{j=1}^N \sum_{j \neq i} \omega^i \omega^j (\text{cov}(m_t^i, m_{t-1}^j) - \text{cov}(cm_t^i, cm_{t-1}^j))/\text{var}(l_t)$	-0.024	0.026	0.050
4. $\text{cov}(s_t, s_{t-1})/\text{var}(l_t)$	-0.014	-0.020	-0.006
$\sum_{j=1}^N \omega^j \omega^j \text{cov}(s_t^j, s_{t-1}^j)/\text{var}(l_t)$	-0.017	-0.005	0.012
$\sum_{j=1}^N \sum_{j \neq i} \omega^i \omega^j \text{cov}(s_t^i, s_{t-1}^j)/\text{var}(l_t)$	0.003	-0.015	-0.018
5. $\text{cov}(s_t, cm_{t-1})/\text{var}(l_t)$	0.087	0.094	0.007
$\sum_{j=1}^N \omega^j \omega^j \text{cov}(s_t^j, cm_{t-1}^j)/\text{var}(l_t)$	0.016	0.005	-0.011
$\sum_{j=1}^N \sum_{j \neq i} \omega^i \omega^j \text{cov}(s_t^i, cm_{t-1}^j)/\text{var}(l_t)$	0.071	0.089	0.018
6. $\text{cov}(cm_t, s_{t-1})/\text{var}(l_t)$	-0.097	-0.048	0.048
$\sum_{j=1}^N \omega^j \omega^j \text{cov}(cm_t^j, s_{t-1}^j)/\text{var}(l_t)$	0.002	0.000	-0.002
$\sum_{j=1}^N \sum_{j \neq i} \omega^i \omega^j \text{cov}(cm_t^i, s_{t-1}^j)/\text{var}(l_t)$	-0.099	-0.048	0.050

Notes: Indexes are constructed to approximate the S&P 500, using NYSE stocks only. See footnote 21 in the text for more details. l_t represents the last-trade return index, m_t the last-trade-midquote return index, cm_t the current midquote return index, and $s_t = m_t - cm_t$ is measurement error in the last-trade-midquote index due to nontrading staleness.

The second line reports estimates of the autocovariance of the current midquote index, cm_t . Most striking is that *its decline of -0.270 is greater than that for m_t* . In other words, nontrading staleness does not explain a positive portion of the reduction in the autocorrelation of l_t —it actually makes the decline in index autocorrelation even more striking. How could it be that, all else equal, as staleness due to nontrading is reduced, the autocorrelation of l_t actually rises?

The answer lies in the elimination of strong cross-stock covariation in quoted prices. For example, suppose that the current quotes are set somewhat inefficiently, in that price changes for GM, while being serially uncorrelated, always lead by one day those of Ford. In this case, an index of *current* quotes will show positive autocorrelation. However, now add the assumption that GM trades continuously,

but that Ford happens to have traded only very early in the trading day. In that case, Ford's last-trade quotes lag behind current Ford quotes by almost one day (and behind current GM quotes by almost two days). Because of this asymmetry, the resulting last-trade index is *less* positively autocorrelated than the current midquote index. Therefore, when trading volume picks up, the autocorrelation in the last-trade index *rises*. This cross-stock asymmetry can explain the results in the first two lines of Table VII for 1983, and the fact that the decline in the autocovariance of cm_t is *larger* than that of m_t .

To further explore this notion of asymmetric predictive power across stocks, another version of the value-weighted cm_t index is computed—this time using equal weights. Defining the equally-weighted current midquote index as eq_t and $z_t \equiv cm_t - eq_t$, the autocovariance of cm_t can be decomposed into four terms:

$$\begin{aligned} \frac{\text{cov}(cm_t, cm_{t-1})}{\text{var}(l_t)} &= \frac{\text{cov}(eq_t, eq_{t-1})}{\text{var}(l_t)} + \frac{\text{cov}(eq_t, z_{t-1})}{\text{var}(l_t)} + \frac{\text{cov}(z_t, eq_{t-1})}{\text{var}(l_t)} \\ &+ \frac{\text{cov}(z_t, z_{t-1})}{\text{var}(l_t)} \end{aligned} \quad (12)$$

Loosely speaking, the terms on the right-hand side of eq. (12) can be interpreted as follows: the first is a measure of small stocks' ability to predict the return on other small stocks; the second is a measure of large stocks' ability to predict returns on small stocks; the third is a measure of small stocks' ability to predict returns on large stocks; and the fourth is a measure of large stocks' ability to predict returns on other large stocks. If stocks respond symmetrically to market-wide information, one would expect that the -0.27 decline in the autocorrelation of cm_t would be distributed equally across these four components. In fact, the change of -0.27 is made up of declines of -0.12 , -0.11 , -0.01 , and -0.04 , respectively, of these four terms. It follows that the overall decline in autocorrelation has come mostly (and about equally) from a fall in the ability of small stocks to predict returns on other small stocks and a fall in the ability of large stocks to predict returns on small stocks.³¹

The fourth, fifth, and sixth lines of Table VII decompose the difference between the autocovariances of m_t and cm_t into three

³¹Lo and MacKinlay (1990) show that the predictability of small stock returns accounts for a large portion of index autocorrelation.

components, similar to those in eq. (10):

$$\frac{\text{cov}(m_t, m_{t-1})}{\text{var}(l_t)} - \frac{\text{cov}(cm_t, cm_{t-1})}{\text{var}(l_t)} = \frac{\text{cov}(s_t, s_{t-1})}{\text{var}(l_t)} + \frac{\text{cov}(s_t, cm_{t-1})}{\text{var}(l_t)} + \frac{\text{cov}(cm_t, s_{t-1})}{\text{var}(l_t)} \quad (13)$$

The changes in lines 4 and 5 are negligible, so that the only important source of net change is measured by the last term on the right-hand side of eq. (13), which is reported on line 6. This term measures the covariance of the new information in quotes beyond that reflected in the last trade, $s_{t-1} = m_{t-1} - cm_{t-1}$, and the return on the current midquote index, cm_t . One might expect this covariance to be negative and rising over time because the autocovariance of cm_t is positive and declining over time. Line 6 of Table VII shows that the covariance between s_{t-1} and cm_t indeed increases from 1983 to 1988 by 0.048, with own- and cross-components of -0.002 and 0.050 , respectively. This cross-covariance can be interpreted as a measure of the responsiveness of current quotes to information which comes out between time $t - 1$ and the last trade as of time $t - 1$. It is in this sense that a decline in the cross-covariance of s_{t-1} and cm_t is evidence of more rapid dissemination of market-wide information.

Note that this effect suggests a more rapid response of quotes to other stocks' quote revisions, not necessarily triggered by trading. This complements both the SRFI hypothesis above (which suggests more rapid response of quotes to other stocks' *order flows*) and the SRPI hypothesis (which suggests more rapid response of *order flows* to changes in other stocks' quotes).

In sum, when a current midquote index, cm_t , which has been purged of the effects of both the bid-ask spread and staleness is computed, it accounts for about -0.270 of the -0.332 decline from 1983 to 1988 in the first-order autocorrelation of the last-trade index, l_t . If one adds the -0.045 decline due to slow response to information (the SRFI and SRPI hypotheses given above), the result is -0.315 of the -0.332 change in the autocorrelation of l_t . Tables VI and VII also show that this decline is entirely due to cross-stock effects. This means that the results are best explained by more rapid processing of market-wide information.

DAILY AND WEEKLY AUTOCORRELATIONS

The focus to this point has been on returns over holding periods of 15 minutes for as long as 15-minute data are available. This section

looks at the first-order autocorrelation of daily and weekly returns to determine whether the decline in autocorrelation also applies to longer horizon returns and whether it is part of a longer-term trend.

Figure 3 shows the first-order autocorrelation of daily returns in each year since 1926. Three different indices are used: Dow-Jones Industrials, the S&P 500, and the NYSE value-weighted index (the last being available through CRSP only since 1962).³² Several striking observations come out of Figure 3.

First, it is clear that the decline in autocorrelation documented in the foregoing sections is evident in daily returns, and that the 1980s are part of a longer-term secular decline in serial correlation which began around 1969. At that time, the daily autocorrelation coefficients were between 0.3 and 0.4—very high when compared to the more




FIGURE 3
Autocorrelation of daily returns on stock indices.

³²The unusual observation for 1963 seems related to the assassination of President John F. Kennedy on Friday, November 22, 1963. On that day the market fell by almost 3%, then rebounded upward on the next trading day (Tuesday) by 3%. When those days are removed from the data the autocorrelation coefficients jump up to about 0.1.

recent years, when daily autocorrelations have, on average, been slightly negative.³³

Second, the three indices tell essentially the same story. Their parallel behavior is important because it indicates that nontrading is unlikely to explain much of the variation in autocorrelations. To see this, note that the Dow-Jones Industrials includes only 30 stocks, all of which are traded very frequently, in contrast with the broader S&P 500 and NYSE indices. Because the Dow-Jones is more actively traded, its serial correlation is expected to be lower, which is indeed the case during the postwar period. Note, however, that the differential between the Dow-Jones and other indices has not changed much during the recent period. It does not appear to have increased—indeed, it has *decreased*—between 1969 and 1990, a period during which index autocorrelations fell by almost 0.4.

Third, there seem to be three distinct regimes since 1926. The first, which corresponds roughly to the interwar period, shows autocorrelations to be about zero. The second, beginning with the war and lasting until the late 1960s, seems (with the exception of 1963) to be constant at about 0.15. The most recent period is associated with a large, but remarkably steady, decline from 0.4 to zero.

Finally, note that the variation in autocorrelations is not predominantly due to changes in the average autocorrelation of individual stock returns. To demonstrate this, Figure 4 graphs the first-order autocorrelation on the NYSE value-weighted index along with the *average* own-stock autocorrelation—the first term on the right-hand side of eq. (4).³⁴ The figure clearly shows that the decline in autocorrelation that began in 1969 is due to cross-stock returns.

The last piece of evidence comes from Figures 5 and 6. They show autocorrelations of weekly returns for the S&P 500, Dow Jones, and value- and equally-weighted NYSE indices, respectively. Because the standard error of each year's autocorrelation coefficient is large, the figures include seven-year moving averages of the coefficients. The hump in autocorrelation beginning in the early 1960s remains evident in these graphs. It is also clear that the positive autocorrelation often found in weekly returns comes primarily from this hump, and that the

³³Monte Carlo simulations suggest that standard errors for these coefficients (allowing for heteroskedasticity) are about 0.12.

³⁴The average own-autocorrelation is estimated by taking the simple average autocorrelation of returns on the 150 largest capitalization stocks for each year and dividing by 150. Since these stocks represent only a fraction of the NYSE's capitalization, this estimate is likely to overstate the magnitude of the own-stock contribution to autocorrelation.




FIGURE 4
First-order autocorrelation of daily NYSE returns.

autocorrelation has not been strongly positive since the mid-1970s. The exception to this is the equally weighted NYSE index return in Figure 6. Its autocorrelation has fallen the least remaining relatively high. This suggests that high-frequency portfolio trading does not yet include a large number of small stocks and therefore cannot fully discipline their prices.

What could explain the episodic behavior of serial correlation seen in Figures 3, 5, and 6? Could the trading practices of the day explain why autocorrelations were so high in the late 1960s and early 1970s, and so low in the 1930s? One possibility is that the relative importance of institutional versus investors has changed over time, and that these investors exhibit very different trading behavior. This is clearly a question for future research.

CONCLUSIONS

The main empirical finding is that the predictability of short-term stock returns has declined markedly in 15-minute data, and somewhat less




FIGURE 5
First-order autocorrelations of weekly returns.

markedly in daily and weekly data. These changes seem concurrent with rapid growth in new institutional trading practices like portfolio and index futures trading. The possibility that technical explanations, such as increases in bid-ask bounce and decreases in nontrading, are responsible for the decline in autocorrelation of 15-minute returns is examined, but the data do not support such explanations. In addition, little evidence is found to support the overreaction hypothesis, which would suggest decreases in both own- autocorrelations and higher-order index autocorrelations, neither of which is found. It is argued, therefore, that the reduction in autocorrelation, which are overwhelmingly due to cross-stock effects, are a result of improved efficiency with which market-wide information is impounded into the prices of individual stocks.

Of course, these results do not imply that new trading practices are beneficial, nor that prices are now closer to the present value of dividends. Futures markets could still produce negative externalities if, in the process of making the index more efficient, futures siphon off




FIGURE 6
First-order autocorrelations of weekly returns.

order flow from individual stocks, and thereby lead to greater inefficiency with respect to stock-specific information.³⁵

BIBLIOGRAPHY

- Atchison, M. D., Butler, K. C., and Simonds, R. R. (1987, March): "Nonsynchronous Security Trading and Market Index Autocorrelation," *Journal of Finance*, 42:111-118.
- Cochran, J. (1988, October): "How Big Is the Random Walk in GNP?" *Journal of Political Economy*, 96:893-920.
- Cutler, D., Poterba, J., and Summers, L. (1990, May): "Speculative Dynamics and the Role of Feedback Trades," *American Economic Review*, 80: 63-69.
- French, K., and Roll, R. (1986): "Stock Return Variances: The Arrival of Information and the Reaction of Traders," *Journal of Financial Economics*, 17:5-26.
- Gammill, J., and Perold, A. (1989, Spring): "The Changing Character of Stock Market Liquidity," *Journal of Portfolio Management*, 15:13-18.
- Glosten, L. R., and Milgrom, P. R. (1985): "Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14:71-100.

³⁵See Gammill and Perold (1989) and Subrahmanyam (1989) for an elaboration of this argument.

- Harris, L. (1989, December): "S&P 500 Cash Stock Price Volatilities," *Journal of Finance*, 44:1155–1176.
- Harris, L., Sofianos, G., and Shapiro, J. E. (1990, February): "Program Trading and Intraday Volatility," NYSE Working Paper.
- Kyle, A. S. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53:1315–1336.
- Lo, A. W., and MacKinlay, A. C. (1988): "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test," *Review of Financial Studies*, 1:41–66.
- Lo, A. W., and MacKinlay, A. C. (1990): "When Are Contrarian Profits Due to Stock Market Overreaction?" *Review of Financial Studies*, 3:175–205.
- MacKinlay, A. C., and Ramaswamy, K. (1988, Summer): "Index-Futures Arbitrage and the Behavior of Stock Index Futures Prices," *Review of Financial Studies*, 1:137–158.
- Poterba, J., and Summers, L. (1988): "Mean Reversion in Stock Prices: Evidence and Implications," *Journal of Financial Economics*, 22:27–59.
- Roll, R. (1984): "A Simple Implicit Measure of the Effective Bid–Ask Spread in an Efficient Market," *Journal of Finance*, 39:1127–1139.
- Scholes, M., and Williams, J. (1977, December): "Estimating Betas from Nonsynchronous Data," *Journal of Financial Economics*, 5:309–327.
- Schwert, W. (1989, December): "Why Does Stock Market Volatility Change over Time?" *Journal of Finance*, 44:1115–1154.
- Sirri, E. (1990): "Price, Volume, and Bid–Ask Spread in a Specialist Market," Mimeo, Harvard Business School.
- Subrahmanyam, A. (1989, April): "A New Rationale for Markets in Baskets of Stocks," Working Paper 10-89, Anderson School of Management, UCLA.
- White, H. (1980, May): "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," *Econometrica*, 48:817–838.