Journal of Financial Economics 17 (1986) 5–26. North-Holland

STOCK RETURN VARIANCES
The Arrival of Information and the Reaction of Traders*

Kenneth R. FRENCH

University of Chicago, Chicago, IL 60637, USA

Richard ROLL

University of California, Los Angeles, CA 90024, USA
Goldman, Sachs & Co., New York, NY 10004, USA

Received March 1985, final version received January 1986

Asset prices are much more volatile during exchange trading hours than during non-trading hours. This paper considers three explanations for this phenomenon: (1) volatility is caused by public information which is more likely to arrive during normal business hours; (2) volatility is caused by private information which affects prices when informed investors trade; and (3) volatility is caused by pricing errors that occur during trading. Although a significant fraction of the daily variance is caused by mispricing, the behavior of returns around exchange holidays suggests that private information is the principle factor behind high trading-time variances.

1. Introduction

Equity returns are more volatile during exchange trading hours than during non-trading hours. For example, the variance of returns from the open to the close of trading on an average day is over six times larger than the variance of close-to-open returns over a weekend, even though the weekend is eleven times longer. On an hourly basis, the variance when the exchanges are open is between 13 and 100 times larger, depending on the non-trading period being considered.

The phenomenon has been pointed out by several authors including Fama (1965), Granger and Morgenstern (1970), Oldfield and Rogalski (1980), and Christie (1981), but it has not generated much attention. We believe it is important. It represents an empirical puzzle whose solution may provide a deeper understanding of information processing in financial markets.

*This paper has benefited from the comments of seminar participants at Boston College, the University of British Columbia, the University of Chicago, Dartmouth College, Harvard University, Northwestern University, Purdue University and Stanford University. We are also grateful to Craig Ansley, Merton Miller, Steven Ross, Robert Stambaugh, William Schwert, Jerold Warner (the referee), and especially, Douglas Diamond and Eugene Fama for comments on an earlier draft.

We consider three possible explanations for the observed variance pattern. The first possibility is that more public information arrives during normal business hours. Under this hypothesis, most return volatility is caused by things like judicial decisions and tender offers and these announcements are clustered during the trading day. The second explanation assumes that most return volatility is caused by private information and that this information only affects prices through the trading of informed investors. If the informed investors are more likely to trade when the exchanges are open, return variances will be high during this period.

The third possibility we consider is that the process of trading introduces noise into stock returns. For example, perhaps investors over-react to each other's trades. This trading noise would increase return variances when the exchanges are open.

To determine the relative importance of these three explanations, we examine the behavior of returns around business days when the New York and American Stock Exchanges were closed. If high trading-time variances are caused by the arrival of public information during the business day, return variances should not fall simply because the exchanges are closed. On the other hand, both the trading noise hypothesis and the private information hypothesis predict that return variances will be unusually low around exchange holidays. We find that the two-day return variance around exchange holidays is only slightly larger than the variance of a normal one-day return.

Our exchange holiday results are consistent with both the private information hypothesis and the trading noise hypothesis. To discriminate between these hypotheses we compare daily return variances with variances for longer holding periods. If daily returns are independent, the variance for a long holding period will equal the cumulated daily variances within the period. However, if daily returns are affected by trading noise, the longer holding period variance will be smaller than the cumulated daily variance.

These tests suggest that, on average, between 4% and 12% of the daily return variance is caused by mispricing. However, even if we assume that all of the mispricing occurs during the trading day, it has a small impact on the relation between trading and non-trading variances. It appears that the large difference between these variances is caused by differences in the arrival and incorporation of information during trading and non-trading periods.

2. Trading and non-trading variances

If hourly stock return variances were constant across trading and non-trading periods, the variance of weekend returns (i.e., Friday close to Monday close) would be three times the variance of weekday returns (e.g., Tuesday close to Wednesday close). In this section we examine this proposition and we

Table 1

Average ratios of multiple-day variances relative to single-day variances for all NYSE and AMEX stocks and for quintiles of stocks sorted by equity value.[a]

		All stocks	Smallest quintile[a]	2	3	4	Largest quintile
Two-day holidays	Average ratio[b]	1.247	1.301	1.199	1.239	1.217	1.281
	Standard error[c]	0.066	0.068	0.054	0.052	0.097	0.100
	Standard deviation[d]	1.354	1.446	1.270	1.371	1.149	1.351
	Number of firms[e]	1962.5	390.3	392.3	392.4	393.5	394.1
	Average sample size[f]	10.0	10.0	10.0	10.0	10.0	10.0
Weekends	Average ratio	1.107	1.122	1.108	1.119	1.105	1.082
	Standard error	0.012	0.010	0.016	0.014	0.014	0.017
	Standard deviation	0.385	0.412	0.379	0.435	0.337	0.286
	Number of firms	2055.3	411.2	410.8	410.6	411.2	411.5
	Average sample size	92.8	92.5	92.8	92.9	93.0	93.0
Holiday weekends	Average ratio	1.117	1.111	1.122	1.099	1.122	1.130
	Standard error	0.092	0.053	0.085	0.071	0.106	0.151
	Standard deviation	1.219	1.176	0.992	1.276	1.232	1.014
	Number of firms	2055.7	411.3	410.8	410.9	411.2	411.5
	Average sample size	11.1	11.0	11.1	11.1	11.1	11.1

[a] Stocks are sorted into quintiles based on their equity values at the beginning of ten two-year subperiods between 1963–1982.

[b] The average ratio comparing the variance of two-, three-, and four-calendar-day returns with the variance of one-day returns. This estimate is the average of the ten subperiod averages.

[c] The standard error of the reported average ratio. This standard error is based on the distribution of the ten subperiod average ratios.

[d] The average cross-sectional standard deviation. The ratios for individual firms are used to estimate the standard deviation for each subperiod. The reported standard deviation is the average of the ten subperiod standard deviations.

[e] The average number of firms in each subperiod.

[f] The average number of multiple-day returns for each stock in each subperiod.

report on the relation between firm size and the trading/non-trading variance differential.

Our tests use the daily returns provided by the Center for Research in Security Prices for all common stocks listed on the New York and American Stock Exchanges between 1963 and 1982. We break this twenty-year period into ten two-year subperiods. For each stock, we calculate return variances for weekdays, weekends, holidays, and holiday weekends during each subperiod. These estimates are used to compute multiple-to-single-day variance ratios for each stock in each subperiod.

The first column of table 1 reports grand averages of the estimated variance ratios. The grand averages are calculated by first averaging the variance ratios across the stocks within each subperiod and then averaging the ten subperiod averages. The grand averages are consistent with the evidence in earlier papers.

The variance of the total return over a weekend or a holiday is only slightly higher than the variance of the total return over a normal weekday. For example, the variance for a three-day weekend return is only 10.7% higher than the variance for a normal one-day return.

Table 1 also reports standard errors of the grand averages. These standard errors, which are based on the distribution of the ten subperiod averages, range from 0.04 for weekends to 0.29 for holiday weekends. Under the assumption that the subperiod averages are independent and identically distributed, the grand averages are many standard errors below 2.0, 3.0, or 4.0.

One can imagine many factors that might affect the way investors acquire and react to information about particular firms. For example, perhaps firms in some industries are closely monitored by financial analysts, while little private information is collected about firms in other industries. In this study, we concentrate on firm size as a potential factor because it is easy to observe and because the rewards from acquiring and using firm-specific information are probably a function of this variable.

To examine whether the relation between trading and non-trading variances is a function of firm size, we sort firms into quintiles based on their equity values at the beginning of each subperiod. The averages of the subperiod averages for the quintiles are reported in columns 2 through 6 in table 1. There is no obvious relation between the estimated variance ratios and firm size. For example, the average two-day variance ratio for the smallest firms (column 2) is 1.30 with a standard error of 0.07 and the average ratio for the largest firms (column 6) is 1.28 with a standard error of 0.10.

To see what the estimated variance ratios imply about the difference between trading and non-trading variances, assume that

(a) returns are intertemporally uncorrelated,
(b) the exchange is open six hours per day (the present situation),
(c) there are just two uniform regimes, trading and non-trading hours; returns are identically distributed within these regimes but have different variances between them.

(These assumptions are made at this point merely for temporary illustrative convenience. We relax them later.)

Let σ_T^2 be the variance of returns *per hour* during trading and let σ_N^2 be the variance per hour at other times. Since there are 66 non-trading hours over the weekend and 18 non-trading hours in a normal business day, the average weekend-to-weekday variance ratio for all firms implies

$$66\sigma_N^2 + 6\sigma_T^2 = 1.107\left(18\sigma_N^2 + 6\sigma_T^2\right). \tag{1}$$

Thus, based on the weekend variance ratio,

$$\sigma_T^2 = 71.8\sigma_N^2;$$

the hourly variance when the New York exchanges are open is roughly seventy times the hourly variance when they are closed. We can make similar transformations with the average variance ratios for two- and four-day holidays. Using the averages for all stocks in table 1 gives:

Non-trading interval	Hourly trading to non-trading variance ratio
Mid-week holidays	13.2
Weekends	71.8
Holiday Weekends	99.6

Trading hours are more volatile than non-trading hours. Among non-trading hours, weekends have lower volatility than normal holidays and holiday weekends have the lowest volatility of all.

3. Possible explanations

There seem to be two general explanations for the empirical phenomenon that prices are more variable during exchange trading hours. The obvious possibility is that information arrives more frequently during the business day. The second possible explanation is that trading somehow induces volatility.

To examine the first possibility, it is useful to divide information into two categories: public information and private information. Public information is information that becomes known at the same time that it affects stock prices. Examples of this information include changes in the weather, Supreme Court decisions, and the outcome of the World Series. Information produced by firms, such as financial reports, or by the government, such as United States Department of Agriculture crop forecasts, is included in this category if no one trades on the information before it is released.

Private information is at the other end of the spectrum. While public information affects prices before anyone can trade on it, private information only affects prices through trading. Much of the information produced by investors and security analysts is in this category.

Obviously, most information falls in the continuum between public and private information. However, this artificial dichotomy is useful because it allows us to develop and test several hypotheses about the variance pattern we observe.

Our first hypothesis is that the higher trading-time volatility occurs because public information is more likely to appear during normal business hours. This explanation is plausible since most public information is probably a by-product of normal business activities.

The private information hypothesis is similar. Under this hypothesis, return variances are higher during trading hours because most private information is incorporated into prices during this period. There are two possible reasons for this. First, the production of private information may be more common when the exchanges are open. For example, security analysts are more likely to work at this time. Activities such as visiting corporate headquarters, examining company documents, and making recommendations to clients are all easier to do during the business day. In addition, the benefits of producing private information are larger when the exchanges are open and the information can be acted on quickly and conveniently.

Even if private information is produced at a constant rate during both trading and non-trading periods, trades based on this information could lead to high trading-time variances. Consider the effect of private information that is produced after the New York exchanges close. Since this information can only affect prices through the trading of informed investors, the price reaction is delayed until this trading occurs. If the informed investors trade on the New York exchanges, their information cannot affect prices until the exchanges open.

The fact that private information only affects prices when markets are open appears to offer a simple, yet general, explanation for high trading-time variances. However, this story will not explain the results in table 1 unless we assume that private information affects returns for more than one trading day. All of the estimates in table 1 are based on close-to-close returns, which include both a non-trading period and a trading period. If non-trading information is completely revealed in prices during the next trading day, it will affect the 'right' close-to-close return. For example, if private information produced during the weekend only affects Monday's return and information produced during a weeknight only affects the next day's return, the weekend-to-weekday variance ratios in table 1 accurately reflect the private information produced during each period. Unless private information affects prices for more than one trading day, the hypothesis that informed investors only trade when the exchanges are open cannot explain the low variance ratios in table 1.

To summarize, the private information hypothesis says that the variance pattern we observe occurs either because most private information is produced during normal business hours or because informed investors usually trade when the exchanges are open and they trade on their information for more than one day.

The second general explanation for high trading-time variances is that the process of trading introduces noise into stock returns. Suppose each day's return can be broken into two components: an information component that reflects a rational assessment of the information arriving that day, and an independent or positively correlated error component.[1] If the daily pricing error occurs during the trading period, it will increase the trading-time variance. It is important to note that under this hypothesis at least some trading noise (the error component in the daily return) is not corrected during the trading day in which it occurs. If all trading noise were corrected quickly, the noise would increase intra-day return variances, but it would not affect our close-to-close returns.

In summary, the hypotheses to be examined are:

(H.1) High trading-time volatility is caused by public information which is more likely to be observed during normal business hours.

(H.2) High trading-time volatility is caused by private information which is more likely to affect prices when the exchanges are open.

(H.3) High trading-time volatility is caused by pricing errors that occur during trading.

4. Tests of the hypotheses

In this section we examine the predictions of the three hypotheses. It is important to recognize that the hypotheses are not mutually exclusive. In fact, the observed variance pattern might be caused by all three factors simultaneously. Our goal is to provide some sense of the empirical importance of each explanation.

4.1. Exchange holidays

The New York and American Stock Exchanges were closed on Wednesdays during the second half of 1968 because of a paperwork backlog. The exchanges were also closed on many of the election days in our sample period. These exchange holidays give us an excellent opportunity to examine the relative importance of our three hypotheses.[2]

Under the public information hypothesis, the return variance for a business day should not depend on whether the exchanges are open or closed. Therefore, this hypothesis predicts that stock return variances will not be reduced by the exchange holidays in 1968. The prediction of the public information hypothesis for election holidays is less clearcut. One might expect unusually high variances on election days since election results are publicly observable

[1] In the discussion below, we add a third component that arises because of the bid/ask spread.

[2] French (1980) also uses returns around the 1968 exchange holidays to make inferences about the exchange holidays.

information. However, perhaps the exchanges close on election days because less public information is available.

The private information hypothesis predicts that return variances will be reduced by both the election day closings and the exchange holidays in 1968. The size of this reduction depends on the interval used to compute returns. Since private information only affects prices when informed investors trade, the reduction in the variance should be large during the period that the exchanges are closed. For example, the variance of the return from the close of trading on Tuesday to the open on Thursday should be unusually low if the exchanges are closed on Wednesday.

Much of the reduction in the variance will be eliminated if the next day's trading is included in the return. The information that would have affected prices on Wednesday will affect prices during trading on Thursday instead. However, the variance for the two-day close-to-close return from Tuesday to Thursday should still be less than twice the variance of a normal one-day return. This difference may persist for two reasons. First, private information may affect prices for more than one trading day. The information that would have been revealed through trading on Wednesday and Thursday may not be fully incorporated in prices if trading is limited to Thursday. Second, less private information may be produced when the markets are closed. Exchange holidays reduce the value of private information. Informed investors either must delay acting on their information – and run the risk that someone else will discover it – or they must find a less convenient way to trade. Because of its reduced value, less private information will be produced when the exchanges are closed.

If we increase the holding period to one week, the private information hypothesis predicts that the effect of the exchange holiday on the total variance should be even smaller. Equivalently, the variance for the days following an exchange holiday should be larger than normal. Adding more trading days to the return interval allows more time for the private information to affect prices. Also, with less information produced on the exchange holiday, more will be produced on succeeding days. There are two reasons for this production effect. First, with more information available to produce, the cost of generating any particular amount should fall. Second, some of the information that is not produced privately because the exchanges are closed might become publicly observable after a few days.

Hypothesis H.3 makes a simpler prediction. If high trading-time variances are caused by trading noise, the variance should fall when the exchanges are closed and the variance that is lost should not be recovered.

Table 2 presents evidence to test these predictions. The first section of this table reports daily variance ratios comparing the two-day returns for exchange holidays in 1968 (from Tuesday close to Thursday close) with a normal one-day variance estimated between January 1963 and December 1982. The

Table 2

Daily and weekly variance ratios for exchange holidays.

		All stocks	Smallest quintile[a]	2	3	4	Largest quintile
		Daily variance ratios					
Exchange holidays in 1968	Average ratio[b]	1.145	1.077	1.043	1.180	1.239	1.274
	Standard deviation[c]	0.882	0.857	0.647	0.979	0.944	1.001
	Number of firms	2083	597	455	374	342	315
	Average sample size[d]	22.7	22.8	22.7	22.6	22.7	22.85
Election holidays	Average ratio[e]	1.165	1.131	1.073	1.186	1.159	1.332
	Standard deviation	1.079	1.222	1.065	1.040	0.799	1.118
	Number of firms	2026	572	426	367	347	314
	Average sample size	8.5	8.2	8.2	8.1	8.7	9.5
		Weekly variance ratios					
Exchange holidays in 1968	Average ratio[f]	0.821	0.901	0.802	0.772	0.793	0.784
	Standard deviation	0.559	0.667	0.484	0.422	0.511	0.612
	Number of firms	2093	600	457	376	344	316
	Average sample size	20.6	20.7	20.6	20.5	20.7	20.8
Election holidays	Average ratio	0.839	0.876	0.776	0.889	0.779	0.868
	Standard deviation	0.614	0.707	0.627	0.678	0.501	0.527
	Number of firms	1188	278	221	192	229	268
	Average sample size	8.5	8.3	8.3	8.3	8.6	8.8

[a] Firms are sorted into quintiles based on their relative equity values when they are first listed in the CRSP daily master file.
[b] Average variance ratio comparing two-day exchange holiday returns with single-calendar-day returns between January 1963 and December 1982.
[c] Cross-sectional standard deviation of the individual firm ratios.
[d] Average number of exchange holidays for each firm.
[e] Average variance ratio comparing two-day exchange holiday returns with single-calendar-day returns from 1962–1969, 1972, 1976, and 1980.
[f] Average ratio comparing the return variance for weeks containing exchange holidays with the return variance for weeks containing five trading days.

results are surprising. The average ratio across all stocks is 1.145. The averages for the size portfolios range from 1.043 for the second quintile to 1.274 for the fifth quintile. In other words, these point estimates indicate that, on average, the variance for the *two-day* exchange holiday returns is only 14.5% higher than the variance for normal one-day returns.

To get an idea about the reliability of these estimates, we construct similar ratios using the returns for Wednesday and Thursday during each half year from 1963 to 1982. For example, we compute a two-day variance for each stock using all of the Wednesday–Thursday returns observed during the first half of 1963. This variance is compared to the one-day variance estimated between July of 1963 and December of 1982. The ratio of these variances is averaged across stocks to get the average ratio for the first half of 1963. This process is repeated for each of the 39 half years in our sample. (The second half of 1968 is not included because it contains the Wednesday holidays.) The averages (which are not shown) range from 1.18 for the second half of 1964 to 4.32 for the second half of 1974, with a grand average of 2.00. It appears that the low 1968 variance ratio, 1.14, is not caused by chance, but by the exchange holiday.

The first section of table 2 also reports average daily variance ratios for election days. During our sample period, the exchanges closed for elections in 1962–1969, 1972, 1976, and 1980. Therefore, we compare the two-day election returns with one-day returns from those years. The average variance ratio for all stocks is 1.165. The portfolio averages range from 1.073 for the second quintile to 1.333 for the fifth quintile.

To check the reliability of the daily election ratios in table 2, we construct similar ratios using combined Tuesday–Wednesday returns for non-election weeks. Each replication involves one observation from each of the eight election years. For example, the first Tuesday–Wednesday pair of each election year is used in the first replication and the second pair is used in the second replication. This procedure generates a total of 45 replications, with average variance ratios ranging from 1.61 for the thirtieth Tuesday–Wednesday pair each year to 2.62 for the first pair. The grand average is 1.98. Again, it does not appear that the election holiday variance ratio of 1.17 is caused by chance. There appears to be a strong relation between the low variance ratios and the exchange holidays.

The daily variance ratios for election holidays and exchange holidays in 1968 are consistent with both the private information hypothesis and the trading noise hypothesis. However, these ratios provide little support for the public information hypothesis, which predicts that the two-day exchange holiday variance should be twice the one-day variance.

Weekly variance ratios in the second section of table 2 offer some evidence about the relative importance of private information and trading noise. Under the trading noise hypothesis, exchange holidays should cause a permanent reduction in the cumulated return variance. On the other hand, the private information hypothesis predicts that most of the lost variance will be recovered; when the holding period is increased there is more time to incorporate private information into prices and to discover information that was not produced on the exchange holiday.

To test these predictions, we compare the returns for weeks that include exchange holidays with the returns for normal five-trading-day weeks. For example, the weekly return for a Wednesday holiday in 1968 is measured from the close of trading on Tuesday to the close of trading on the following Tuesday. The five-trading-day variance is estimated using returns from Tuesday close to Tuesday close over all five-trading-day weeks in the full 1962–1982 sample period. The election week returns are measured from Monday close to Monday close and they are compared with weekly returns for 1962–1969, 1972, 1976, and 1980.

The weekly variance ratios in table 2 are consistent with the trading noise hypothesis. Across all stocks, the average weekly ratio for exchange holidays in 1968 is 0.82, and the average election week ratio is 0.84. However, neither of these estimates is very reliable. Simulated weekly variance ratios for the exchange holidays in 1968, which are constructed like the simulated daily ratios above, vary between 0.54 and 2.04. Simulated election week variances range from 0.76 to 1.53. The standard deviations of the simulated average ratios are 0.35 and 0.14, respectively. It is difficult to draw meaningful inferences from the weekly exchange holiday ratios.

4.2. *Autocorrelations*

The exchange holiday results support both the private information hypothesis and the trading noise hypothesis. We can obtain more information about the relative importance of these hypotheses by examining the autocorrelations of the daily returns. Neither public information nor private information will generate observable serial correlation. In principle, information may induce autocorrelation by changing the level of expected returns. However, the variance of expected returns is almost certainly so small that autocorrelation from this source is unobservable in realized returns for individual stocks.

Under the trading noise hypothesis, stock returns should be serially correlated. It is difficult to characterize short-run autocorrelations without a specific mispricing model. However, unless market prices are unrelated to the objective economic value of the stock, pricing errors must be corrected in the long run. These corrections would generate negative autocorrelations.

Two other factors may induce serial correlation under all three hypotheses. Close-to-close returns, such as those reported by CRSP, contain measurement error because each closing trade may be executed at any price within the bid/ask spread. If these measurement errors are independent from day to day, they will induce negative first-order autocorrelation. For example, suppose today's closing price is on the bid side of the market. Then today's observed return is negatively biased and tomorrow's observed return is positively biased. If today's price is on the ask side, the pattern is reversed but the observed returns are still negatively correlated.[3]

[3] If daily bid/ask errors are not independent, they can induce negative autocorrelations beyond lag 1. The autocorrelations in table 3 use all of the prices in the CRSP daily master file. These prices include both trade prices and the mean of bid and ask prices when a stock did not trade during a day. To control for one potential source of dependence, we have also estimated the autocorrelations using just trade prices. Deleting returns involving bid/ask prices has only one noticeable effect – the first-order autocorrelations increase slightly. For example, the average first-order autocorrelation across all stocks increases from 0.003 to 0.009.

Systematic variation in expected returns can also induce serial correlation. For example, the day of the week effects documented by French (1980) induce positive autocorrelations at every fifth lag (5, 10, 15, etc.) and negative autocorrelations at all other lags. Day of the month effects documented by Ariel (1984) also imply non-zero autocorrelations. However, since the variance of daily realized returns is much larger than the variance of daily expected returns, autocorrelation from this source will have little effect on our results.[4]

Because the predictions of the trading noise hypothesis are not precise, we are not interested in a detailed study of the autocorrelation structure of daily returns. However, the general behavior of the autocorrelations can help us discriminate between the trading noise hypothesis and the information hypotheses. In summary, we expect that measurement error from the bid/ask spread will lead to negative first-order autocorrelation under all three hypotheses. Neither the public nor the private information hypothesis predicts any other serial correlation, while the trading noise hypothesis predicts that daily returns will be negatively correlated beyond lag one.

Table 3 shows average autocorrelations for lags between one and fifteen days. The general procedure used to compute these averages is similar to the procedure used in table 1. Autocorrelations are estimated for individual stocks during each two-year subperiod. The first column of table 3 reports grand averages that are calculated by averaging the autocorrelations across all of the stocks within each subperiod and then averaging the ten subperiod averages. Columns 2 through 6 report the average autocorrelations for firms that have been sorted into quintiles based on their equity values at the beginning of each subperiod. Table 3 also includes standard errors of the autocorrelation estimates. These standard errors are based on the distribution of the ten subperiod averages, under the assumption that these averages are independent and identically distributed.[5]

The results in table 3 are generally consistent with the predictions of the trading noise hypothesis. All of the estimated autocorrelations from lag 2 to lag 12 are negative. Although the estimates are small in absolute magnitude, many are more than three standard errors from zero. The persistence of the negative autocorrelations suggests that trading noise is not completely corrected for at least two weeks.

The behavior of the first-order autocorrelations in table 3 is surprising. We expected measurement error within the bid/ask spread to induce negative

[4] To examine this issue in more detail, we have recomputed the autocorrelations reported in table 3 below using returns which are adjusted for day-of-the-week effects. This adjustment does not alter any of our inferences.

[5] Under the assumption that returns are serially independent, the expected value of the estimated autocorrelations for each firm is $-1/(T-1)$, where T is the number of observations used in the estimate. [See Moran (1948).] Therefore, we increase the individual autocorrelation estimates by $1/(T-1)$ before computing the subperiod and full period averages.

Table 3

Average daily autocorrelations in percent for all NYSE and AMEX stocks and for quintiles of stocks sorted by equity value.[a]

Lag	All stocks	Smallest quintile	2	3	4	Largest quintile
1	0.33	−6.42	−1.66	1.17	2.49	5.44
	(0.87)	(1.55)	(1.07)	(0.83)	(0.84)	(1.01)
2	−1.15	−1.94	−1.43	−1.28	−0.75	−0.40
	(0.15)	(0.24)	(0.20)	(0.15)	(0.14)	(0.26)
3	−1.15	−1.35	−1.39	−1.26	−1.00	−0.81
	(0.15)	(0.16)	(0.23)	(0.19)	(0.21)	(0.22)
4	−0.68	−0.85	−0.66	−0.62	−0.49	−0.78
	(0.23)	(0.16)	(0.19)	(0.29)	(0.28)	(0.34)
5	−0.28	−0.44	−0.15	−0.05	−0.15	−0.59
	(0.23)	(0.17)	(0.28)	(0.22)	(0.27)	(0.39)
6	−0.95	−0.84	−0.72	−0.84	−0.92	−1.38
	(0.24)	(0.24)	(0.25)	(0.29)	(0.26)	(0.29)
7	−0.64	−0.53	−0.53	−0.50	−0.63	−0.98
	(0.25)	(0.19)	(0.30)	(0.27)	(0.26)	(0.32)
8	−0.37	−0.14	−0.42	−0.31	−0.25	−0.73
	(0.24)	(0.15)	(0.23)	(0.30)	(0.25)	(0.40)
9	−0.45	−0.34	−0.47	−0.27	−0.57	−0.60
	(0.19)	(0.13)	(0.16)	(0.17)	(0.20)	(0.40)
10	−0.26	−0.12	−0.10	−0.15	−0.38	−0.51
	(0.15)	(0.12)	(0.15)	(0.18)	(0.19)	(0.24)
11	−0.52	−0.27	−0.54	−0.58	−0.57	−0.65
	(0.18)	(0.10)	(0.13)	(0.21)	(0.23)	(0.32)
12	−0.20	−0.35	−0.18	−0.21	−0.10	−0.19
	(0.22)	(0.12)	(0.20)	(0.24)	(0.28)	(0.31)
13	−0.15	−0.11	0.02	−0.03	−0.14	−0.47
	(0.20)	(0.12)	(0.18)	(0.20)	(0.28)	(0.35)
14	0.15	0.28	0.05	0.18	0.14	0.10
	(0.25)	(0.13)	(0.17)	(0.24)	(0.27)	(0.50)
15	0.42	0.28	0.38	0.49	0.43	0.52
	(0.15)	(0.12)	(0.17)	(0.17)	(0.21)	(0.28)

[a] The autocorrelations and standard errors (in parentheses) are reported in percent. Autocorrelations are estimated for individual firms during each of ten two-year subperiods between 1963 and 1982. These autocorrelations are averaged to compute subperiod averages. Each reported autocorrelation is the average of ten subperiod averages. The standard error is based on the distribution of the ten subperiod averages. Approximately 500 returns are used to estimate the autocorrelations for each firm in each subperiod. On average, there are about 380 firms in each quintile during each subperiod.

are not. In fact, the first-order autocorrelation for the largest quintile of stocks is 5.4%, with a standard error of 1.0%.

Although they are surprising, the positive autocorrelations also support the trading noise hypothesis. If we rule out the possibility that the reported prices contain positively correlated measurement errors, we are unable to imagine any sensible explanation of these results that does not involve trading noise. For example, suppose traders over-react to new information and this over-reaction persists for more than one day. Then tomorrow's pricing error is positively correlated with both today's information component and today's pricing error. Alternatively, suppose the market does not incorporate all information as soon as it is released. Then today's pricing error is negatively correlated with today's information and tomorrow's error is positively correlated with today's information. The positive correlation between today's information and tomorrow's error could generate positively autocorrelated returns. Since negative first-order autocorrelation induced by the bid/ask spread is smaller for the larger firms, it dominates the error-induced positive autocorrelation only in the first and second quintiles.

The results in table 3 are consistent with the trading noise hypothesis. However, since the average autocorrelations are small in absolute magnitude, it is hard to gauge their economic significance. To estimate the importance of the trading noise hypothesis, we compare daily return variances with variances for longer holding period returns. If daily returns were independent, the variance for a long holding period would equal the cumulated daily variances within the period. On the other hand, if daily returns are temporarily affected by trading noise, the longer period variance will be smaller than the cumulated daily variances.

This comparison presumes that the relative importance of both pricing errors and bid/ask errors is reduced as the holding period is increased. For example, suppose mispricing is corrected within three weeks. Then pricing errors that occur during the first ten weeks of each three-month holding period have no effect on the three-month return and errors that occur during the last three weeks have a reduced effect. If pricing errors are corrected within three weeks and bid/ask errors are corrected overnight, most of the three-month return reflects a rational assessment of the information arriving during the three-month period. When the holding period is extended to six months, this approximation becomes even more accurate. By comparing the variance of long holding period returns (which reflect information) with the variance implied by daily returns (which reflect information, pricing errors, and bid/ask errors), we can estimate the fraction of the daily variance that is caused by rational assessments of information.[6]

Table 4 reports average actual-to-implied variance ratios for holding periods of two trading days; one, two, and three weeks; and one, three, and six months. The general procedure used to compute these average ratios is similar to the procedure used in tables 1 and 3. We first compute actual-to-implied

[6] This comparison was suggested to us by Eugene Fama. Perry (1982) uses a similar approach to examine the process generating stock returns.

Table 4

Actual-to-implied variance ratios for all NYSE and AMEX stocks and for quintiles of stocks sorted by equity value.

		All stocks	Smallest quintile[a]	2	3	4	Largest quintile
Two trading days	Average ratio[b]	0.999	0.933	0.979	1.007	1.021	1.048
	Standard error[c]	0.010	0.015	0.011	0.010	0.009	0.013
	Number of firms[d]	1900.2	362.6	374.6	373.9	386.8	402.3
	Average sample size[e]	250.5	249.0	250.3	250.8	251.0	251.0
One week	Average ratio	0.966	0.853	0.928	0.979	1.005	1.053
	Standard error	0.017	0.025	0.019	0.018	0.018	0.019
	Number of firms	1899.6	362.1	374.5	373.9	386.8	402.3
	Average sample size	103.6	102.7	103.6	103.9	104.0	104.1
Two weeks	Average ratio	0.943	0.803	0.900	0.959	0.995	1.045
	Standard error	0.024	0.026	0.025	0.025	0.027	0.025
	Number of firms	1899.5	362.0	374.5	373.9	386.8	402.3
	Average sample size	51.5	50.9	51.5	51.7	51.8	51.8
Three weeks	Average ratio	0.929	0.784	0.888	0.953	0.985	1.024
	Standard error	0.022	0.026	0.026	0.023	0.024	0.024
	Number of firms	1899.5	362.0	374.4	373.9	386.8	402.2
	Average sample size	34.0	33.5	34.0	34.1	34.2	34.2
One month	Average ratio	0.906	0.773	0.874	0.931	0.959	0.983
	Standard error	0.022	0.023	0.027	0.025	0.020	0.024
	Number of firms	1898.9	361.7	374.3	373.9	386.8	402.2
	Average sample size	23.6	23.0	23.6	23.8	23.8	23.9
Three months	Average ratio	0.894	0.752	0.876	0.949	0.942	0.941
	Standard error	0.045	0.032	0.043	0.051	0.055	0.066
	Number of firms	1895.3	359.8	373.6	373.5	386.6	401.8
	Average sample size	7.7	7.3	7.7	7.8	7.9	7.9
Six months	Average ratio	0.883	0.731	0.862	0.931	0.929	0.907
	Standard error	0.102	0.062	0.086	0.109	0.117	0.129
	Number of firms	1554.2	203.4	291.8	324.8	350.5	383.7
	Average sample size	4.0	4.0	4.0	4.0	4.0	4.0

[a] Stocks are sorted in quintiles based on their equity value at the beginning of ten two-year subperiods between 1963–1982.

[b] Average ratio comparing the actual holding period variance with the variance implied by single-trading-day returns under the assumption that the one-day returns are independent. The reported ratio is the average of ten subperiod averages.

[c] The standard error of the reported average ratio. This standard error is based on the distribution of the ten subperiod average ratios.

[d] The average number of firms in each subperiod.

[e] The average number of multiple-day returns for each stock in each subperiod.

variance ratios for each stock in each two-year subperiod. This is done in four steps. For example, to estimate the weekly actual-to-implied variance ratio for a given stock in a particular two-year subperiod, we first calculate the average trading day return during the 104 weeks in that period. Next, we cumulate the daily squared deviations around this average. Then under the assumption that the daily returns are independent, we estimate the implied weekly variance by dividing this total by 104. Finally, we divide the actual weekly variance by the implied variance. The same procedure is used to estimate variance ratios for other holding periods.[7]

The first column of Table 4 reports grand averages that are calculated by averaging the estimated variance ratios across all the stocks within each subperiod and then averaging the ten subperiod averages. The averages for stocks that have been sorted into quintiles based on their equity values at the beginning of each subperiod are reported in columns 2 through 6. Table 4 also includes standard errors that are based on the distribution of the ten subperiod averages, under the assumption that the subperiod averages are independent and identically distributed.

The results in table 4 indicate that a significant fraction of the daily variance is caused by pricing and bid/ask errors. The six-month actual-to-implied variance ratio for all firms is 0.88. The six-month averages for the smallest and largest quintiles are 0.73 and 0.91, respectively. Based on these point estimates, 27% of the daily variance for the first quintile and 9% of the daily variance for the fifth quintile is eliminated in the long run. One would draw nearly identical inferences from the three-month variance ratios. This supports the assumption that bid/ask and pricing errors have relatively little effect on three- and six-month holding period returns.

Since both the pricing errors and the bid/ask errors are temporary, the six-month ratios in table 4 only allow us to make an estimate of their combined effect. However, by assuming that the variance of the bid/ask errors is zero, these ratios place an upper bound on the point estimate of the relative variance of the pricing errors. We can estimate a lower bound for this variance by combining the results in tables 3 and 4.

Suppose each day's return is made up of three independent components: a rational information component (X_t), a mispricing component (Y_t), and a bid/ask error (Z_t),

$$R_t = X_t + Y_t + Z_t. \tag{2}$$

[7]These ratios may be affected by two sources of bias. Both the actual and implied variances are estimated with error. Since we are using the ratio of these estimates our measure is biased upward. However, simulations suggest that this bias is negligible. The second source of bias may be more important. We are assuming that the expected returns are constant over each estimation period. Violations of this assumption will have little effect on the implied variances since they are based on daily variance estimates. However, changing expected returns could positively bias the actual long-term variance estimates. To reduce this effect, we limit each estimation period to two years and we limit the holding periods to a maximum of six months.

Also, suppose that the daily information components are independent and identically distributed with variance $\mathrm{var}(X_t)$. The bid/ask error in the daily return depends on the error in the current price (e_t), and the error in the previous day's price (e_{t-1}),

$$Z_t = e_t - e_{t-1}. \tag{3}$$

If the daily price errors (e_t) are independent and identically distributed, the variance and first-order autocovariance of the bid/ask errors equal

$$\mathrm{var}(Z_t) = 2\,\mathrm{var}(e_t) \tag{4}$$

and

$$\mathrm{cov}(Z_t, Z_{t-1}) = -\mathrm{var}(e_t) = -\mathrm{var}(Z_t)/2. \tag{5}$$

Therefore, the first-order autocorrelation of the bid/ask errors is -0.5.[8]

If pricing and bid/ask errors have a negligible effect on six-month returns, the six-month variance ratios in table 4 can be written as

$$V_6 = \mathrm{var}(X_t)/\mathrm{var}(R_t), \tag{6}$$

where $\mathrm{var}(X_t)$ and $\mathrm{var}(R_t)$ are the variances of the daily information component and the total daily return, respectively. Using eq. (5) and the assumption that the daily information components are serially independent, the first-order autocorrelation of the daily returns is

$$\begin{aligned} \rho_R &= \mathrm{cov}(R_t, R_{t-1})/\mathrm{var}(R_t) \\ &= \left[\mathrm{cov}(Y_t, Y_{t-1}) + \mathrm{cov}(Z_t, Z_{t-1})\right]/\mathrm{var}(R_t) \\ &= \left[\rho_{Y1}\mathrm{var}(Y_t) - \mathrm{var}(Z_t)/2\right]/\mathrm{var}(R_t). \end{aligned} \tag{7}$$

Eqs. (6) and (7) can be combined to obtain an expression for the relative variance of the pricing errors,

$$\mathrm{var}(Y_t)/\mathrm{var}(R_t) = (1 - V_6 + 2\rho_R)/(1 + 2\rho_{Y1}). \tag{8}$$

Unfortunately, we cannot observe ρ_{Y1}, the autocorrelation of the pricing errors. However, since this autocorrelation must be less than 1.0, eq. (8) gives a lower bound for the point estimate of the relative variance,

$$\mathrm{var}(Y_t)/\mathrm{var}(R_t) > (1 - V_6 + 2\rho_R)/3. \tag{9}$$

[8] This bid/ask spread phenomenon is examined in more detail by Cohen et al. (1983) and Roll (1984).

Using the average first-order autocorrelations in table 3 and the average six-month variance ratios in table 4, the upper and lower bounds on our point estimates of the relative pricing error variance for all stocks and for each quintile are:

	All stocks	Smallest quintile	2	3	4	Largest quintile
Upper bound	11.7%	26.9%	13.8%	6.9%	7.1%	9.3%
Lower bound	4.1%	4.7%	3.5%	3.1%	4.0%	6.7%

The lower bound is roughly constant across the five portfolios. This similarity is not limited to the lower bound. Differences in the relative pricing error variances will be small as long as the autocorrelation of these errors is approximately the same across portfolios.

4.3. *Implications*

The estimates in tables 3 and 4 suggest that a non-trivial fraction of the daily variance is caused by mispricing. However, pricing errors have a negligible effect on the weekend-to-weekday variance ratios in table 1. Suppose we adjust those ratios under a set of assumptions that magnifies the impact of mispricing. Specifically, assume that the variance of the weekday pricing errors and the variance of the weekday bid/ask errors are as large as the variance of weekend errors. Then the weekday and weekend returns (R_{1t} and R_{3t}, respectively) can be written as

$$R_{1t} = X_{1t} + Y_t + Z_t, \tag{10}$$

$$R_{3t} = X_{3t} + Y_t + Z_t. \tag{11}$$

Based on the average ratio for all firms in table 1, the variance of R_{3t} is 10.7% larger than the variance of R_{1t},

$$\operatorname{var}(X_{3t} + Y_t + Z_t) = 1.107 \operatorname{var}(X_{1t} + Y_t + Z_t). \tag{12}$$

The average six-month variance ratio for all firms in table 4 is 0.88. To magnify the effect of mispricing further, assume that this ratio applies to the weekend variance,

$$\operatorname{var}(X_{3t}) = 0.883 \operatorname{var}(X_{3t} + Y_t + Z_t). \tag{13}$$

Under the assumption that the information and mispricing components are independent, eqs. (12) and (13) can be combined to eliminate the bid/ask and pricing error variances,

$$\mathrm{var}(X_{3t}) = 1.123\,\mathrm{var}(X_{1t}). \tag{14}$$

Eliminating the effect of these errors increases the average weekend-to-weekday variance ratio for all firms by less than 2%. This effect varies from less than 1% for the largest quintile of stocks to less than 6% for the smallest quintile. Bid/ask and pricing errors also have a negligible effect on the two- and four-day variance ratios in table 1 and on the exchange holiday ratios in table 2.

It appears that the low daily variance ratios are caused by a reduction in the arrival of information when the exchanges are closed. Moreover, the exchange holiday variances suggest that private information causes most stock price changes.[9]

5. Summary and conclusions

Asset returns display a puzzling difference in volatility between exchange trading hours and non-trading hours. For example, we estimate that the per hour return variance was about 70 times larger during a trading hour than during a weekend non-trading hour, on average, over all stocks listed on the New York and American Exchanges from January 1963 through December 1982.[10]

We consider three factors that might explain the high trading-time variances. First, the arrival of public information may be more frequent during the business day. Second, private information may be much more likely to affect prices when the New York exchanges are open. Third, the process of trading may induce volatility.

Our results indicate that, on average, approximately 4 to 12% of the daily variance is caused by mispricing. However, even if we assume that pricing errors are generated only when the exchanges are open, these errors have a trivial effect on the difference between trading and non-trading variances. We conclude that this difference is caused by differences in the flow of information during trading and non-trading hours. Moreover, small return variances over exchange holidays suggest that most of this information is private.

Appendix

If we are willing to ignore sampling error (and just assume that sample estimates are population values), we can deduce additional information about

[9] In the appendix we develop some implications under the assumption that the information and error components are not independent.

[10] This estimate is based on the variance ratio for weekends in table 1.

the correlation between information and mispricing and about the quantity of information produced on a non-market business day (such as a Wednesday in 1968).

First, define $W_t = Y_t + Z_t$ as the sum of the mispricing component and the bid/ask error. Define X_t as the information-induced return for one day. The variance ratio for table 4 can be written as

$$a = \mathrm{var}(X)/\mathrm{var}(X+W). \tag{A.1}$$

Solving (A.1) with a = 0.883 (from table 4),

$$\sigma_X/\sigma_W = \frac{\rho_{XW} + \sqrt{\rho_{XW}^2 + 0.1325}}{0.1325}, \tag{A.2}$$

where ρ_{XW} is the contemporaneous correlation between X and W, and σ is the standard deviation. Thus, if there is *no* correlation between information and mispricing ($\rho_{XW} = 0$),

$$\sigma_X/\sigma_W = 2.75.$$

In principle, we could have a low information-to-mispricing variance ratio. For example, if $\rho_{XW} = -1$, σ_X/σ_W is only 0.48. At the other extreme, if $\rho_{XW} = +1$, $\sigma_X/\sigma_W = 15.6$.

The variance ratios for business days which are not trading days can be written as

$$b = \mathrm{var}(kX+W)/\mathrm{var}(X+W), \tag{A.3}$$

where $k^2 - 1$ $(1 \leq k \leq \sqrt{2})$ is the information produced during a business-day-exchange holiday. For $b = 1.145$ (from table 2) and $\sigma_{XW} = 0$, (A.1) and (A.3) imply $K = 1.079$; i.e., only about 16 percent of a normal business day's information was produced on the 1968 Wednesday business days which were exchange holidays.

Going one step further, we can combine (A.1) and (A.3) to eliminate σ_{XW}. This provides an expression for the ratio of mispricing to information variance as a function of k,

$$q = \frac{\mathrm{var}(W)}{\mathrm{var}(X)} = \frac{ak^2 + (1-a)k - b}{a(k-1)} = k + \frac{1}{a} + \frac{1-b}{a(k-1)}. \tag{A.4}$$

We note that $\partial q/\partial k = 1 + ((b-1)/a(k-1)^2)$, which is positive because $b > 1$. The function q has a zero at $k_L = \{(a-1) + [(a-1)^2 + 4ab]^{1/2}\}/2a$. For values of k greater than $k_U = \{(2a-1) + [(2a-1)^2 + 4a(b-a)]^{1/2}\}/2a$, the mispricing variance exceeds the information variance, i.e., $q > 1$.

Table A.1

Exchange holiday information consistent with information variance exceeding mispricing variance (% of normal day).

Actual-to-implied variance ratio, a	Ratio of two-day exchange holiday return variance to normal day return variance, b					
	1.05	1.10	1.15	1.20	1.25	1.30
	Information produced on exchange holiday as a percentage of information produced, on a normal day, consistent with information variance exceeding mispricing variance[a]					
0.75	5.78 to 9.88	11.59 to 19.56	17.41 to 29.07	23.26 to 38.44	29.12 to 47.68	34.99 to 56.82
0.80	5.60 to 9.86	11.21 to 19.48	16.84 to 28.90	22.48 to 38.15	28.13 to 47.26	33.80 to 56.25
0.85	5.43 to 9.84	10.87 to 19.40	16.32 to 28.73	21.78 to 37.87	27.24 to 46.86	32.72 to 55.70
0.90	5.27 to 9.82	10.55 to 19.32	15.84 to 28.87	21.13 to 37.60	26.43 to 46.46	31.73 to 55.17
0.95	5.13 to 9.79	10.27 to 19.24	15.40 to 28.40	20.54 to 37.34	25.69 to 46.08	30.83 to 54.66

[a] The lower and upper bounds are given by
$K_L = 100\{(a-1)+[(a-1)^2+4ab]^{1/2}\}/2a$ and $K_U = 100\{(2a-1)^2+4a(b-a)]^{1/2}\}/2a$.

The information variance exceeds the mispricing variance only when k is between k_U and k_L. For our estimated parameters, $k_L = 1.074$ and $k_U = 1.130$. Under the assumption that the estimated values are population values, the information variance would exceed the mispricing variance if the information produced on a 1968 Wednesday was between 15.3 and 27.7 percent of that produced on a normal day. In fact, depending on the values of a and b, the range can be even narrower than $k_U - k_L$ because the implied correlation coefficient in (A.3) must lie between -1 and $+1$. From (A.1), the correlation between the information and mispricing components is obtained using the solution to (A.4),

$$\rho_{XW} = (1-a-aq)/2a\sqrt{q}. \tag{A.5}$$

For instance, $q = 0$ is clearly ruled out by (A.5) unless $a = 1$. Thus, the lower bound on k must exceed k_L.

In general, the restriction on the correlation, $-1 \leq \rho_{XW} \leq 1$, implies from (A.5)

$$1+\sqrt{1/a} > \sqrt{q} > \sqrt{1/a}-1. \tag{A.6}$$

For our estimates $a = 0.883$, $b = 1.145$, $\sqrt{q}$ has a lower bound of 0.06419, implying a lower bound on k of 1.0745 (which is slightly higher than k_L).

There are sampling errors in the estimates of a and b. Thus, although the range of k where the information variance is larger than the mispricing variance is rather narrow for our point estimates, it could be much larger for other values of a and b. Table A.1 gives the range $100(k_U^2 - 1)$ to $100(k_L^2 - 1)$

for other values of a and b which could be conceivable given the sampling error.

As the other variance ratio b increases, the interval widens. But even for $b = 1.30$, the information variance is larger than the mispricing variance only if the amount of information produced on a 1968 Wednesday is *less* than about 56 percent of the information produced on a normal business day.

The results are rather insensitive to the variance ratio a. Also, for all the values in the table, the correlation between X and W is negative. It ranges from -0.33 for $a = 0.75$, $b = 1.05$, to -0.47 for $a = 0.95$, $b = 1.30$.

References

Ariel, Robert A., 1986, A monthly effect in stock returns, Journal of Financial Economics, forthcoming.

Christie, Andrew A., 1981, On efficient estimation and intra-week behavior of common stock variances, Unpublished working paper, April (University of Rochester, Rochester, NY).

Cohen, Kalman J., Gabriel, A. Hawawini, Steven F. Maier, Robert A. Schwartz and David K. Whitcomb, 1983, Friction in the trading process and the estimation of systematic risk, Journal of Financial Economics 12, Aug., 263–278.

Fama, Eugene F., 1965, The behavior of stock market prices, Journal of Business 38, Jan., 34–105.

French, Kenneth R., 1980, Stock returns and the weekend effect, Journal of Financial Economics 8, March, 55–70.

Granger, Clive W.J. and Oskar Morgenstern, 1970, Predictability of stock market prices (Heath-Lexington, Lexington, MA).

Moran, P.A.P., 1948, Some theorems on time series, Biometrika 35, 255–260.

Oldfield, George S., Jr. and Richard J. Rogalski, 1980, A theory of common stock returns over trading and non-trading periods, Journal of Finance 35, June, 729–751.

Perry, Phillip R., 1982, The time-variance relationship of security returns: Implications for the return generating stochastic process, Journal of Finance 37, June, 857–870.

Roll, Richard, 1984, A simple implicit measure of the bid/ask spread in an efficient market, Journal of Finance 39, Sept., 1127–1139.