Division of Research
School of Business Administration

April 1989

POST-EARNINGS-ANNOUNCEMENT DRIFT: DELAYED PRICE RESPONSE, OR RISK PREMIUM?

Working Paper #592

Victor L. Bernard
The University of Michigan
and
Jacob K. Thomas
Columbia University

ABSTRACT

Prior studies have documented that subsequent to the announcement of unexpectedly high (low) earnings, estimated cumulative excess stock returns continue to drift up (down). This study is intended to help assess whether the post-earnings-announcement drift reflects a delayed price response, or a failure to adjust returns fully for risk.

Several of the results are difficult to reconcile with explanations based on incomplete risk adjustment. They are consistent with a delayed response to information.

The authors appreciate the suggestions and comments of Ray Ball, Bruce Bublitz, Werner DeBondt, George Foster, Robert Holthausen, Jack Hughes, Gene Imhoff, Roger Kormendi, S. P. Kothari, Richard Leftwich, Laurentius Marais, Rick Mendenhall, Jim Noel, Bob Officer, Steve Penman, Chris Olsen, Jay Ritter, Abbie Smith, Katherine Schipper, Tom Stober, Wanda Wallace, Ross Watts, Paul Zarowin, and participants at workshops at the University of California at Berkeley, the University of Chicago, Columbia University, the University of Connecticut, New York University, and the Ohio State University. We also thank Carol Frost, Kathy Petroni, and Jim Wahlen for their research assistance, and Gautum Kaul for supplying data on certain macroeconomic variables. Professor Bernard is grateful for the financial support of the University of Michigan and Dow Corning Corporation.

An ancestor of this paper was discussed at the Universities of Melbourne, New South Wales, Queensland, and Sydney, as well as Macquarie University and Kurring-gai College. We appreciate comments received during those discussions.

1. INTRODUCTION

Beginning with the pioneering work reviewed by Fama [1970], an abundance of evidence was accumulated in support of the hypothesis that stock prices reflect all publicly available information. Recently, however, Summers [1985] has questioned the power of most prior tests to reject that hypothesis, and a variety of newer studies have found evidence that could be inconsistent with that hypothesis.[1] One puzzling piece of evidence is commonly referred to as "post-earnings-announcement drift": stock prices appear to react with a lag to the announcement of earnings news. Foster, Olsen, and Shevlin [1984] examine this phenomenon over the 60 days subsequent to quarterly earnings announcements. They find that a long position in stocks with unexpected earnings in highest decile, combined with a short position in stocks in the lowest decile, yields a positive mean annualized "excess"[2] return of about 25 percent, before transactions costs, from 1974 through 1981. The same strategy yields an even higher average "excess" return (40 percent) when restricted to stocks for small firms (even after controlling for the well-known Banz-Reinganum size effect). The evidence is particularly puzzling because it is difficult to understand how the market could fail to react quickly and completely to information as visible and as readily available as earnings data.[3]

---

[1]Poterba and Summers [1987], Ou and Penman [1988,1989], Landsman and Ohlson [1987], DeBondt and Thaler [1985,1987], Harris and Ohlson [1988], Lehmann [1988], and Lakonishok and Vermalean [1988] present evidence consistent with inefficiency. Although failure to have fully adjusted for risk is always an alternative explanation for such evidence, most of these studies take steps to mitigate that concern.

[2]"Excess" returns represent the difference between actual returns and expected returns, where the expectation is conditioned on the actual return on a control portfolio (e.g., the market portfolio or a size-matched portfolio).

[3]Patell and Wolfson [1984] and Brown, Clinch, and Foster [1987] show that unusual price activity surrounding earnings announcements lasts only minutes, or perhaps hours. However, Patell and Wolfson examine only large firms, and, in any case, these research designs would not capture the slow post-earnings-announcement drift discussed here.

The problem in interpreting evidence of this kind is that there are two alternative explanations for the apparent excess returns. One class of explanations suggests that, especially for small firms, prices do not respond fully and immediately to new information. This might occur either because certain costs (such as the costs of transacting, or the cost of imperfect diversification) exceed gains from immediate exploitation of information for a sufficiently large number of rational traders,[4] or because traders do not act rationally. A second class of explanations suggests that, because the capital asset pricing model (CAPM) used to calculate excess returns is either incomplete or misestimated, researchers fail to adjust raw returns fully for risk. As a result, the so-called excess returns are nothing more than fair compensation for bearing risk that is priced but not captured by the CAPM estimated by researchers. In the case of post-earnings-announcement drift, firms with unexpectedly high (low) earnings may become more (less) risky on some unrecognized dimension.[5]

Given our current understanding of capital asset pricing, it is difficult to distinguish between the two alternative explanations (see Ball [1988]). However, this paper attempts to specify what forms of CAPM misspecification (including either omission or misestimation of risk factors), on the one hand, or delayed price response, on the other, would be consistent with the data. Such information should be useful in assessing the plausibility of the alternative explanations.

The paper presents a variety of evidence concerning the behavior of the drift for NYSE, AMEX, and NASDAQ firms from 1974 through 1986. We find several features of the evidence difficult to attribute to a failure to adjust fully for risk:

---

[4]This explanation also begs the question of why market makers (e.g., specialists) would not move prices immediately in response to new information, even when transactions cost impede most traders from acting.

[5]Finally, a third explanation-- bias resulting from research design problems other than CAPM misspecification-- is always possible. However, FOS do much to dismiss this possibility.

1) Shifts in systematic risk (beta) are too small to explain the drift. Ball, Kothari, and Watts [1988] suggest that beta shifts explain post-earnings-announcement drift, and that beta estimation procedures used in prior research failed to detect the beta shifts. Using the Ball-Kothari-Watts estimation procedure, we also detect beta shifts around the time of earnings announcements. However, the magnitude of the shifts falls far short of what would be necessary to explain post-announcement drift.

2) Other risk factors also fail to explain the drift. Chen, Roll, and Ross [1986] identify five factors other than beta that appear to affect asset pricing. However, we find no evidence that a trading strategy based on post-announcement drift is risky along any of these five dimensions.

3) Consistency of drift over time. If post-announcement returns on a zero-investment portfolio represent a premium for an unidentified risk, then that risk must appear periodically in the form of a loss; moreover, the expected cost of the loss (in terms of utility) must be commensurate with the value of the premium. However, a zero-investment trading strategy based on publicly available earnings information generates a gain (i.e., positive excess return) in 13 of 13 years, and in 45 of 50 quarters. Furthermore, in the five quarters where losses are experienced, they are not large; cumulative losses are exceeded by cumulative gains by a factor of 35 to one. Unless these losses are many times more "costly" (in terms of utility) than the gains of other quarters are "valuable," these data are difficult to reconcile with capital asset pricing theory.

4) Low raw returns on "bad news" stocks. According to capital asset pricing theory, expected raw (total) returns on risky assets can be less than risk-free returns only under special conditions. However, over the five days after the announcement of low unexpected earnings, mean raw returns are *negative* for small firms, and less than Treasury bill rates for the overall sample. Over the

two months subsequent to the announcement, mean raw returns are only slightly in excess of Treasury bill rates. Thus, if the post-announcement drift is to be explained as a product of CAPM misspecification, one must find it plausible that the riskiness of "bad news" firms declines so much (or their value as hedges increases so much) that, after the earnings announcement, their cash flows deserve to be priced at something close to the Treasury bill rate, and for short periods, at a rate less than the Treasury bill rate.

5) The drift is bounded, as if the response to earnings is inhibited once prices are within a certain range of the "appropriate" level. As unexpected earnings grows larger, the contemporaneous response to earnings also increases. The same is true of the delayed response, but only to a certain point, after which the drift remains constant, no matter how large is the unexpected earnings. This feature of the data is consistent with what would occur if some cost inhibits the response to earnings, once prices are within a certain range of the "correct" price. Interestingly, the implied range varies inversely with firm size, as do transactions costs, and is within the bounds of transactions costs for individual investors (including bid-ask spreads and commissions). This feature of the data can be reconciled with CAPM misspecification only if there is a "kink" in the relation between unexpected earnings and the unidentified risk factor for which it proxies. That is, unexpected earnings would have to be correlated with the risk factor up to a certain point, but then additional increases in unexpected earnings could no longer be correlated with increases in risk.

6) Estimated excess returns on "bad news" firms are more than twice as large as those for "good news" firms. This feature of the data is not predicted by existing explanations based on CAPM misspecification. However, it is consistent with impediments to the immediate impounding of "bad news," due to the difficulties involved with short sales.

7) Other features of the data. We confirm FOS's finding that the drift is largest for small firms, and Watts' [1978] finding that the drift is temporary. (Fifty to 75 percent of the drift occurs within three months; it lasts about six months for large firms and nine months for small firms.) Although this evidence could potentially be consistent with CAPM misspecification, it begs two unanswered questions. First, why would unidentified risk factors be more important for small firms? Second, if a risk shift occurs, why is it reversed more slowly for small firms?

Our conclusion is that the evidence cannot plausibly be explained by the failure to adjust returns fully for risk. On the other hand, the evidence is consistent with delays in the response to earnings information, perhaps as the result of costs that impede trading for some investors.

The remainder of this paper is organized as follows. Section 2 summarizes the current state of understanding of post-earnings-announcement drift, and presents arguments for delayed price response and CAPM misspecification as explanations for the drift. Section 3 describes the sample and some of the methods used in our empirical tests, and presents the results of those tests. A discussion of the evidence and some conclusions are presented in Section 4.

2. POST-EARNINGS-ANNOUNCEMENT DRIFT: THE NATURE OF THE PHENOMENON

Post-earnings-announcement drift was documented as early as Ball and Brown [1968], and has been studied many times (examples include Watts [1978] and Rendleman, Jones, and Latane [1982]). Foster, Olsen, and Shevlin [1984] (FOS) replicated the essential results of prior papers, and also presented a number of new findings. The well-documented finding replicated by FOS is duplicated here in Figure 1. Figure 1 summarizes the results of a trading strategy, labeled by FOS as the EBM (earnings-based-model) strategy. To implement the strategy, FOS first estimate unexpected earnings for a sample of NYSE and AMEX firms. The unexpected earnings, scaled by the standard deviation of prior forecast errors, is then compared to the cross-sectional distribution of scaled unexpected earnings for the prior quarter. Based on their standing relative to that distribution, firms are assigned to one of ten portfolios. The excess (size-adjusted) returns on those ten portfolios are then plotted over the 120 trading days surrounding the earnings announcement date. The prediction is that, even after the earnings are announced, the firms with positive (negative) unexpected earnings will generate positive (negative) excess returns.

Figure 1 confirms that the estimated post-earnings-announcement excess returns vary monotonically with unexpected earnings. A long position in EBM portfolio 10 (that with the highest unexpected earnings), combined with a short position in EBM portfolio 1, yields an estimated excess return of 6.31 percent over the 60 trading days after the earnings announcement, or about 25 percent on an annualized basis. The question is whether this estimated excess return reflects an incomplete adjustment for risk, or a delayed price response.

2.1 THE CASE FOR CAPM MISSPECIFICATION

Ball [1988] argues that there is good reason to believe stock markets to be efficient on *a priori* grounds, because such markets are "paradigm examples of competition." Ball [1978] explains why, even if prices respond quickly and completely to information, trading strategies based on earnings numbers might give the *appearance* of profitability, as a result of misspecifications in the CAPM used to measure the excess returns. There is some evidence consistent with this explanation in Ball, Kothari, and Watts [1988] and FOS [1984].

Ball, Kothari, and Watts [1988] (hereafter, BKW) suggest that beta (systematic risk) shifts upward (downward) for firms with high (low) unexpected earnings. Since some prior studies assumed for purposes of estimation that betas were stationary, this causes an upward (downward) bias in estimated excess returns. To overcome the difficulty, BKW use an estimation approach that permits betas to shift annually.

BKW do detect beta shifts in the suggested directions during the year of the unexpected earnings. However, most of the shift is reversed by the post-announcement period, which is what matters for purposes of explaining the phenomenon of interest. Moreover, there is little difference between the post-announcement betas of high- and low-unexpected earnings firms, which is what would be required to explain the drift generated in a design like that used in Foster, Olsen, and Shevlin [1984].

When BKW measure excess returns while allowing for beta shifts, the observed post-announcement drift is no longer significant. However, since their sample includes primarily large firms, and focuses on annual, rather than shorter post-announcement periods, their failure to document a significant drift may in part reflect a lack of power. Whether the BKW methodology can explain drifts of the magnitude apparent in a more powerful design, or whether risk measures other than beta could explain the drift, will be addressed later in this paper.

A second source of evidence consistent with CAPM misspecification is the major result in Foster, Olsen, and Shevlin [1984]. FOS contrasts two alternative approaches to analyzing the post-announcement behavior of stock returns. The first approach is that used to generate Figure 1: the earnings-based-model (EBM) strategy. The second approach assigns firms to portfolios on the basis of a different proxy for unexpected earnings: the estimated excess stock returns over the 60 days prior to and including the earnings announcement day.[6] This approach was labeled the SRM (security-return model) strategy. For convenience, the FOS SRM results are presented in Figure 2. The essential result is that there is no indication of post-earnings-announcement drift. Thus, post-announcement drift is evident only for a subset of measures of unexpected earnings.

The results of the SRM tests in FOS have been interpreted as indicating post-announcement drift reflects some problem in risk measurement: "Using the (SRM) method of forming portfolios yields no unusual return behavior following the earnings announcement and suggests again that the results of previous studies are caused by a misspecified pricing model" (Dyckman and Morse [1986, p. 58]). The same conclusion is not drawn explicitly by FOS, but could be inferred from some of their discussion. FOS explain that the EBM tests are vulnerable to certain problems in risk adjustment discussed by Ball [1978]; the SRM tests were motivated as an approach to mitigate these problems. Given that the drift vanishes when one moves from the EBM tests to the SRM tests, it is understandable that readers would infer that the drift in the EBM tests reflects some problem in risk adjustment.

Bernard and Thomas [1989] suggest that any such inference is unwarranted. The reason is that the FOS results are consistent not only with certain explanations under which the drift represents a risk premium, but also with certain other explanations under which the drift is a delayed price response.

---

[6]FOS also examined tests based on excess returns over the two-day window ending on the earnings announcement day, and obtained similar results. We do not focus on these short-window tests, however, because they suffer from a bias that would tend to obscure the drift, as discussed in the next footnote.

The intuition in Bernard and Thomas [1989] is as follows. Assume that stock prices react completely and immediately to earnings news upon its announcement for some firms, but not for others. For the latter firms, the post-announcement drift observed in the EBM tests represents the completion of that response. The question is whether that "delayed" price response might not be observed in the SRM tests, where portfolios are formed based on a ranking of excess returns prior to and including the earnings announcement day. First, note that the extreme SRM portfolios, with large stock price movements prior to and including the announcement day, are most likely to include firms whose "timely" response to earnings is complete; therefore, one would not expect to find a detectable post-announcement drift for those portfolios. Firms for which there *is* a "delayed" response would be scattered across the less extreme SRM portfolios, where there is some mixing of "good" and "bad" earnings news. Since the average unexpected earnings is small for such non-extreme portfolios, it is not clear that any drift for these portfolios would be empirically detectable. The conclusion is that there may be no detectable drift for any of the SRM portfolios, even if stock prices respond with a delay to earnings news for some firms.[7]

Bernard and Thomas suggest that an appropriate interpretation of FOS's SRM test is *not* that it discriminates between CAPM misspecification and delayed price response. Instead, it imposes restrictions on the nature of CAPM misspecifications, on the one hand, and delayed price response, on the other. The permissible delayed price response is one that is largest, for those firms whose "timely" response is least consistent with the earnings

[7]There are other reasons why no drift would be observed in the SRM tests, even if prices respond to earnings announcements with a lag. One possible reason is that the SRM tests lack statistical power, because they do a poor job of partitioning firms in terms of unexpected earnings. Bernard and Thomas [1989] indicate that while this argument may have considerable merit, it could at best be a partial explanation. Another reason is that, given non-zero bid-ask spreads, the portfolio formation procedure in the SRM tests introduces a bias that would tend to obscure the drift. This bias is important in interpreting the results of SRM tests where firms were ranked on excess returns in the two days surrounding the earnings announcement, but not necessarily for the SRM tests discussed above. For more detail, see Bernard and Thomas [1989].

signal. Evidence in Bernard and Thomas [1989] is consistent with that depiction, as well as forms of CAPM misspecification with the same empirical manifestation.

2.2 THE CASE FOR A DELAYED PRICE RESPONSE

That post-earnings-announcement drift could represent a delayed response to information has been viewed as plausible by some academics. For example, Lev and Ohlson [1982, p. 284] describe the evidence of post-earnings-announcement drift as the "most damaging to the naive and unwavering belief in market efficiency." However, it is difficult to explain why the market could fail to respond immediately to earnings information.

One possibility, that information processing costs impede the adjustment process, is difficult to believe in this case. Earnings information is freely available and quickly disseminated. In addition, some firms supply weekly reports to investors, summarizing precisely the information needed to implement the strategies studied in this paper; the reports are based directly on the academic research.

A second possibility is that the market fails to appreciate the full implications of earnings information. For example, the market may fail to form an unbiased expectation of future earnings, immediately upon revelation of current earnings; some portion of the response may not occur until analysts forecasts are revised, or the future earnings are realized.[8] This too is difficult to believe, since it would imply that the market has failed to "learn" from its past valuation errors, even over many years. Nevertheless, it must be admitted that, while Kormendi and Lipe [1987] indicate that responses to current earnings

[8]Clearly, an efficient market may resolve uncertainty about the implications of a previously-released earnings number when future earnings are released. Nevertheless, regardless of how much uncertainty surrounds current earnings, stock prices in an efficient market should immediately reflect an unbiased expectation of future earnings. If information uncertainty is not "priced out," this implies no predictable post-announcement drift. If information uncertainty is priced out and investors are risk averse, this implies positive post-announcement drift for both good and bad earnings news, which is inconsistent with the data.

reflect at least some of the implications for future earnings, there is no strong evidence that the immediate response to current earnings is complete.

A third possibility is that some costs inhibit a complete and immediate response to earnings news. Examples could include direct transactions costs, the costs of selling short, and the costs of implementing a strategy (including opportunity costs). When direct transactions costs are defined to include both bid-ask spreads and commissions, they are about 5 percent (2 percent) for small (large) stocks (Stoll and Whaley [1983]); these amounts are larger than the 60-day excess returns to a single position (long or short) documented by FOS. One difficulty with this explanation is that transactions costs for large investment houses are can be only a fraction of one percent.[9] Nevertheless, since post-announcement drift is most clearly evident for small stocks, where large investment firms may not represent the marginal transactor, and because direct transactions costs are not the only potential costs, we will return to a more detailed discussion of this possibility.

2.3 THE CURRENT STATE OF UNDERSTANDING: SUMMARY

Although there are good reasons to believe *a priori* that post-earnings-announcement drift could be explained as the product of research design flaws, including the failure to account fully for risk, there is no compelling evidence indicating that to be the case. Ball, Kothari, and Watts [1988] suggest that beta shifts may explain the drift, but their suggestion has yet to be tested within the context of a powerful research design. Foster, Olsen, and Shevlin [1984] provide evidence that has been interpreted to indicate that the drift reflects CAPM misspecification, but Bernard and Thomas [1989] suggest that such an interpretation is unwarranted. Thus, whether post-announcement drift represents a risk premium or a delayed price response remains an open question.

---

[9] Lehmann [1988] assumes that for such traders, the transactions costs are only 1/10 of one percent, equal to the amount of the transfer tax.

3. EMPIRICAL TESTS

3.1 OVERVIEW

This section presents a series of empirical tests designed to assess the relative plausibility of CAPM misspecification as an explanation for post-earnings announcement drift. We first present some descriptive evidence about the magnitude of the drift, its longevity, and its relation to firm size. We then examine (1) whether shifts in betas might explain the drift, (2) whether other measures of risk used in tests of arbitrage pricing could explain the drift, (3) whether any risk associated with trading strategies based on earnings information ultimately surfaces in the form of a loss, and (4) whether raw returns for firms with negative unexpected earnings are less than the risk-free rate (or even negative) in the post-announcement period.

Since much of the evidence is difficult to explain as a product of incomplete risk adjustment, we examine (5) whether the data are consistent with behavior predicted when the drift represents a delayed price response caused by transactions costs, or other costs with certain characteristics, and (6) whether the drift is larger for "bad news," which can be exploited by some traders only through short-selling.

3.2 SAMPLE AND ESTIMATION PROCEDURES

3.2.1 Sample selection.

The sample includes 84,792 firm-quarters of data for NYSE/AMEX firms for 1974-1986. We also conduct some supplementary tests based on 15,457 firm-quarters of data for over-the-counter stocks on the NASDAQ system for 1974-1985. Criteria for inclusion in the sample are the same as those used by FOS, who studied NYSE/AMEX firms for the period 1974-1981. We require that the firm be listed on the CRSP files, and have at least 10 consecutive quarters of earnings data on Compustat. Our NYSE/AMEX sample includes only firms that appeared on any of the Compustat files released from

1982 through 1987.[10] Since firms included in earlier files but dropped from Compustat before 1982 are excluded from the sample, there is a potential for a survivorship bias in the first half of our data set. However, FOS conducted tests which indicated that post-announcement drift is not sensitive to this form of bias. Moreover, our conclusions are insensitive to whether we include or exclude "nonsurvivors" dropped from the Compustat files between 1982 and 1987.

3.2.2 <u>Estimation of excess returns.</u>

For the NYSE/AMEX sample, cumulative excess returns are calculated using an approach like that of FOS. FOS use a companion portfolio approach designed to control for the Banz–Reinganum size effect.[11] Under this approach, excess returns are calculated as follows:

$$ER_{jt} = R_{jt} - R_{pt}$$

where ER_{jt} = excess return for firm j, day t;

R_{jt} = raw return for firm j, day t;

R_{pt} = equally–weighted mean return for day t on the NYSE/AMEX firm size decile that firm j is a member of at the beginning of the calendar year.

In our tests based on excess returns, we preserve comparability with FOS, and sum excess returns over time to obtain cumulative excess returns (CERs). A problem with summing excess returns over time that it implicitly assumes daily rebalancing, and leads to an upward bias in the returns cumulated over long periods (Blume and Stambaugh [1983],

---

[10] The NASDAQ sample was selected from the 1987 Compustat file.

[11] This approach to measuring excess returns makes no attempt to control for systematic risk. Since our conclusions are based on comparisons of excess returns on high and low unexpected earnings portfolios, this introduces a bias if systematic risk differs between those two. We test for such a possibility in section 3.3.2.

Roll [1983]). However, since this bias affects both our "treatment" and "control" portfolios, there is no reason to expect any bias in our estimated excess returns. Nevertheless, we have also conducted some analyses patterned after the above approach, but based on compounded returns; those analyses indicate that the difference between excess returns on extreme "good news" and "bad news" firms is not sensitive to the choice between compounded and summed excess returns.[12] In addition, we describe in section 3.2.5 an alternative excess return calculation that is free from the bias described by Blume and Stambaugh [1983].

Observations were excluded from the analysis if the return for the earnings announcement day was missing on CRSP, or if the CRSP returns series did not encompass the 160 trading days surrounding the earnings announcement.

3.2.3 Estimation of standardized unexpected earnings (SUE).

Procedures for estimating unexpected earnings were patterned after those used by FOS for their EBM Model 2. That is, earnings were forecasted by estimating the Foster [1977] model with historical data.[13] The difference between actual and forecasted earnings

---

[12] We first compounded raw returns over time for individual firms, and then subtracted the compounded market return for the value-weighted NYSE index over the same period. We then formed ten portfolios, based on rankings of standardized unexpected earnings (SUE), within groups of large, medium, and small firms. Finally, we calculated the excess returns to a long position in the highest SUE decile and a short position in the lowest SUE decile. The excess returns over 60 post-announcement days for large, medium, and small firms were 2.7 percent, 4.8 percent, and 5.0 percent, respectively. Comparable amounts to be reported in the primary analysis are 2.8 percent, 4.5 percent, and 5.3 percent.

[13] The Foster model assumes that earnings follow a first-order autoregressive process in seasonal differences. FOS indicate [p. 582] that they used a maximum of 20 observations to estimate the Foster model. We used a maximum of 24 observations. FOS indicate [p. 581] that firms were included in the sample even if only 10 consecutive quarters of data were available. We retained such firms also, but where fewer than 16 observations were available, we did not attempt to estimate the Foster model. Instead, we assumed that earnings followed a seasonal random walk. FOS indicate [p. 582] that they obtained essentially the same results when this model was substituted for the Foster model.

was then scaled by the standard deviation of forecast errors over the estimation period. We refer to this scaled amount as standardized unexpected earnings, or SUE.

3.2.4. Portfolio assignment.

Holthausen [1983] and FOS describe a bias that is introduced when firms are assigned to portfolios. When those assignments are based on rankings of unexpected earnings within the distribution for *all* firms, including some that have *not yet announced* earnings for the quarter, there is a hindsight bias that tends to magnify the drift. We adopt the procedure used by FOS to overcome that bias. It involves assigning firms to portfolios, on the basis of their standings relative to the distribution of unexpected earnings in the *prior* quarter.

3.2.5 Alternative excess return calculation: continuously balanced SUE strategy. The strategy studied by FOS and replicated here offers the advantage of controlling for the Banz–Reinganum size effect, and maintaining zero net investment at all times. However, it could be implemented precisely only by taking new positions in size control portfolios almost every day; each control portfolio contains hundreds of stocks, and changes in content each year. This is not necessarily a serious issue, since the control portfolios could be viewed simply as a benchmark for performance evaluation, as opposed to positions that would literally be assumed by the trader. However, if the trader does not literally take positions in the control portfolios, then the FOS strategy is no longer a zero-investment strategy at all points in time. That is, the trader will at times have more long positions in "good news" stocks than short positions in "bad news" stocks, or vice versa.

The specific earnings number used was earnings before extraordinary items and discontinued operations.

This raises the question of whether the FOS strategy represents an implementable zero-investment strategy.

To assess the sensitivity of our results to these issues, we report some results based on a zero-investment strategy that would be easier to implement. Since it involves having the same amount invested in "good news" and "bad news" firms at all points in time, we label this strategy the "continuously balanced SUE strategy." (To differentiate it, we will sometimes label the FOS approach the "FOS control portfolio SUE strategy.")

The "continuously balanced SUE strategy" works as follows. On a given trading day, we identify any firms that announced earnings, and where standardized unexpected earnings falls in the upper quintile ("good news") or lower quintile ("bad news") or the prior quarter distribution. If both "good news" and "bad news" firms exist for that day, we assume a long position in the former, and a short position in the latter. The long (short) positions are initially equally-weighted across the available good (bad) news firms, with the total amount of the long position exactly offsetting the total amount of the short position. We then compute buy-and-hold (i.e., continuously compounded) returns on each of the stocks in the long and short position, over the 60 trading days subsequent to the earnings announcement.

On 14 percent of trading days, there were either no new "good news" or no "bad news" firms available, and so no match could be created. In such a case, we "wait" until a match becomes available. For example, if two "good news" firms announced earnings on day 1, but no "bad news" firms announced, we would wait until at least one "bad news" firm announced earnings. If the first available "bad news" firm announced on day 4, it would be matched with all "good news" firms announcing from days 1 through 4, and we would then compound returns from day 5 through day 64. In 97 percent of all cases, a match became available within two days.

To provide some control for the Banz-Reinganum size effect, this matching process was always conducted within groups of small, medium, and large firms.[14] Since we segregate firms into only three size groups, while the FOS control portfolio strategy utilizes ten size groups, the size control used here is not as precise. However, assuming that smaller firms are as likely to announce "bad news" as "good news," this introduces no bias in the results.[15] Moreover, by using only three size groups, we increased the probability of finding matches of "good news" and "bad news" firms within a short period of time.

The "continuously balanced SUE strategy" is much easier to implement than that used by FOS, but would still be costly to the extent that short selling must be used. However, for investors who already own the stocks that announce "bad news," this strategy should be easy to implement.

3.3 EMPIRICAL RESULTS

3.3.1 Descriptive results

Magnitude of the drift. FOS [1984] provide estimates of the magnitude of post-earnings announcement drift, and show that the drift varies inversely with firm size. In this and the following section, we replicate those results and demonstrate that they persist in more recent data. Unless otherwise specified, the results in this section are based on the procedures used by FOS, to maintain comparability; results based on the "continually balanced SUE strategy are reported only as supplemental information.

[14]Firm size was measured each year based on January 1 market values of equity. Small firms include those among the lowest four deciles of the NYSE/AMEX; large firms are those among the highest three deciles.

[15]If "bad news" firms are more likely to be small, due to price declines in anticipation of the earnings announcement (and vice versa for "good news" firms), then the Banz-Reinganum size effect would impart a *downward* bias in our estimated excess returns. That is, the bias would tend to offset any post-announcement drift.

Figure 3 presents CER plots for the NYSE/AMEX sample, after assigning firms to portfolios on the basis of standardized unexpected earnings. In contrast to the format in Figures 1 and 2, Figure 3 separates CER plots for the pre- and post-announcement periods, to make the post-announcement excess returns easier to guage. The results are similar to those obtained by FOS; there is a pronounced post-earnings announcement drift, increasing monotonically in unexpected earnings. A long position in the highest unexpected earnings decile and a short position in the lowest decile would have yielded an estimated excess return of approximately 4.2 percent (5.4 percent) over the 60 days (80 days) subsequent to the earnings announcement, or about 18 percent on an annualized basis. For the 1974-1981 period studied by FOS, we obtain an annualized return of 19 percent, less than the 25 percent implied by their results.[16]

Relation of drift to firm size. Figures 4 and 5 indicate how the drift varies by firm size, by presenting results for large and small firms.[17] As noted by FOS, the post-announcement drift is larger for smaller firms. Among small firms, a long position in the highest unexpected earnings decile and a short position in the lowest decile yields an excess return of approximately 5.3 percent over the 60 days subsequent to the earnings announcement. Comparable excess returns for medium-sized firms (not shown) and large firms are 4.5 percent and 2.8 percent, respectively.

Results based on the "continuously balanced SUE strategy" described in section 3.3.3 are similar. For 60-day holding periods, mean excess returns for small, medium, and large firms are 5.1 percent, 4.3 percent, and 2.8 percent.

---

[16]Differences between our results and those of FOS are most pronounced for small, "good news" firms. A possible explanation for the difference involves how control portfolios were constructed. It appears as if FOS included only NYSE firms in their control portfolios (see FOS, p. 585), whereas we included both NYSE and AMEX firms.

[17]"Small" firms are those for which the market value of common equity lies in the lower 40 percent of the distribution for the NYSE/AMEX. "Large" firms lie in the upper 30 percent of the distribution, and "medium-sized" firms lie in the middle 30 percent.

In regressions not reported here, we use the approach of FOS [see p. 595] to test the statistical significance of the post-announcement drift, and the effect of firm size. The results confirm that the magnitude of the drift is related to the magnitude of unexpected earnings, and that the absolute magnitude of the drift is related to firm size, at significance levels less than .01.

We do not present comparable plots for NASDAQ firms. However, the same phenomenon observed for NYSE/AMEX firms is found in that sample. The magnitude of the drift for NASDAQ firms lies between that observed for small and medium-sized firms on the NYSE/AMEX.

Longevity of the drift. Table 1 summarizes information about the longevity of the post-announcement drift. It presents the magnitude of the drift for stocks ranked in the lowest and highest quintile of the unexpected earnings, over periods extending two years beyond the earnings announcement date.

Most of the drift occurs during the first 60 trading days (about three months) subsequent to the earnings announcement, and there is little evidence of statistically significant drift beyond 180 trading days. If we assume all of the drift occurs within 480 days, then the fraction of the drift experienced within 60 days is 53 percent, 58 percent, and 76 percent for small, medium, and large firms, respectively. Approximately 100 percent of the drift occurs within nine months for small firms, and within six months for large firms. This result is consistent with the findings of Watts [1978], who finds a significant drift for six months in a sample of primarily large firms.

Table 1 suggests that, if the drift is explained by an incomplete adjustment for risk, the risk must exist only temporarily and must persist longer for small firms than large firms.

3.3.2. Shifts in betas as a potential explanation.

We now turn to a battery of tests designed to assess the plausibility of incomplete risk adjustment as an explanation for post-announcement drift. We first consider Ball, Kothari, and Watts' [1988] suggestion that beta shifts may account for post-announcement drift. They indicate that betas increase for firms with high unexpected earnings, and decrease for firms with low unexpected earnings.

Beta shifts are obviously a concern in a design that estimates betas in one period, and then uses those estimates in a different period. That was the case in much of the early research on post-announcement drift. However, that is not the case in the FOS design that we adopt. In this design, betas are never estimated; we *assume* that betas for our long and short positions are equal during the post-announcement period. Under this assumption, the combined long and short positions have zero systematic risk. Thus, while we will examine the Ball-Kothari-Watts hypothesis that betas *shift* around the time of earnings announcements, our ultimate concern is with differences in the *levels* of betas for high- and low-SUE firms in the post-announcement period.

Beta estimates are presented in Table 2, for firms in each SUE decile, and for a series of 60-day windows surrounding the earnings report date. There is a distinct indication of a beta shift in the direction suggested by Ball-Kothari-Watts, during the 60-day post-announcement window. In that window, there is a nearly perfect rank-order correlation between betas and SUE deciles. However, in contrast to Ball-Kothari-Watts, we find no evidence of beta shifts prior to the earnings announcement.

Although a post-announcement beta shift is evident in our data, it is much smaller than would be necessary to explain fully the magnitude of the drift. The difference in betas between the extreme high and low SUE deciles is .22 in the 60 day post-announcement period. Given an annualized return of 18 percent in the [1,60] window from a strategy of buying firms in the highest SUE decile and selling those in the lowest decile, and noting that the expected market portfolio return exceeded the risk free rate by 13 percent during this period, a difference in betas of approximately 1.4 (=.18/.13) would be required to explain the anomaly. That is more than six times larger than the difference we estimate (and 38 times larger than the post-announcement difference estimated by BKW).

Another difficulty with the beta-shift hypothesis is that the drift continues for 120 to 180 trading days (Table 1), but the beta shifts are almost completely reversed within 60 trading days after the earnings announcement. In the (61,120) window, for example, the betas for high SUE firms are only .03 higher than those for low SUE firms-- not nearly large enough to explain the magnitude of the drift in that window.

Finally, we should note that the beta in the post-announcement period for the "continuously balanced SUE strategy" is even smaller than the difference (.22) estimated in Table 2. Specifically, using a time series regression in quarterly data for 50 quarters, we obtain a beta estimate of .06, with an associated t-value of 1.79.[18]

We conclude that there is some merit to the Ball-Kothari-Watts claim that betas shift around earnings announcements, but that the magnitude and duration of the shifts fall far short of the amounts necessary to explain the magnitude and duration of the drift. If risk shifts are to explain post-earnings announcement drift fully, they must involve shifts in risk other than beta.

3.3.3. Other commonly-discussed asset pricing factors as potential explanations.

APT risk factors as potential explanations. In this section, we test for the possibility that trading strategies based on SUEs are risky on dimensions not captured by systematic risk. The risk factors we consider are those found in the literature on arbitrage pricing theory. Chen, Roll, and Ross [1986] found evidence that risks associated with industrial production, changes in default risk premia, and changes in term structure appeared to be

[18] Note that the "continuously balanced strategy" offsets positions in stocks in the extreme *quintiles* of the SUE distribution, rather than the extreme *deciles* that served as the basis for the previous discussion. Given that less extreme SUE portfolios experience smaller beta shifts, a smaller estimated beta for the "continuously balanced SUE strategy" is not surprising.

priced. They found weaker evidence that risks associated with unanticipated inflation and changes in expected inflation also affect asset pricing.

In Table 3, we regress the quarterly returns on our "continuously balanced SUE strategy" against quarterly measures of the five risk factors studied by Chen, Roll, and Ross.[19] In addition, we consider a regression that also includes the return on the NYSE index (net of the Treasury bill rate). Table 3 indicates whether a positive or negative correlation with a particular factor would indicate that the portfolio is "risky," as opposed to offering a "hedge" against risk. The evidence of Chen, Roll, and Ross suggests that assets that are *positively* correlated with unanticipated growth in industrial production (QP), and unanticipated changes in the default risk premium (UPR) are risky and have correspondingly higher required returns. On the other hand, assets that are *negatively* correlated with changes in expected inflation (DEI), unanticipated inflation (UI), and unanticipated changes in the term structure (UTS) are risky, and have higher required returns.

Table 3 provides no evidence that the returns on the SUE strategy are significantly correlated with any of the five factors studied by Chen, Roll, and Ross. Moreover, the five factors as a group do not explain a significant fraction of the variance in the strategy's return. Finally, note that for two of the five factors (QP and UTS), the sign of the coefficient would indicate that the SUE strategy provides a hedge against risk, rather than an increased exposure to risk.

---

[19]The variables were measured using the procedures of Chen, Roll, and Ross [1986] as they would be applied to quarterly data, with the following exceptions. First, our measure of inflation was the GNP deflator, rather than the Consumer Price Index; we used the GNP deflator because ASA-NBER forecasts of that series were available in machine-readable form. (Chen, Roll, and Ross used the Fama-Gibbons inflation forecasting model.) Second, our measure of the unanticipated default risk premium was the difference between the return on low-grade and high-grade corporate bonds, rather than the difference between low-grade corporates and government bonds. See Table 3 for further information.

We thank Gautum Kaul for supplying the data on default risk premia. Data on term structure were taken from Ibbotson and Associates [1988]. All other data were available from the Citibase tapes.

If the right-hand-side variables in Table 3 accurately measure ex post premia on all risk factors that are priced, then the intercept in the regression provides a test of market efficiency. Given that the dependent variable is the return on a zero-investment portfolio, the intercept should be zero under the efficient markets hypothesis. However, the estimated intercept indicates an excess return of 4 percent per quarter, with t-values of 7.55 and 8.00.

Dividend yield as a potential explanation. We also examine changes in dividend yields on "good news" and "bad news" portfolios. If dividend yields affect asset pricing, as predicted by the Brennan [1970] "after-tax" CAPM, then they could conceivably explain post-earnings announcement drift. What would be required is a sufficiently large increase in the difference between dividend yields on "good news" and "bad news" stocks. We do detect such a change, but the magnitude (4/10 of one percent of price) would imply a trivial impact on expected returns, given economically plausible dividend yield effects.

Our conclusion is that, if post-announcement drift represents a risk premium (or reflects another factor potentially affecting asset pricing, such as dividend yields), then the explanation must go beyond factors commonly discussed in the asset pricing literature. In the following sections, we examine whether some other, unidentified risk factor could plausibly explain the drift.

3.3.4. Consistent profitability of the strategy.

In this section, we examine how frequently a zero-investment SUE trading strategy generates a negative return. If a zero-investment strategy yields a positive *mean* return because it is risky, that risk must periodically surface in the form of a loss.

Figure 6 presents the 60-day excess returns on the EBM strategy, for each quarter from 1974:III through 1986:IV. Panel A presents returns to the FOS control portfolio strategy, where we assume a long (short) position in the firms whose unexpected earnings are ranked in the highest (lowest) decile. Panel B presents the returns to the "continually balanced SUE strategy."[20]

The interesting feature of Figure 6 is the consistency with which the zero-investment portfolios generate positive returns. The returns in Panel A are positive in 45 of 50 quarters, and in 13 of 13 years. In Panel B, the returns are positive in 44 of 50 quarters, and in 13 of 13 years. FOS present similar evidence; their Figure 2 (page 594) implies a positive excess return in 31 of 32 quarters.[21]

If indeed the post-announcement drift is explained by asset pricing model misspecification, the consistently positive returns are difficult to understand. If the returns on a zero-investment portfolio represent compensation for risk, then there must occur losses with an expected cost (in terms of utility) that is equal to the expected value of the risk premium. However, for the overall sample, returns of more than 200 percent (before compounding) have been generated over the 50 quarters, with negative returns in only five or six quarters-- and those negative returns sum to less than six percent.

To better appreciate how surprising the consistency is, consider the behavior of the *ex post* risk premium for a widely accepted measure of risk: beta. Fama and MacBeth [1973] present returns on zero-investment, unit-beta portfolios for the period 1935-1968. That portfolio generates a mean annualized return of about 10 percent, but among the 134 quarters represented there, returns on this portfolio were negative 39 percent of the time. The mean annualized return on the zero-investment portfolios described in Figure 6 is

[20] In Panel A, returns to a position are assigned entirely to the calendar quarter in which the position is assumed, even though that return generally spans portions of two calendar quarters. In Panel B, returns to positions that span two quarters are allocated properly to those calendar quarters. A future draft will present Panel A on a basis directly comparable to Panel B.

[21] However, FOS do not discuss the implications of this result for distinguishing among alternative explanations for the drift.

higher (18 percent), and yet the portfolios experience a loss only 10 or 12 percent of the time.

Some readers of prior drafts have questioned whether the consistent profitability depicted in Figure 6 could reflect some problem in the benchmark we use to measure excess returns. Note that, to the extent that our benchmark fails to control for some risk that is priced in the market, the results are that much more surprising.[22] For example, if we have failed to control for systematic risk and our combined long and short position has a positive beta, one would expect the excess return on the strategy to be negative when the overall market declines. However, over the 50 quarter horizon, the equally-weighted NYSE index declined 16 times, and yet the excess return on the strategy in Panel A was positive in 15 of those 16 quarters (11 of 16 quarters for Panel B).

We are able to reconcile this evidence with CAPM misspecification (i.e., failure to control fully for risk) only under the following conditions:

(1) the infrequency of losses in the 1974-1986 period is extremely unusual, relative to what would be observed in a longer period, or

(2) the risk premium earned on the EBM strategy represents compensation for the risk of infrequent but catastrophic losses, none of which was observed within this 13-year time span, or

(3) the disutility of the losses we observe is commensurate with the utility of the gains, because the losses occur during periods when a $1 decline in wealth is 35 times more important than the average $1 increase in wealth. (Cumulative gains are 35 times larger than cumulative losses.)

We find (1) and (3) implausible; those who choose to defend (2) must do so without the aid of any evidence.

---

[22]The results could conceivably be explained by a failure to control for some factor that increases expected returns consistently through time. However, to our knowledge, the only asset pricing models developed to date that include such a factor are the Brennan [1970] "after-tax" CAPM (which includes a dividend yield effect), and the Amihud-Mendelson [1986] CAPM, which includes a term linked to the bid-ask spread. Section 3.3 dismissed Brennan's dividend yield effect as an explanation. The Amihud-Mendelson CAPM could explain the result only if announcement of "good news" (bad news) caused a long-run increase (decrease) in the proportional bid-ask spread. But this is the *opposite* of what one would expect, given that the proportional bid-ask spread varies inversely with price, and that "good news" ("bad news") firms tend to experience price increases (decreases).

3.3.5. Raw returns on "bad news" firms.

The large estimated negative *excess* post-announcement returns for firms with extreme negative unexpected earnings suggests that the *total (raw)* post-announcement returns for those firms could be less than the risk-free rate, or even negative. Such predictably low raw returns on risky assets can be reconciled with most modern capital asset pricing models, but only under special conditions that many would find implausible as applied to a broad cross-section of stocks. Essentially, the stocks would have to offer some hedge, the value of which more than offsets the cost of any other risk to which the asset is exposed.

Table 4 summarizes the total returns, compounded over various periods, for "bad news" stocks: those ranked in the lowest decile of the unexpected earnings distribution. The total annualized returns on the "bad news" stocks (averaged over firms of all sizes) are 1.5 percent, 12.6 percent, and 10.4 percent for periods ended 5, 20, and 40 trading days subsequent to the earnings announcement. Total annualized returns for "good news" firms are 36.5 percent, 29.3 percent, and 26.7 percent over the same periods. These returns were generated during 1974-1986, when the average annualized return on 3-month Treasury bills was 8.8 percent, and the return on the equally-weighted NYSE index was approximately 22 percent (13 percent for the value-weighted index).

The total annualized returns for *small* "bad news" stocks over the 5 days after the earnings announcement are not only less than the average Treasury bill rate, but are actually negative. The total returns for *medium* firms over the same 5-day window are zero, and over the 40 days subsequent to the announcement are less than the T-bill rate. All other total returns are in excess of the T-bill rate. However, for two months following the announcement, the difference is small; for the overall sample, the 40-trading-day return is only 10.4 percent, or 1.6 percent higher than the average T-bill rate.[23] This contrasts with a 26.8 percent total return for the "good news" firms.

In order to reconcile this evidence with CAPM misspecification, one must believe either that (1) betas on the "bad news" stocks are near zero (and negative for small and medium stocks shortly after the announcement), or (2) the value of these stocks as hedges against some unidentified risk causes their cash flows to be discounted at rates less than Treasury bill rates during the 5-day post-announcement period, and at rates nearly that low for two months thereafter. Condition (1) is inconsistent with evidence in Table 2, and condition (2) is, for us, difficult to believe.[24]

3.3.6. Tests of a hypothesis about post-announcement drift and costs that may impede trading.

Much of the above evidence is difficult to explain as the result of an incomplete adjustment for risk. Therefore, we turn to the possibility that the drift might occur because some costs create sufficient impediments to trading to prevent a complete and immediate response to earnings announcements. The tests in this section were inspired by Ball [1978, p. 110], who argued that:

---

[23] A comparison of raw returns to the average Treasury bill rate is imprecise, in that it assumes the event periods are evenly distributed in calendar time. We also calculated the difference between raw returns and contemporaneous returns on Treasury bills that were within one week of maturity (assuming that the T-bills are rolled over weekly). For the overall sample, the difference was negative for the first 5 days of the post-announcement period (-7.0 percent), and positive for the first 40 days (2.1 percent). The average return on T-bills one week from maturity was 8.5 percent during 1974-1981.

[24] Some readers of prior drafts have questioned whether the low raw returns to the "bad news" firms can be explained by a deviation of observable *ex post* realizations from the unobservable expected returns modeled in the CAPM. First, note that the expected returns at issue are those existing *after* the announcement of the bad news. Thus, if those expected returns do *not* reflect the bad news, but the *ex post* realized returns *do*, then the expectations immediately after the announcement must have ignored available information. That is, the argument is equivalent to labeling the drift a delayed price response. Second, imprecision in estimates is unlikely to cause much difference between expected returns and *ex post* realizations. The standard errors of the annualized raw returns over the intervals [1,5], [1,20], and [1,40] are all less than 1 percent.

> "...if the "slow" market reaction is explained in terms of transactions costs (or costs of "professionals" operating in the market), then *small* deviations from expectations are those which imply market disequilibrium. Large deviations presumably attract more investors and are promptly incorporated in prices because (under this hypothesis) the net gain, after costs, is higher. The consistent interpretation of this hypothesis is that the excess returns persist up to, but not beyond, the level of marginal transactions and information processing costs."

Under Ball's depiction, a post-announcement drift would be observed only when unexpected earnings is small. Alternatively, the drift may be observed for all levels of unexpected earnings, but would never exceed a threshold (equal to the cost of exploiting the information), regardless of how large is unexpected earnings.[25] That is, regardless of whether the *total* stock price response implied by an earnings announcement is 2 percent, 5 percent, or 20 percent, the price might move immediately to within (say) 2 percent of the implied level. At that point, incentives to exploit the earnings information would be eliminated for many traders, and the remainder of the response would occur only with some delay. In such a market, the post-announcement drift would increase as unexpected earnings increases, but only to some upper bound; beyond that bound, the drift would remain constant, regardless of how large is unexpected earnings.

Ball [1978, p. 110] notes that existing evidence does not appear consistent with this characterization: "...the evidence...is that extreme-rank earnings and dividend changes are associated with *larger* estimated excess returns, contrary to the "transactions cost" and "private cost" explanations." However, we consider here whether we (and prior researchers) have failed to observe this upper bound because we have not yet examined sufficiently extreme values of unexpected earnings.[26] Our approach is to divide our sample into progressively smaller portfolios, based on rankings of unexpected earnings. That is, we first divide the sample into halves, then thirds, quintiles, deciles, and so on, until

---

[25]Although we initially inferred that Ball's depiction was consistent with the second alternative, he has indicated to us that he intended to imply the first.

[26]We are extremely grateful to Jim Noel, who suggested the tests in this section.

finally we divide the sample into 100 portfolios, based on rankings of unexpected earnings. At each of these steps, we calculate the excess return from a long position in the portfolio with the highest unexpected earnings, and a short position in the portfolio with the lowest unexpected earnings. Thus, at each step, the value of unexpected earnings in our portfolios becomes more extreme. If post-announcement drift is caused by a cost that impedes trading, we should observe that, at a point bounded by that cost, the drift should increase no more, even though unexpected earnings continues to increase.

The results are presented in Figure 7, Panel A. We find that the drift (over 60 days) grows larger, up to the point where the sample is split into deciles. Beyond that point, no matter how extreme is unexpected earnings, the drift does not increase. Note that the upper bound for the drift is about 4 percent, or 2 percent per position. That amount is within the bounds of transactions costs for the average firm in the post-1975 era of negogiated commissions, as estimated by Stoll and Whaley [1983], where such costs include both commissions and the bid-ask spread. Figure 7 Panel B shows that the drift is bounded at approximately 5 percent, 4.3 percent, and 3 percent, for small, medium, and large firms, respectively. This is consistent with Stoll and Whaley's [1983] evidence that transactions costs vary inversely with firm size; when their sample is segregated into thirds, transactions costs are 3.9 percent, 2.6 percent, and 2.0 percent for small, medium, and large firms. When these amounts are doubled to account for a combined long and short position, they exceed the bounds implied by Figure 7, Panel B.

One potential alternative explanation for the result is that the more extreme values of unexpected earnings simply reflect estimation error. That is, beyond some upper bound, any additional increases in our measures of unexpected earnings represent nothing more than noise. However, the data indicate that this is not the case. Figure 7 also presents the *pre*-announcement excess returns for portfolios with varying levels of unexpected earnings. Note that even though the *post*-announcement drift reaches a maximum when the sample is split into deciles, the *pre*-announcement drift continues to increase to the point where the sample is split into 40 portfolios. Thus, increases in unexpected earnings (at least to that point) have stock price impacts, and are not purely the result of noise.

The evidence in Figure 7 is (for us) difficult to reconcile with arguments based on CAPM misspecification. To accommodate this result, such an argument would have to introduce a "kink" in the relation between unexpected earnings and risk. That is, unexpected earnings would have to proxy for an omitted risk factor up to some point, but then additional increases in unexpected earnings could no longer be correlated with increases in risk.

On the other hand, the trading-cost-based explanation for the results in Figure 7 also raises some difficult questions:

Why would trading costs necessarily cause underreaction to new information, as opposed to simply introducing noise in prices? A possible answer to this question rests on the assumption that there are some uninformed traders who stand willing to buy when prices fall, or sell when prices rise. Then prices will tend to be "sticky," and can be driven to the "appropriate" level only by speculators who trade on the new earnings information. If costs prevent those speculators from fully exploiting the new information, prices will tend to underreact.

Why don't specialists or other market makers move the price to the "appropriate" level upon the first trade after the earnings announcement? A possible answer is that "noise" in trading patterns, requirements to maintain an orderly market, and limits on a market makers' own capital may make it difficult for market makers to know what the full response to earnings should be, and to move the bid-ask spread to the "appropriate" level single-handedly. Note also that, so long as the price adjusts slowly and the market maker always buys at the bid and sells at the ask, he/she profits throughout the adjustment process. Therefore incentives to bear the risk of driving prices to a perceived "appropriate" level may be limited.

Why can't the drift be eliminated by traders who face no commissions and can bypass the specialist's bid-ask spread, thus facing trivial transactions costs? A partial answer to this question may be that the drift is most clearly evident for small firms, and that for such stocks, the bid-ask spread is difficult to avoid for those who are not market makers. In addition, investment houses may be unable to take large positions in small stocks without exerting price pressure. Finally, unless an investment house already owns a "bad news" stock, the costs of short-selling may prevent exploitation of negative unexpected earnings.

3.3.7 Comparison of returns to long and short positions.

If the costs of trading do play some role in explaining post-announcement drift, then one might expect the excess returns to short positions in "bad news" firms to exceed those for long positions in "good news" firms. The reason is that restrictions on short sales would make it prohibitively expensive for some traders to act on "bad news."

Evidence in Table 1 is consistent with this hypothesis. The estimated excess returns to short positions in "bad news" firms are larger, and last longer, than the estimated excess returns on "good news" firms. For small, medium, and large "bad news" firms, the estimated CERs over the [1,180] window are 7.3 percent, 6.4 percent, and 2.7 percent, respectively. Comparable amounts for "good news" firms are 2.6 percent, 2.4 percent, and 2.2 percent. For the combined sample, the CERs for the "bad news" firms are more than twice as large as those for the "good news" firms.[27]

---

[27]We are in the process of recalculating these excess returns using a method that simultaneously eliminates the problems discussed by Blume and Stambaugh [1983], and controls for the Banz-Reinganum size effect separately for the "good news" and "bad news" firms. The calculations in Table 1 satisfy the latter condition, but not the former. Since the Blume-Stambaugh effect should cause upward bias in cumulated returns for both "treatment" and "control" portfolios, it would not necessarily cause any bias in the estimated excess return itself. Nevertheless, the Blume-Stambaugh effect is likely to introduce noise, and thus elimination of this problem should lead to more precise comparisons.

4. DISCUSSION AND CONCLUSIONS

Much of the evidence presented here casts doubt on CAPM misspecification as an explanation for post-earnings announcement drift. In section 4.1, we summarize implications of the evidence for various forms of misspecification. Section 4.2 then reviews the plausibility of alternative explanations.

4.1 IMPLICATIONS OF THE EVIDENCE FOR CAPM MISSPECIFICATION

CAPM misspecification can assume several different forms. We divide them into (1) risk mismeasurement and (2) other misspecifications. In turn, risk mismeasurement can include (a) misestimation of systematic risk and (b) exclusion of risk factors other than systematic risk.

4.1.1.a. Risk mismeasurement: misestimation of beta.

Several features of the evidence cast doubt on the Ball, Kothari, and Watts [1988] contention that beta shifts explain post-earnings announcement drift. The key features are:

1) Estimated beta shifts are not nearly large enough and do not last long enough to explain the magnitude and longevity of the drift (section 3.3.2).

2) The Ball-Kothari-Watts hypothesis predicts that the EBM strategy has a positive beta, and thus should perform poorly in bear markets. However, the EBM strategy yields consistently positive returns in both bull and bear markets (section 3.3.4).

3) The Ball-Kothari-Watts hypothesis offers no explanation for why the drift should be larger and last longer for small firms than large firms, or why they should be larger and last longer for "bad news" firms than "good news" firms.

4.1.1.b. <u>Risk mismeasurement: exclusion of risk factors other than systematic risk.</u>

Key features of the data that cast doubt on the relevance of this potential explanation for post-earnings-announcement drift are:

1) We find no evidence that an SUE trading strategy is risky along any of the five dimensions identified by Chen, Roll, and Ross [1986] as important factors in asset pricing (section 3.3.3).

2) If the EBM strategy is risky on some unidentified dimension, then there is little evidence of that risk surfacing in the form of losses whose cost (in terms of utility) could plausibly be commensurate with the value of the supposed risk premium (section 3.3.4). The consistent profitability of the EBM strategy raises the question, "Where's the risk?"

3) Subsequent to earnings announcements, "bad news" firms yield mean *total* returns that are less than T-bill yields during the first week, and only slightly greater than T-bill yields during the first two months (section 3.3.5). Unless the riskiness of "bad news" stocks declines so much (or their value as hedges increases so much) that their cash flows deserve to be discounted at something close to the riskless rate, this feature of the data cannot be reconciled with an argument based on CAPM misspecification.

4) The drift is initially increasing in unexpected earnings, but appears to reach an upper bound, beyond which the drift remains constant as unexpected earnings rises (section 3.3.6). In order to reconcile this result with CAPM misspecification, one would have to believe that unexpected earnings proxy for an unidentified risk factor only to some point; beyond that point, further increases in unexpected earnings would have to be uncorrelated with the unidentified risk.

4.1.2. Other forms of CAPM misspecification.

CAPM misspecification could also involve a failure to allow for market imperfections such as taxes. If the difference between ordinary and capital gains tax rates affects pricing, then a "dividend yield effect" would exist in stock returns. Such an effect could be consistent through time, potentially explaining the consistently positive excess returns on the EBM strategy. However, as indicated in section 3.3.3, differences in dividend yields between the high and low unexpected earnings firms are not nearly large enough to explain the magnitude of the drift.

4.2 DELAYED PRICE RESPONSE AS AN EXPLANATION

If arguments based on CAPM misspecification cannot plausibly be reconciled with the data, we are left with alternative explanations that label the drift as a delayed price response. We are reluctant to entertain delays in price response resulting entirely from the ignoring of information, because earnings information is so widely disseminated and freely available. Moreover, if this were the explanation, we would expect the drift to be always increasing in unexpected earnings, rather than bounded, as indicated in section 3.3.6.

The remaining alternative is a delay in price response that occurs because some costs impede trading. The results in section 3.3.6 and 3.3.7 are consistent with this possibility. Moreover, the magnitude of the drift is such that an SUE strategy would not be profitable for other than large investors with low transactions costs. Given that the SUE strategies studied here would also require short-selling and some effort, they may not be profitable as stand-alone strategies, even for large investors. If so, trading on post-announcement drift may be profitable only as a marginal investment strategy. That is, traders who choose to transact for reasons unrelated to earnings announcements might buy stocks that have recently reported strong earnings, or sell stocks already in their portfolio that have recently experienced weak earnings. It is not clear that this type of marginal strategy would quickly drive prices to the "right" level, and the data are fully consistent with the proposition that it does not. Moreover, this may be the only explanation that is simultaneously consistent with (1) the rational use of "recent earnings surprise" as a buy/sell signal among several institutions and investment houses, and (2) the persistence of the drift, despite this activity.

If indeed trading costs (including direct transactions costs and other costs of implementation) do explain post-earnings-announcement drift, then we should observe drifts for other information events as well. It is interesting to note that drifts are observed after a variety of events, including, for example, 13-D filings to announce the acquisition of at least 5 percent of a firm's stock (Larcker and Lys [1987]), repurchase tender offers (Lakonishok and Vermaelen [1988]), dividend announcements (Charest [1978], bond rating downgrades (Holthausen and Leftwich [1986], and earnings forecast revisions by managers (McNichols [1989]) and analysts (Brown, Foster, and Noreen [1985]).

4.3 CONCLUSIONS

This paper has re-examined prior research on post-earnings-announcement drift, and has offered some new evidence on the issue. We have attempted to discriminate between two alternative explanations for the drift: a delay in the response to earnings reports, and a failure to adjust excess returns fully for risk.

We contend that much of the evidence cannot plausibly be reconciled with arguments built on risk mismeasurement. In contrast, all features of the data are consistent with expectations, if the drift represents a delay in the price response caused by costs that impede trading activity.

If the latter explanation is correct, the implications would be as follows. First, it may be *not* be true that, so long as *any* trading occurs, market makers insure that prices reflect all available information. Second, prices may be "noisy" indicators, rather than precise measures, of values that would exist in a frictionless market. Third, returns over short periods may be "noisy," and perhaps downward-biased indicators of economic consequences to a firm, especially when the consequences are small, relative to firm value. This final implication may be quite important for accounting studies, since so many of them seek to confirm subtle information effects over short return intervals.

Table 1

Longevity of post-earnings announcement drift

Excess returns on High and Low SUE (standardized unexpected earnings) portfolios
(those with unexpected earnings ranked in the highest and lowest deciles)

	Percentage excess returns[1]								
Holding period	small firms			medium firms			large firms		
(trading days, relative to announcement)	high SUE	low SUE	diff. (hi-lo)	high SUE	low SUE	diff. (hi-lo)	high SUE	low SUE	diff. (hi-lo)
-59-0	6.42*	-8.27*	14.70*	4.15*	-6.81*	10.96*	2.71*	-4.13*	6.84*
1-60	2.19*	-3.13*	5.32*	1.93*	-2.58*	4.51*	1.45*	-1.29*	2.74*
61-120	0.38	-2.24*	2.62*	0.33	-2.22*	2.55*	0.51**	-0.76*	1.27*
121-180	0.03	-1.93*	1.95*	0.18	-1.61*	1.79*	-0.20	-0.66*	0.45
181-240	0.20	-0.38	0.58	-0.40	-0.58***	0.18	-0.45**	-0.44***	-0.01
241-300	-1.22*	0.56	-1.77*	-0.88*	0.16	-1.04**	-0.62*	0.23	-0.85*
301-360	-0.54	-0.96*	0.42	-0.36	-0.71**	0.35	-0.64*	-0.16	-0.48
361-420	-0.27	-0.33	0.06	-0.18	-0.37	0.18	-0.55*	-0.71*	0.17
421-480	0.29	-0.51	0.80	-0.43	0.37	-0.80***	-0.42**	-0.73*	0.31
Cumulative post-announcement drift									
1-60			5.32*			4.51*			2.74*
1-120			7.95*			7.06*			4.02*
1-180			9.90*			8.85*			4.47*
1-480			9.99*			7.72*			3.61*
Cumulative drift (as a fraction of 480 day drift)									
1-60			0.53			0.58			0.76
1-120			0.80			0.91			1.11
1-180			0.99			1.15			1.24
1-480			1.00			1.00			1.00

---

1 Small, medium and large firms are size deciles 1 to 4, 5 to 7, and 8 to 10, respectively, based on January 1 market equity values for all NYSE and ASE firms.

Significance levels for two-tailed tests of the hypotheses that excess returns equal zero are coded as follows:

- * significant at the 1% level
- ** significant at the 5% level
- *** significant at the 10% level

Table 2

Beta Estimates by SUE Category, in Periods Surrounding Earnings Announcement

Beta estimates are generated as follows. For each 60-day window, we calculate compounded returns for individual stocks and the equally-weighted CRSP index. This constitutes a single observation in a market model regression estimated for each SUE category. There are approximately 8500 (overlapping and thus non-independent) observations underlying estimates for the (-59,0) and (1,60) windows, are slightly fewer for other windows. The standard error for each estimate in the table is approximately 0.02. Due to lack of independence in the data, the estimated standard errors are likely to be downward biased (Bernard [1987]).

SUE decile (1=low; 10=high)	Pre-announcement period (-119,60)	(-59,0)	Post-announcement period (1,60)	(61,120)	(121,180)	(181,240)
1	.93	1.03	.90	.98	.97	.97
2	.90	.94	.90	.89	.93	.99
3	.95	.94	.92	.90	.92	.94
4	1.00	.95	.97	.94	.95	.91
5	.98	1.04	1.00	1.05	.97	.97
6	1.05	1.03	1.03	1.01	.99	.94
7	1.05	.97	1.01	1.03	.95	.96
8	1.02	1.03	1.06	.99	.96	.93
9	.98	1.00	1.07	.97	.90	.87
10	1.04	1.02	1.12	1.01	.99	.89
Difference: 10 minus 1	.11	-.01	.22	.03	.02	-.08

Table 3

Sensitivity of drift to risk factors used in studies of arbitrage pricing theory

Returns to a zero-investment portfolio, with long positions in "good news" stocks and short positions in "bad news" stocks, are calculated for each calendar quarter. Such portfolios are created within groups of small, medium, and large firms, and returns for the three size groups are averaged. The returns are then regressed against the return on the equally-weighted NYSE index (in risk premium form), and five factors identified by Chen, Roll, and Ross [1986] as potentially influencing asset prices.

	Independent variables[1]:								F-Test[2] of significance of variables other than (Rmt-Rft)
	Intercept	Rmt-Rft	QP	DEI	UI	UPR	UTS	R-Square	
Sign of coefficient, if risky (as opposed to a hedge)		+	+	-	-	+	-		
Coefficient	.04	.03	-.12	-1.86	-.40	.10	.07	.11	.26
T-value	(7.55)	(.57)	(-.52)	(-1.09)	(-.36)	(.38)	(.82)		
Coefficient	.04	---	-.07	-2.11	-.38	.19	.10	.10	.99
T-value	(8.00)		(-.34)	(-1.29)	(-.34)	(.95)	(1.50)		

[1] Independent variables are defined as follows:

Rmt-Rft = return on equally-weighted NYSE index, less 90-day Treasury bill rate;
QP = quarterly growth rate in industrial production, lagged ahead one period;
DEI = change in expected inflation;
UI = unanticipated inflation;
UPR = unanticipated change in the default risk premium (return on high yield bonds (under BBB), less return on AAA bonds)
UTS = unanticipated change in the term structure; (return on long-term government bonds, less Treasury bill rate).

[2] F(1,43) is significant at the .05 level for values in excess of 4.08.

Table 4

Total (Raw) Returns on "Bad News" (Lowest SUE decile) Portfolios

Holding period (trading days, relative to announcement)	Small firms Raw Return	Small firms Cum. Raw Return	Medium firms Raw Return	Medium firms Cum. Raw Return	Large firms Raw Return	Large firms Cum. Raw Return
Pre-announcement period:						
(-79,0)	-1.8%*	-1.8%*	-.4%	-.4%	2.6%*	2.6%*
Post-announcement period:						
(1,5)	-.14	-.14	.00	.00	.23	.23
(6,20)	.89*	.75*	.85*	.85*	1.19*	1.42*
(21,40)	1.31*	2.05*	.46	1.31*	.24	1.66*
(41,60)	2.36*	4.42*	1.87*	3.18*	1.85*	3.51*
(61,80)	1.32*	5.74*	.78*	3.95*	1.07*	4.59*

Annualized post-announcement raw return:

	Small firms	Medium firms	Large firms
(1,5)	-7.0	0.0	11.5
(1,20)	9.4*	10.6*	17.8*
(1,40)	12.8*	8.2*	10.2*
(1,60)	18.4*	13.2*	14.6*
(1,80)	17.9*	12.3*	14.3*

Comparable annualized raw returns for "good news" (highest decile SUE) portfolio:

	Small firms	Medium firms	Large firms
(1,5)	32.5%*	41.6%*	35.5%*
(1,20)	26.6*	33.7*	27.5*
(1,40)	29.7*	28.0*	22.3*
(1,60)	32.9*	27.8*	21.4*
(1,80)	30.5*	26.9*	20.8*

Mean annualized returns across all firm size categories

	Low SUE	High SUE
(1,5)	1.5%	36.5%*
(1,20)	12.6*	29.3*
(1,40)	10.4*	26.7*
(1,60)	15.4*	27.4*
(1,80)	14.8*	26.1*

*Significantly different from zero, .05 level (two-tailed test).

Figure 1

Results of FOS Earnings-based-model (EBM) tests

Firms are ranked according to unexpected earnings, defined as actual earnings minus statistical forecast. Portfolio 10 includes firms with the highest values of unexpected earnings. (From FOS, page 589.)

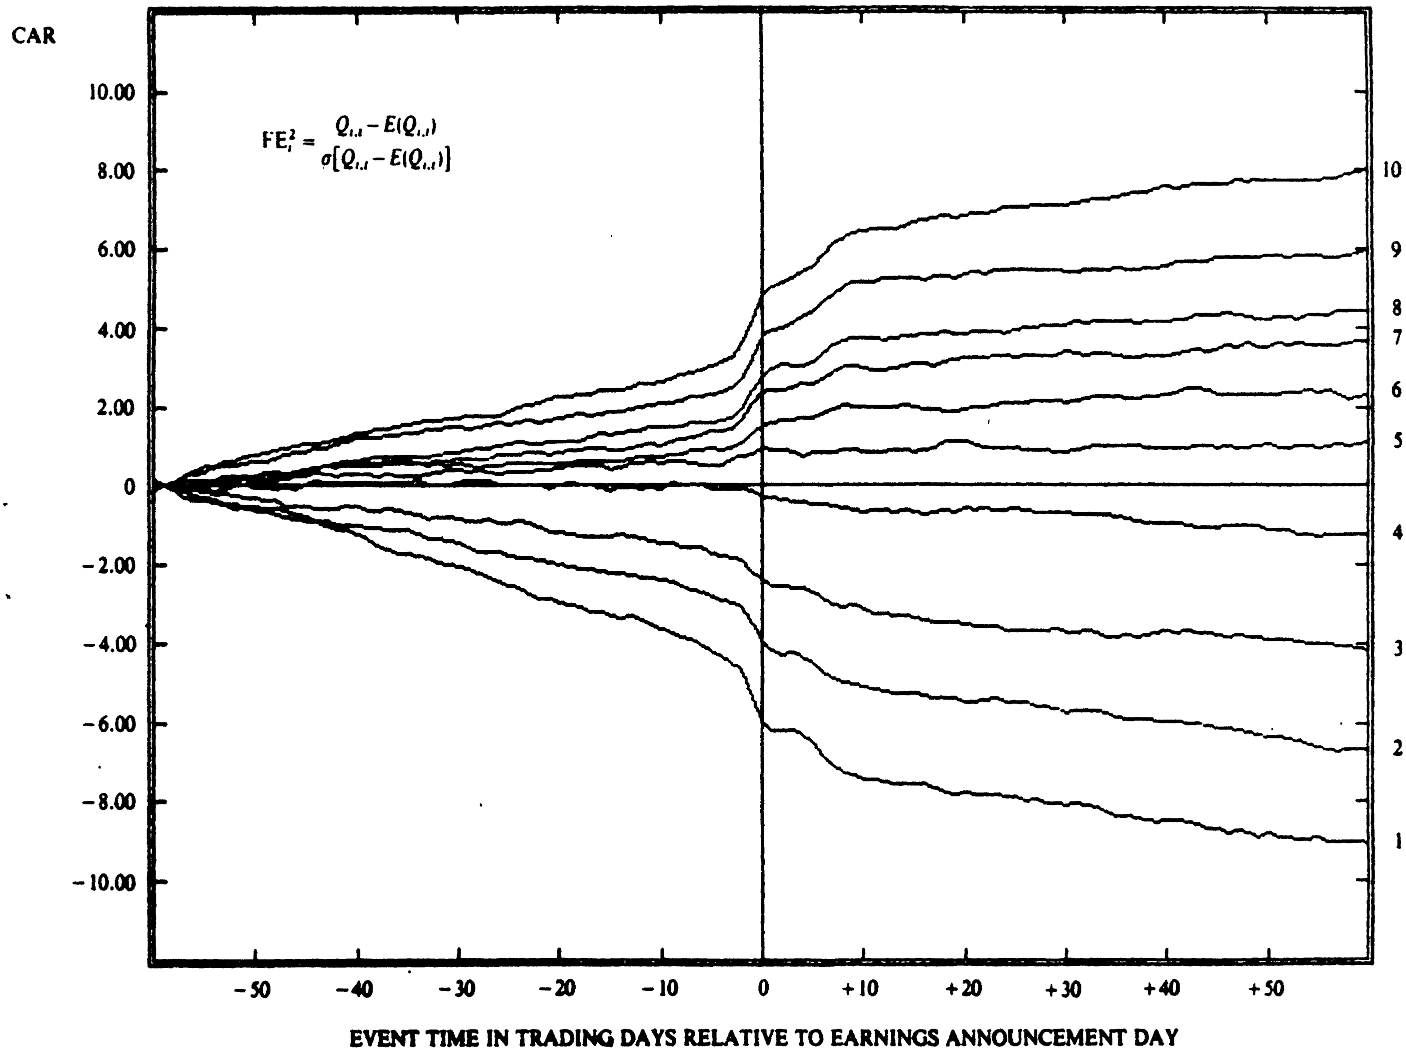

Figure 2

Results of FOS Security-return-model (SRM) tests based on excess returns in (-60,0) interval

Firms are ranked according to excess returns during days −60 through 0, where day 0 is the earnings announcement day. Portfolio 10 includes firms with highest values of excess returns during that period. (Duplicated from FOS, page 591.)

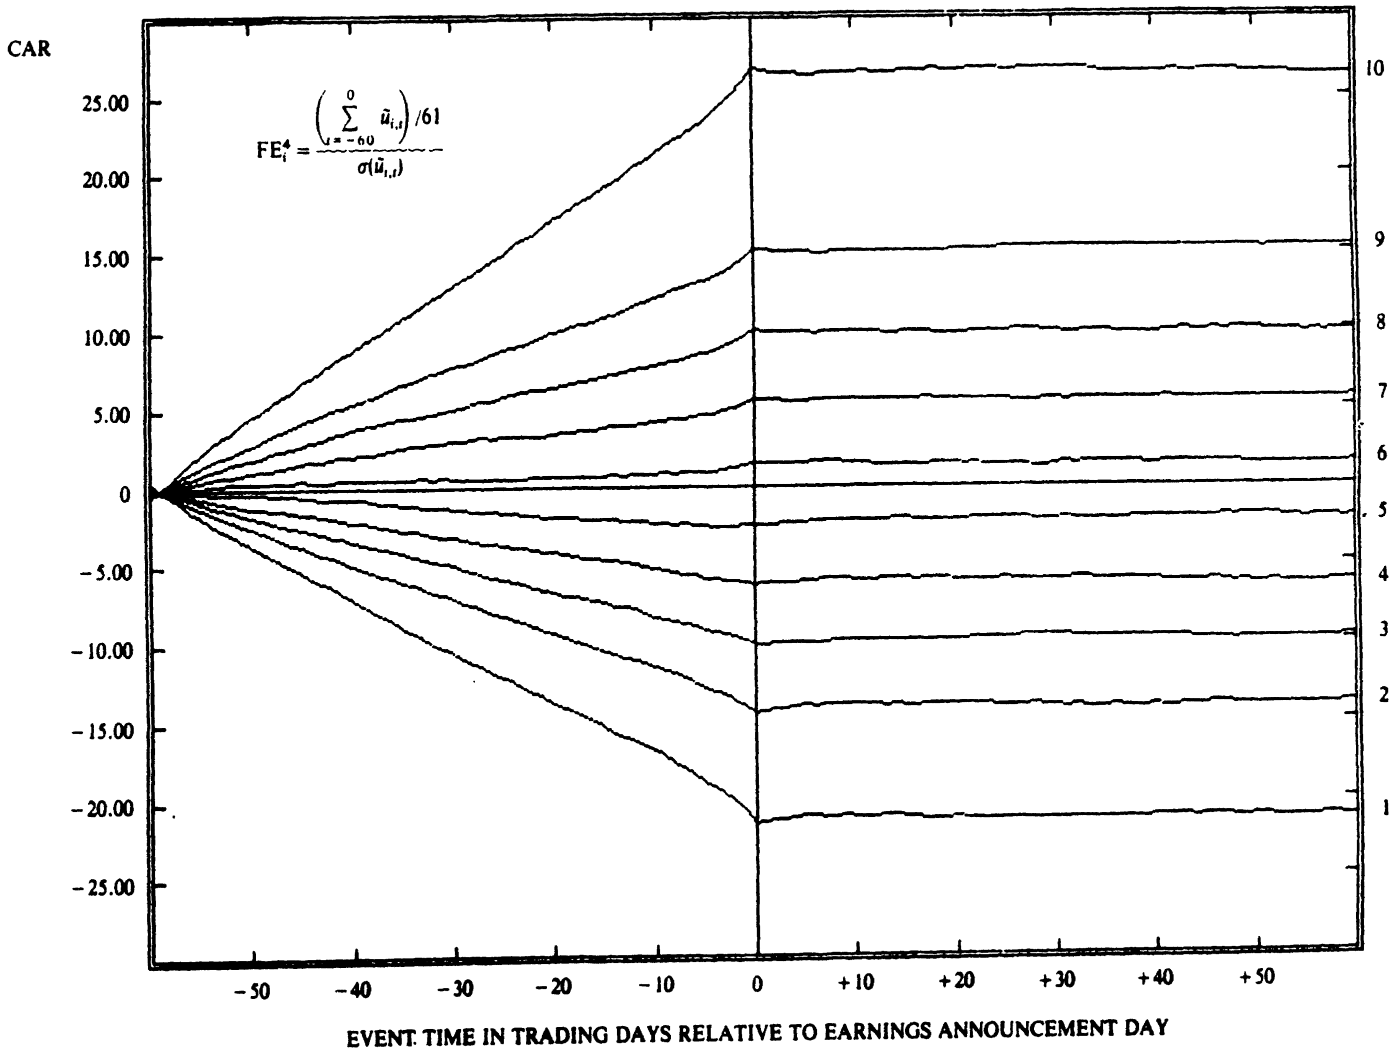

Figure 3

Cumulative excess returns (CER) for SUE portfolios: all announcements
Announcements ranked by deciles of standardized unexpected earnings (SUE)[1]

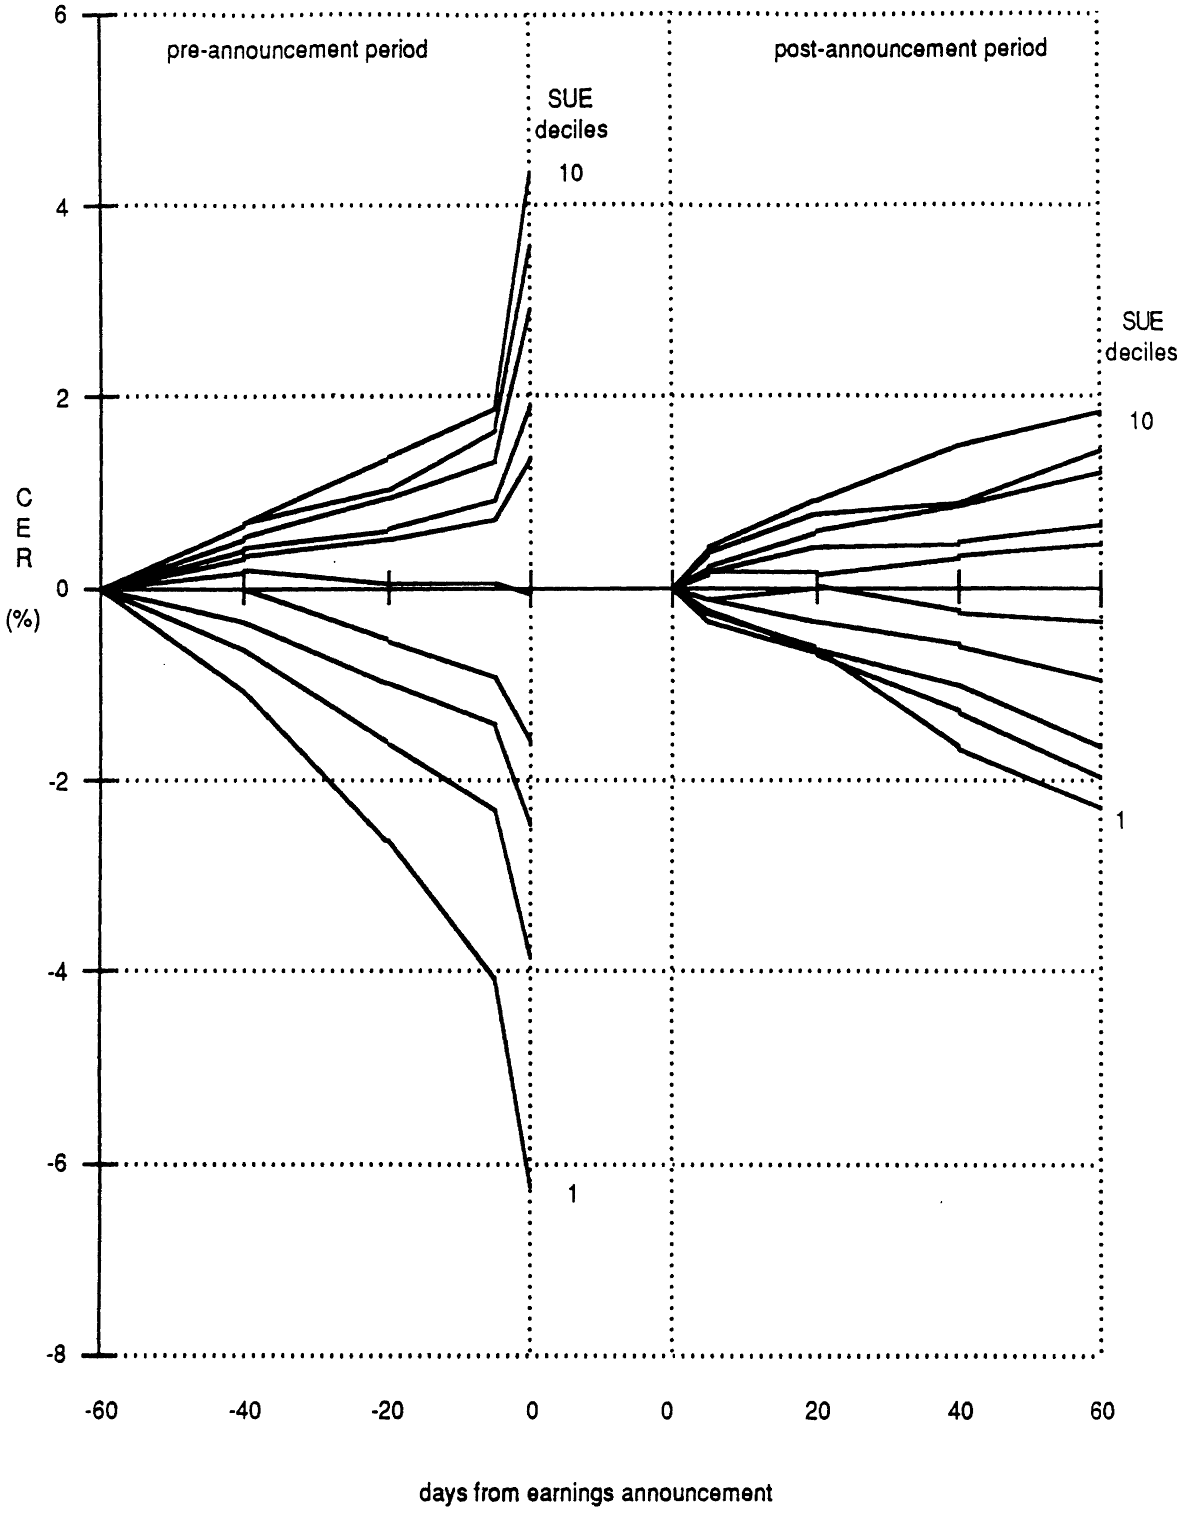

1 See Foster, Olsen and Shevlin (1984) for procedures used to compute cumulated excess returns (based on comparisons with benchmark size decile portfolios) and SUE (based on their Model 2).

Figure 4.

Cumulative excess returns (CER) for SUE portfolios: large firms only[1]
Announcements ranked by deciles of standardized unexpected earnings (SUE) [2]

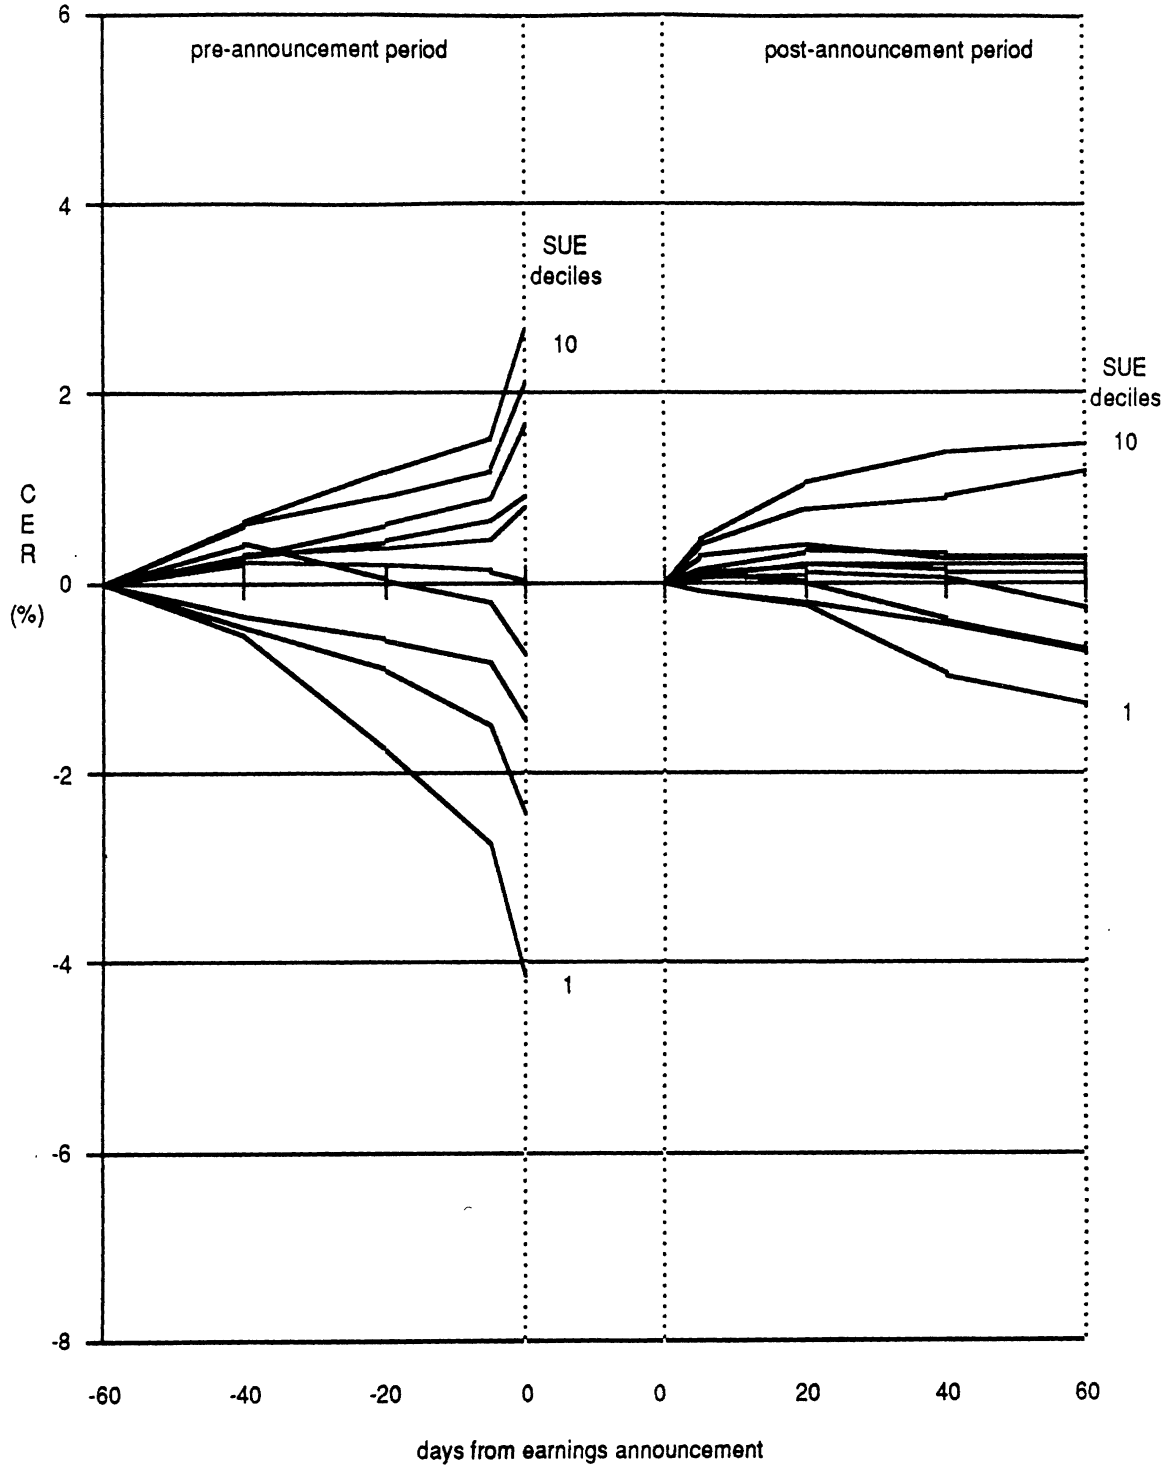

1 Large firms are size deciles 8 to 10, based on January 1 market equity values for all NYSE and ASE firms.

2 See Foster, Olsen and Shevlin (1984) for procedures used to compute cumulated excess returns (based on comparisons with benchmark size decile portfolios) and SUE (based on their Model 2).

Figure 5

Cumulative excess returns (CER) for SUE portfolios: small firms only[1]
Announcements ranked by deciles of standardized unexpected earnings (SUE) [2]

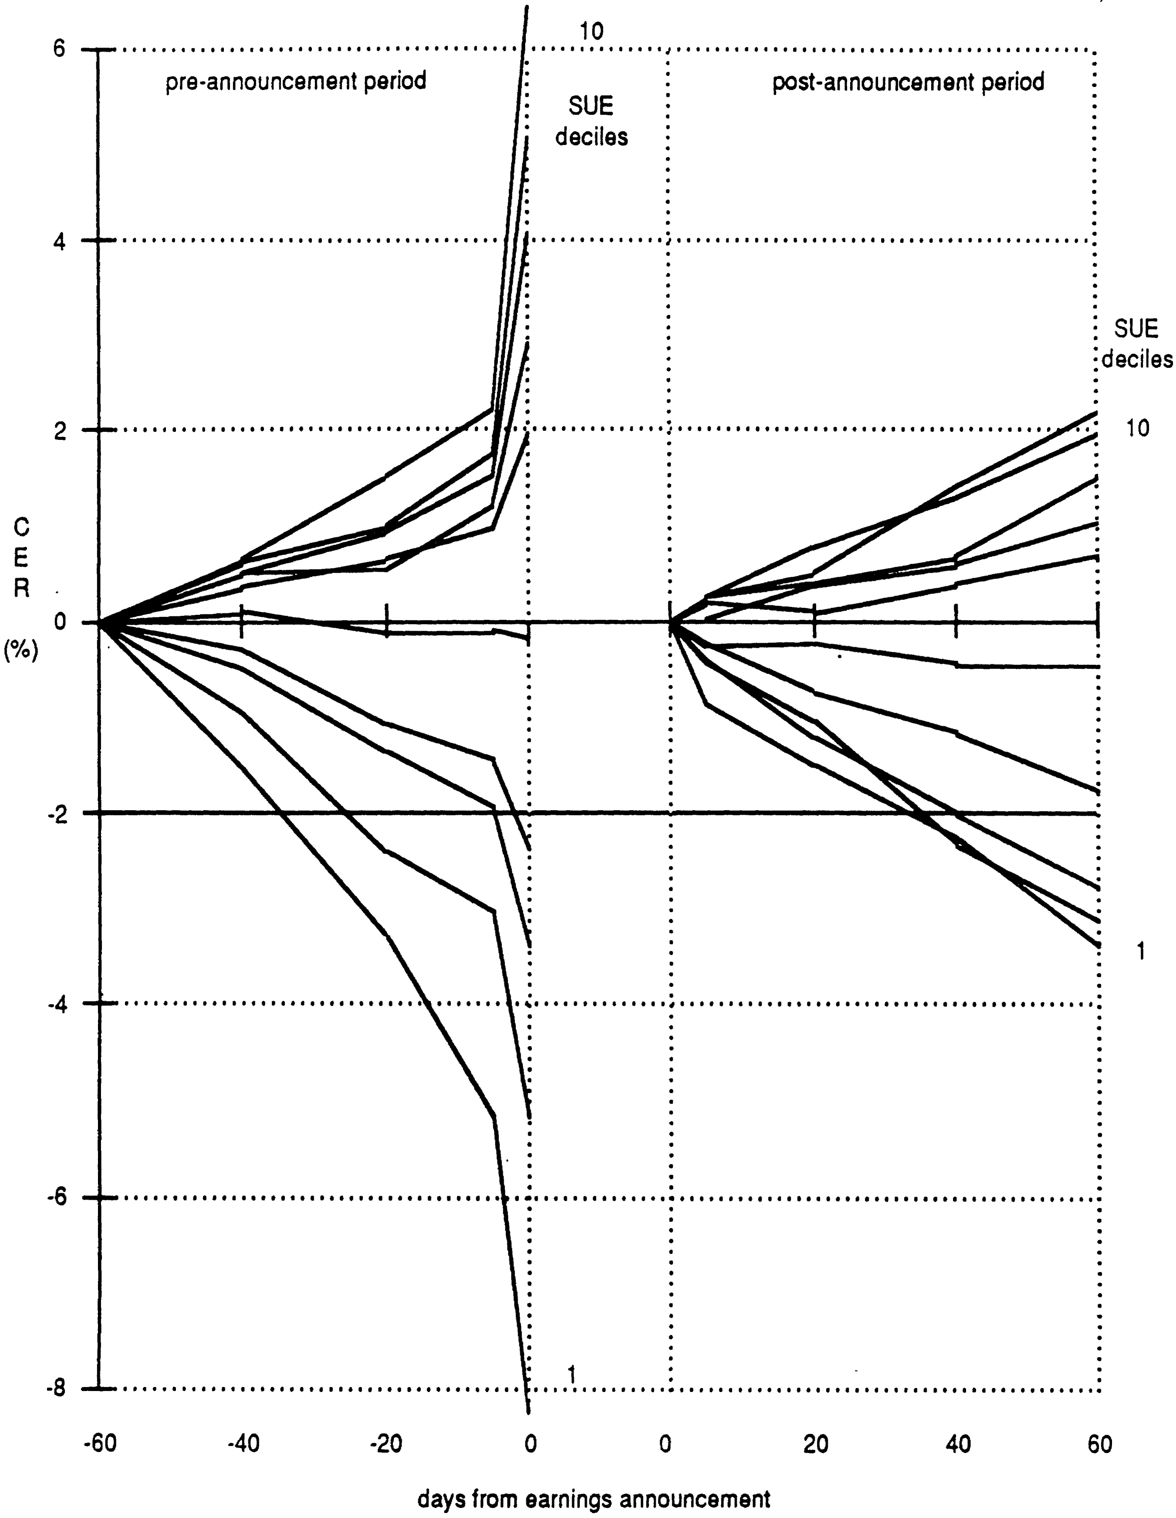

---

1 Small firms are size deciles 1 to 4, based on January 1 market equity values for all NYSE and ASE firms.

2 See Foster, Olsen and Shevlin (1984) for procedures used to compute cumulated excess returns (based on comparisons with benchmark size decile portfolios) and SUE (based on their Model 2).

Figure 6

60-day cumulative excess returns (CER) from SUE strategies
(from long (short) position in firms in highest (lowest) SUE portfolios)[1]

Panel A: FOS control portfolio SUE strategy:[2]

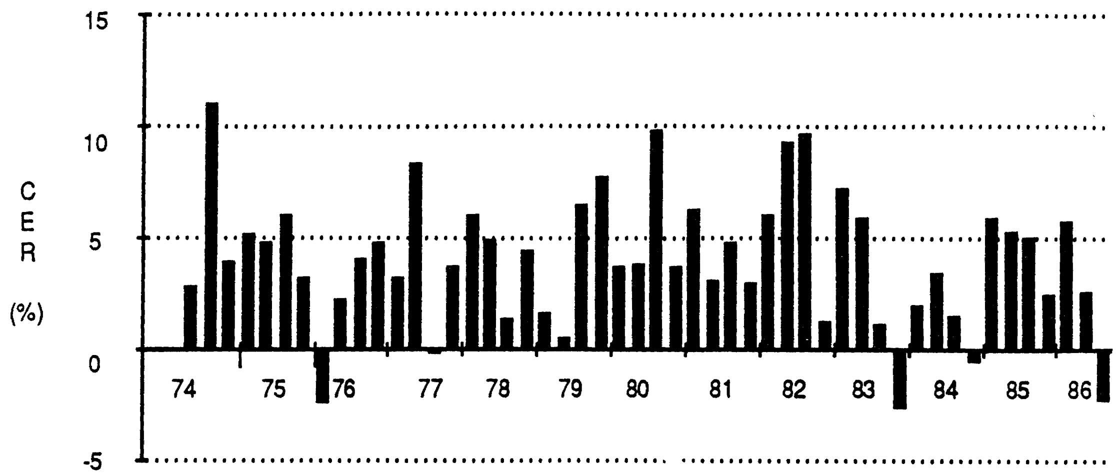

Panel B: Continuously balanced SUE strategy[3]

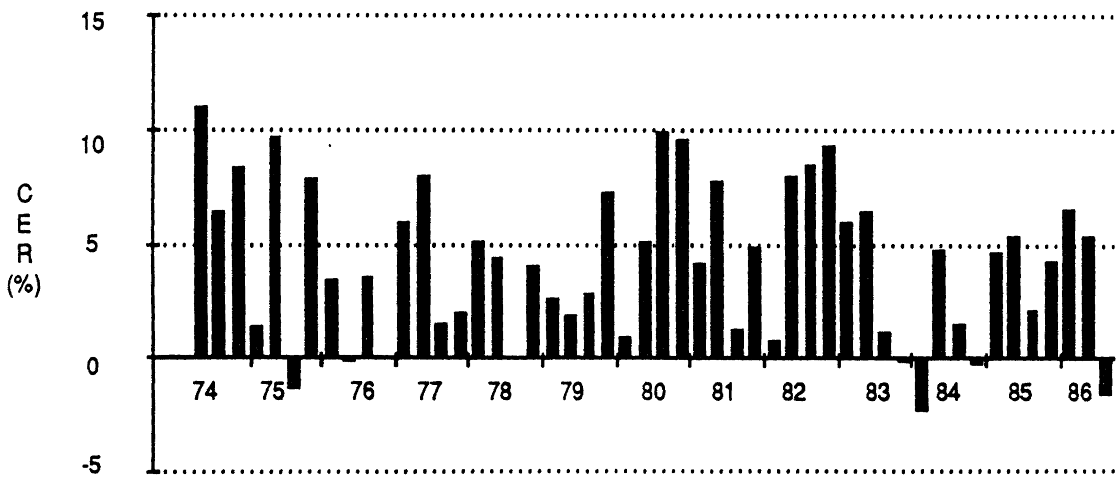

1 See Foster, Olsen and Shevlin (1984) for procedures used to compute SUE (based on their Model 2).

2 Reported by quarter of investment in portfolio, for extreme SUE deciles, using additive CERs (see Foster, Olsen and Shevlin, 1984).

3 Reported by calendar quarter, for extreme SUE quintiles, using compounded returns.

Figure 7

Test of an explanation for the drift, based on costs that impede trading

The plot presents the difference in 60 day cumulative excess returns (CER) between the most positive and most negative extreme SUE portfolios (constructed by splitting the sample into 2, 3, 5, 10,.. 100 portfolios based on SUE). The hypothesis predicts that, if the drift is caused by costs that impede trading, the post-announcement drift should remain less than those costs, regardless of the SUE difference between extreme portfolios. Thus, as the extreme portfolio differences in SUE increase (represented by higher numbers of portfolios towards the right of the graph), the post-announcement CERs should level out, despite increases in the pre-announcement CERs.

Panel A: Overall sample: Pre-announcement and post-announcement excess returns.

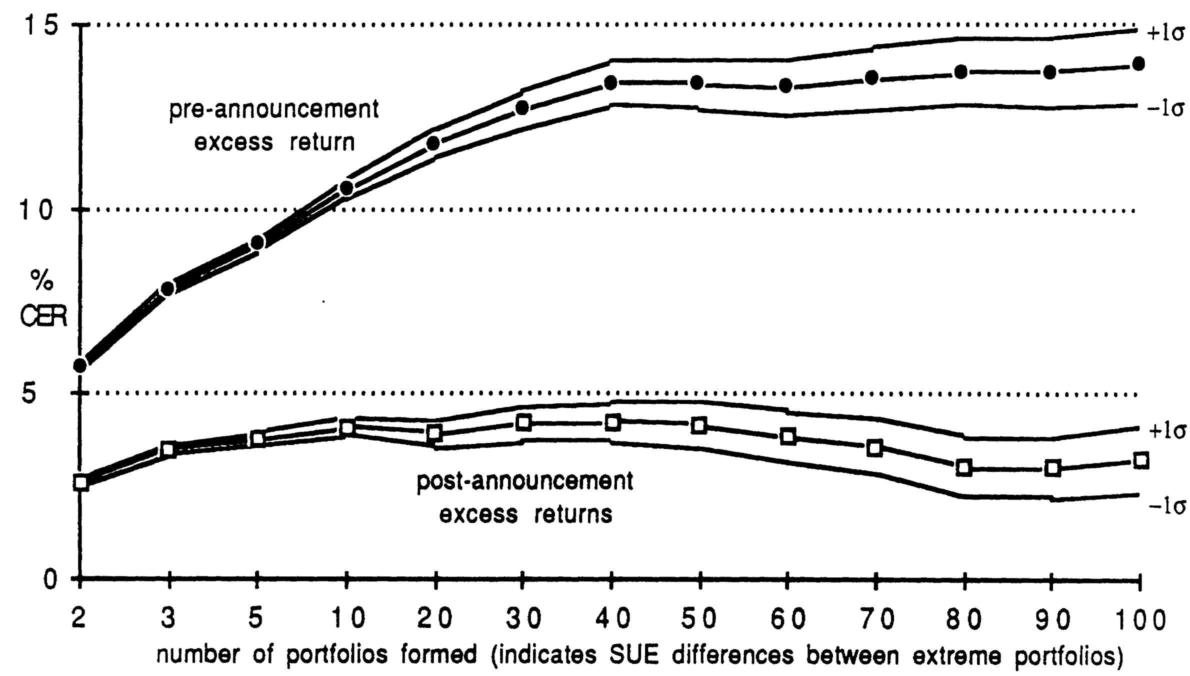

Panel B: Comparison by firm size: post-announcement excess returns.

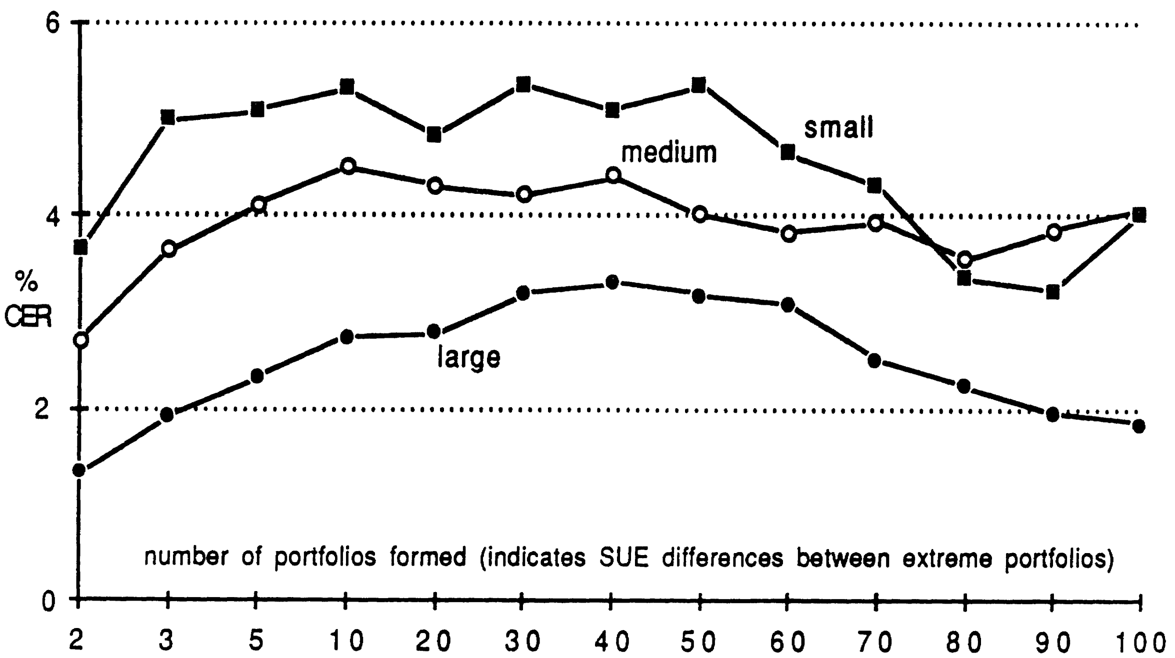

REFERENCES

Amihud, Y., and H. Mendelson, "Asset Pricing and the Bid-Ask Spread," Journal of Financial Economics (December 1986), pp. 223-250.

Ball, R., "Anomalies in Relationships Between Securities' Yields and Yield-Surrogates," Journal of Financial Economics (June/September 1978), pp. 103-126.

Ball, R., "What Do We Know About Stock Market Efficiency?" Working paper, University of Rochester (1988).

Ball, R., and P. Brown, "An Empirical Evaluation of Accounting Income Numbers," Journal of Accounting Research (1968), pp. 159-178.

Ball, R., S.P. Kothari, and R. Watts, "The Economics of the Relation Between Earnings Changes and Stock Returns" Working Paper, University of Rochester (December, 1988)

Beaver, W., R. Lambert, and D. Morse, "The Information Content of Security Prices," Journal of Accounting and Economics (1980), pp. 3-28.

Bernard, V., "Cross-sectional Dependence and Problems in Inference in Market-based Accounting Research," Journal of Accounting Research (Spring 1987), pp. 1-48.

Bernard, V., and J. Thomas, "A Re-Interpretation of Evidence on the Cause of Post-Earnings-Announcement Drift" Working Paper, University of Michigan (1989).

Blume, M. and R. Stambaugh, "Biases in Computed Returns: An Application to the Size Effect," Journal of Financial Economics (June 1983), pp. 387-404.

Brennan, M. J., "Taxes, Market Equilibrium, and Corporate Financial Policy," National Tax Journal (December 1970), pp. 417-427.

Brown, P., G. Clinch, and G. Foster, "Research Issues with Intra-Day Capital Market Data," Working Paper, Stanford University, July 1987.

Brown, P., G. Foster, and E. Noreen, "Security Analyst Multi-year Earnings Forecasts and the Capital Market," Studies in Accounting Research No. 21 (American Accounting Association, 1985).

Charest, Guy, "Dividend Information, Stock Returns, and Market Efficiency," Journal of Financial Economics (September, 1978), pp 297-330.

Chen, N., R. Roll, and S. Ross, "Economic Forces and the Stock Market," Journal of Business (July 1986), pp. 383-404.

DeBondt, W. and R. Thaler, "Does the Stock Market Overreact?" Journal of Finance (July 1985), pp. 793-805.

DeBondt, W. and R. Thaler, "Further Evidence on Investor Overreaction and Stock Market Seasonality," Journal of Finance (July 1987), pp. 557-581.

Dyckman, T. and D. Morse, Efficient Capital Markets and Accounting: A Critical Analysis, Prentice-Hall, 1986.

Fama, E., "Efficient Capital Markets: A Review of Theory and Empirical Work," Journal of Finance (May 1970), pp. 383-417.

Foster, G., C. Olsen, and T. Shevlin "Earnings Releases, Anomolies, and the Behavior Of Security Returns," The Accounting Review (October 1984), pp. 574-603.

Harris, T. and J. Ohlson, "Accounting Disclosures and the Market's Valuation of Oil and Gas Properties: Evaluation of Market Efficiency and Functional Fixation," Working paper, Columbia University (January, 1988).

Holthausen, R., "Abnormal Returns Following Quarterly Earnings Announcements," Proceedings of the Seminar on the Analysis of Security Prices (University of Chicago, May 1983), pp. 37-59.

Holthausen, R., and R. Leftwich, "The Effects of Bond Rating Changes on Common Stock Prices," Journal of Financial Economics (September, 1986), pp. 57-90.

Ibbotson and Associates, Stocks, Bonds, Bills, and Inflation: 1988 Yearbook, Ibbotson and Associates, Chicago, 1988.

Kormendi, R., and R. Lipe, "Earnings Innovations, Earnings Persistence, and Stock Returns," Journal of Business (July 1987), pp. 323-346.

Lakonishok, J., and T. Vermalean, "Anomalous Price Behavior around Repurchase Tender Offers," Working paper, University of Illinois, 1988.

Landsman, W. and J. Ohlson, "Evaluation of Market Efficiency for Supplementary Accounting Disclosures: The Case of Pension Assets and Liabilities," Working Paper, UCLA (November 1987).

Larcker, D., and T. Lys, "An Empirical Analysis of the Incentives to Engage in Costly Information Acquisition: The Case of Risk Arbitrage," Journal of Financial Economics (March 1987), pp. 111-127.

Lehmann, B. "Fads, Martingales, and Market Efficiency," Working Paper, Columbia University (1988).

McNichols, M., "Evidence of Informational Asymmetries from Management Earnings Forecasts and Stock Returns," The Accounting Review (January, 1989), pp. 1-27.

Ou, J. and S. Penman, "Financial Statement Analysis and the Prediction of Stock Returns," Working paper, University of California at Berkeley (1988).

Ou, J., and S. Penman, "Accounting Measurement, Price-Earnings Ratios, and the Information Content of Security Prices," Journal of Accounting Research (forthcoming, Supplement, 1989).

Patell, J. and M. Wolfson, "The Intraday Speed of Adjustment of Stock Prices to Earnings and Dividend Announcements," Journal of Financial Economics (June 1984), pp. 223-252.

Poterba, J., and L. Summers, "Mean Reversion in Stock Prices: Evidence and Implications," NBER Working paper no. 2343 (August 1987).

Rendleman, Jr., R.J., C. P. Jones, and H. A. Latane, "Empirical Anomalies Based on Unexpected Earnings and the Importance of Risk Adjustments," Journal of Financial Economics (1982), pp. 269-287.

Roll, R., "On Computing Mean Return and the Small Firm Premium," Journal of Financial Economics (June 1983), pp. 371-386.

Stoll, H. and R. Whaley "Transactions Costs and the Small Firm Effect," Journal of Financial Economics (June 1983), pp. 57-80.

Summers, L., "Does the Stock Market Rationally Reflect Fundamental Values?" Journal of Finance (July 1985), pp. 591-601.

Watts, R. "Systematic 'Abnormal' Returns after Quarterly Earnings Announcements," Journal of Financial Economics (June-September 1978), pp. 127-150.