

Post-Earnings Announcement Drift: Timing and Liquidity Costs*

Robert H. Battalio and Richard R. Mendenhall[□]

Abstract:

The persistence of the post-earnings announcement drift leads many to believe that trading barriers prevent knowledgeable investors from eliminating it. For example, Bhushan (1994) contends that informed investors quickly exploit the information in earnings surprises driving stock prices to within transactions costs of efficient values and leaving the observed post-earnings price drift unexploitable. We use the exact dates and times of earnings announcements to compare the profits generated by trading immediately after earnings surprises—at quotes actually available to investors—with the profits generated by waiting until the close to trade. We further address the possible implications of commissions, price concession, and arbitrage risk. Under a wide range of assumptions, our results leave little doubt that between 1993 and 2002 an investor could have earned hedged-portfolio returns of *at least* 14% per year after trading costs.

JEL classification: G14

Keywords: Earnings; Post-earnings announcement drift; Anomalies; Bid-Ask Spread; Market microstructure.

*Mendoza College of Business, University of Notre Dame, Notre Dame, IN. 46556-5646. The authors gratefully acknowledge the contribution of Thomson Financial for providing earnings per share forecast data, available through the Institutional Brokers Estimate System. This data has been provided as part of a broad academic program to encourage earnings expectations research. We thank Jeff Bacidore, Kirsten Bergstrand, Shane Corwin, Peter Easton, Tim Loughran, Bill McDonald, Paul Schultz, and seminar participants at the University of Notre Dame.

[□]Corresponding author. Tel.: +1-574-631-6076; fax: +1-574-631-5255.

E-mail address: rmendenh@nd.edu (R. Mendenhall).

1. Introduction

Post-earnings announcement drift, the tendency for cumulative abnormal returns to drift in the direction of an earnings surprise following an earnings announcement, is a well-documented and persistent capital markets anomaly. Early researchers demonstrated that investors initiating positions on *assumed* earnings announcement dates, a specified number of days following the fiscal quarter end, appeared to earn abnormal returns (e.g., Jones and Litzenberger 1970; Litzenberger, Joy, and Jones 1971). Following the availability of machine-readable earnings announcement dates, beginning about twenty-five years ago, researchers continued to find apparent abnormal returns for investors initiating positions at closing prices on *Compustat's* earnings report date. We extend this process of using more precise data to more accurately measure the potential profitability of the drift by documenting the exact day and time that investors could act on the earnings information. Further, we examine the impact of transactions costs on the magnitude of the drift. Researchers are divided on whether the drift exceeds or is bounded by transactions costs. We examine a major component of transactions costs, bid-ask spreads, both immediately following earnings announcements and during non-event times to determine if trading gains exceed these costs. Finally, we assess the possible effects of other trading costs such as commissions and price concessions on our results.

Assuming that transactions occur at closing prices on *Compustat's* earnings report date, as previous studies do, may lead to biased estimates of profitability for several reasons. *Compustat* defines their earnings report date as “the date in which quarterly earnings and earnings per share are first publicly reported in the various news media (such as the *Wall Street Journal* or newswire services)” (p. 234).¹ If the *Compustat* report date is the date the article comes out in the *Wall Street Journal*, then investors are assumed to wait at least six and a half *trading* hours before acting on information that is very time sensitive.² Similar comments

¹ *Compustat* North America User's Guide, 2004.

²After presenting a paper on the drift using traditional methodology, one of the authors of this paper was confronted by academics who claimed the closing price assumption was too heroic. After the same presentation, he was confronted by practitioners (some who claimed to trade on the drift) who claimed the assumption was far too conservative. One practitioner reported that his company initiates trades within “30 seconds of the earnings announcement.”

apply if *Compustat*'s report date is the newswire date and the announcement comes out prior to the close of the market; why would investors hoping to exploit the information in earnings announcements leave money on the table by waiting to trade? On the other hand, if *Compustat*'s report date is the newswire date and the announcement appears on the newswire after the close of trading, then investors are assumed to be clairvoyant—they act prior to receiving the information. Getting the date right is not trivial. For our sample, the difference in hedge returns between assuming investors transact at the close of the *Compustat* report date and assuming they transact one day later is 3.78% per quarter or over 15% per year. One of the contributions of this paper is, for a sample of firm-quarter observations, to carefully document the exact date and time of earnings announcements relative to the *Compustat* earnings report date in order to more accurately assess the potential trading profits of investors attempting to exploit the information in earnings surprises.

Following Livnat and Mendenhall (2006), who find that using two signals of earnings surprise increases the magnitude of the drift, we focus on firm-quarter observations that fall into extreme deciles on the basis of two earnings forecasts—those of analysts and those of a popular time series model. Since most drift studies assume that investors initiate positions at the closing price on the *Compustat* earnings report date and a few use the report date plus one day, we compare returns obtained in these two ways to those generated by initiating positions at the *true* first closing price following the time of the earnings announcement. Through this comparison (and others) we prescribe reasonable assumptions for researchers using daily data from CRSP and *Compustat*.

We next assume that investors initiate their positions at the prevailing quotes (offer for buys, bid for sells) following the exact time of the earnings announcement. If the announcement occurs outside of normal trading hours, we assume the position is initiated at the first firm quote after the market opens the following day. The component of the post-announcement return generated from the actual time of the earnings

announcement to the time of the closing price has previously been ignored by researchers.³ This is understandable due to the high cost of obtaining both exact earnings announcement dates and times and, to a lesser extent, intra-day bid and offer quotes. While investors acting on earnings announcements will not, in every case, obtain this entire return component, those investors attempting to exploit the information in earnings surprises will act as quickly as possible and obtain some fraction of this return. We continue by using our data to examine and compare the profitability of scenarios representing a wide range timing and transactions costs assumptions.

Some of our results can be summarized as follows. Regarding the timing issue, most researchers assume that investors can trade at the closing price on the *Compustat* earnings report date (e.g., Rendleman Jones, and Latané 1982; Foster, Olsen and Shevlin 1984; Bernard and Thomas 1989 and 1990; Bhushan 1994; Ball and Bartov 1996; and Liang 2003). We show that for recent sample periods this assumes that investors trade *prior to* the first public announcement of earnings about one third of the time (33.7%). Assuming investors transact early is much more egregious than assuming they act late. These cases assume investors capture both the last few hours or minutes of the market's anticipation of the surprise and its initial reaction as well as the *post*-announcement drift. Prior research provides no guidance on whether using the closing price on the *Compustat* earnings report date overstates or understates the drift. Our results show that, ignoring transactions costs, relative to trading at the actual first closing price following the public announcement of earnings, assuming investors transact at the closing price on the *Compustat* earnings report date *overstates* the drift by 2.66% per quarter for our sample; assuming investors transact at the close the day *following* the *Compustat* earnings report date (e.g., Livnat and Mendenhall 2006) *understates* the drift by 1.12% per quarter.

³As discussed above, presumably in the majority of cases, correcting for this return component will involve adding the return from the actual time of the earnings announcement to the time of the closing price. In some cases, those when the *Compustat* closing price occurs prior to an announcement made after trading hours that day, correcting for this component involves subtracting a return that investors could not actually obtain. Only empirically can we determine the net effect of the component.

Regarding transactions costs, Bhushan (1994) examines trading costs indirectly and suggests that the drift exists only “up to the level of transactions costs” (p. 50). In other words, the drift is *not* profitable after trading costs. Other researchers (e.g., Ball 1992, p. 333) claim that the drift is too large to be bounded by transactions costs. We suggest that without examining trading costs directly (and using the correct trading day), both conclusions are premature. While trading costs cannot *cause* the drift, it seems important to know if the stock market works in such a way as to render it unprofitable. While many papers present descriptive evidence on the general magnitude of bid-ask spreads (e.g., Huang and Stoll 1996; Bennet and Wei 2006), Lee, Mucklow, and Ready (1993) show that spreads widen in anticipation of earnings announcements. Further, their results suggest that the larger the earnings surprise, the greater the increase in spreads. Without examining the firm-specific spreads at the time of the earnings announcements, we simply do not know how their magnitudes compare to that of the drift.

It turns out that our results contradict those of Bhushan (1994) and show that investors could have profitably traded stocks at quotes that existed at the time of the earnings announcement. Other timing and trading cost assumptions lead to similar conclusions. For example, we compare the scenario above with ones where investors delay their transactions and pay lower–non-announcement period–bid-ask spreads. Taken as a whole, our results strongly suggest the existence of multiple trading strategies that would have allowed investors to earn a hedged-portfolio return of *at least* 14% per year—after considering actual earnings announcement times/dates and actual bid-ask spreads—over the 1993 to 2002 period. Additional results suggest that the existence of commissions, price concessions, and arbitrage risk do not alter our conclusions. Given the results of prior research, these returns are much larger than we would have predicted. In large part this is due to basing earnings surprise portfolios on two types of earnings forecasts (as in Livnat and Mendenhall 2006) and focusing on the most extreme observations—behavior we would expect of knowledgeable investors.

The rest of the paper is organized as follows. The next section briefly reviews the post-earnings announcement literature and motivates our inquiry. The third section describes the sample, data, and variables. The fourth section presents our results. The fifth section discusses why our results may over- or understate the profitability of the drift and the last section concludes.

2. *Literature and Motivation*

2.1 POST-EARNINGS ANNOUNCEMENT DRIFT

The drift of cumulative excess returns in the direction of recent earnings surprises, now referred to as post-earnings announcement drift or the Standardized Unexpected Earnings (SUE) effect, was first documented by Ball and Brown (1968). Since then, the drift has been confirmed by a stream of research using steadily improving research methods which document the anomaly with increasing precision. Ball (1978) provides a comprehensive review of the early literature and discusses the limitations of early studies. He mentions specifically: the failure of some studies to collect earnings announcement dates; *Compustat's* practice of listing updated earnings for those firms that revise earnings reports; computational biases in estimating abnormal performance; and possible errors in estimating the relative risks of stocks. He concludes that researchers' inability to properly specify expected returns is the most likely cause for the drift.

For the most part, the limitations documented by Ball (1978) have been overcome. For example, most studies over the last twenty-five years use earnings dates from *Compustat*. While, as we will show, these dates are not perfect, they represent a vast improvement over assumed dates. Consider, for example, the three large scale studies of the 1980's: Rendleman, Jones, and Latane (1982); Foster, Olsen, and Shevlin (1984); and Bernard and Thomas (1989). Each of these studies uses *Compustat* earnings announcement dates and each reaffirms the drift anomaly on a sample consisting of tens of thousands of firm-quarter observations over dozens of calendar quarters. Further, Livnat and Mendenhall (2006) show that *Compustat's* policy of reporting restated earnings, in lieu of earnings originally observed by investors, has an insignificant effect

on the magnitude of the drift. Finally, Bernard and Thomas (1989) show the drift is robust to a wide range of methods of estimating and concatenating daily abnormal returns.

Ball's (1978) conclusion, that the apparent abnormal returns following earnings announcements are probably due to our inability to properly estimate expected returns (and are, therefore, illusory), is difficult or impossible to completely rule out. But Bernard and Thomas (1989) present a strong case against risk-based explanations for the drift. For example, they show that small firms exhibit *negative* raw returns on average for several days following negative earnings surprises. While not impossible, it seems difficult to justify a return generating model that predicts negative returns on stocks whose returns positively covary with those of market portfolio proxies. Further, Bernard and Thomas show that drift strategies are not risky along any of the dimensions identified by Chen, Roll, and Ross (1986). Finally, since positive-surprise firms exhibit higher raw returns than negative-surprise returns almost every quarter, they ask "Where's the risk?" (p. 32).

In a follow up paper, Bernard and Thomas (1990) show that the pattern of returns around the four earnings announcements following the earnings surprise is what we would expect to observe if the market over weights a particular model of earnings. Specifically, Bernard and Thomas point out that the news media tend to compare announced quarterly earnings figures with earnings from the same quarter of the prior year—a seasonal random walk (SRW) forecast. Bernard and Thomas hypothesize, therefore, that some investors may behave as if earnings follow an SRW pattern. This is important because the SRW model understates the implications of earnings surprises for future earnings levels, i.e., investors using this model of earnings generation will systematically underreact to earnings surprises. This could give rise to the drift. Battalio and Mendenhall (2005) directly examine investors' trading response to earnings and show that those investors who initiate small trades do indeed behave as if they follow the SRW model. That is, investors who initiate small trades respond to SRW errors, but they appear to completely ignore the more economically relevant analyst forecast errors. In other words, Battalio and Mendenhall provide direct evidence that a significant

subset of investors exhibit, on average, precisely the suboptimal investor behavior that Bernard and Thomas (1990) proffer as the cause of the drift.

For many, the Bernard and Thomas (1989, 1990) papers changed the drift from being merely an example of researchers' inability to properly estimate expected returns to being a manifestation of a slow market response to earnings. For example, Ball (1992) again reviews the earnings anomaly literature and concludes that the most likely cause of the drift is not model misspecification as he concluded in 1978, but rather "either market inefficiency or substantial costs of investors acquiring and processing information" (p. 319).

Bhushan (1994) extends Bernard and Thomas (1990) and Ball (1992) by imagining the market dichotomized into two types of investors: those who understand the true time-series properties of earnings and those who do not. The former are professional investors with low information processing costs and the latter are Bernard and Thomas's unsophisticated investors who model earnings as a seasonal random walk. The latter group, as stated above, underestimate the implications of current earnings innovations for future earnings levels. Bhushan, like Bernard and Thomas (1990), hypothesizes that these investors, by systematically underreacting to earnings surprises, cause market prices following earnings announcements to be biased. Bhushan goes on to hypothesize that the trading actions of the former group, e.g., "arbitrageurs and other fundamentally-minded investors" (p.49), tend to undo the pricing effects of these biases, but only to within transactions costs. That is, sophisticated investors will undo the potentially biasing actions of unsophisticated investors only to the point that their actions are profitable *after transactions costs*.

Bhushan (1994) uses price as a proxy for bid-ask spread and commissions and recent dollar trading volume for indirect trading costs. He provides empirical evidence that the magnitude of the drift is positively associated with both proxies for transactions costs. Three limitations of Bhushan's study are that he uses *Compustat* earnings report dates, he relies on proxies for trading costs—he does not directly observe bid-ask spreads—and he provides no evidence that the drift is actually bounded by trading costs. The proxies that he

uses for transactions costs, share price and recent trading volume, are significantly positively correlated with measures related to average investor sophistication such as firm size and institutional holding. Bartov, Krinsky, and Radhakrishnan (2000) provide results suggesting that when institutional holding is added to Bhushan's analysis, trading cost proxies are no longer significant. Their results suggest that trading costs are not binding. We explicitly consider the impact of transactions costs on the drift an open question.

2.2 THE TIMING ISSUE—WHEN CAN INVESTORS ACT ON EARNINGS INFORMATION?

Since studies started using announcement dates from *Compustat*, the issue of when investors can act on earnings information has been largely ignored. As stated above, some early studies assume earnings have been announced by a certain date relative to the end of the fiscal quarter. For example, Jones and Litzenberger (1970) assume investors act two months following the fiscal quarter end and Litzenberger, Joy, and Jones (1971) assume investors act with a three month lag. The three large-scale studies of the 1980's, Rendleman et al (1982), Foster et al (1984), and Bernard and Thomas (1989), however, all assume that investors trading on earnings information can transact at the closing price on the date reported by *Compustat*.⁴

Evidence provided by Berkman and Truong (2006) suggests that the assumption that investors transact at the closing price on the *Compustat* earnings report date may have been reasonable for the sample periods used in the studies cited above, but is probably not reasonable for more recent samples.⁵ They show that since about the mid 1990s much of the response to earnings occurs on the day following the *Compustat* report date. Using 2000-2004 data, they show that a large fraction of earnings announcements now occur after the close of trading. For their sample, the fraction of earnings announcements that occur after trading

⁴ At least two other studies, Ball and Bartov (1996) and Bhushan (1994), use the Bernard and Thomas (1989, 1990) data and, therefore, also assume buying at the close of trading on the *Compustat* announcement date.

⁵ Some more recent studies assume investors can trade at the close on the *Compustat* report date, e.g., Liang (2003).

hours goes from 42.3% in 2000 to 48.8% in 2004. They compare these numbers to Patell and Wolfson (1982) who find 15% of their announcements occur after trading.

We address this issue by using the exact date and time of the announcement as reported on *Factiva* to infer when investors might reasonably trade on the information. Our primary objective regarding this issue is to determine the impact of different trading time assumptions on the magnitude of the drift. An additional objective is to provide descriptive evidence on the most appropriate return window when using daily data, so we compare our intra-day results to those assuming transactions at closing prices on *Compustat* days 0 and +1.

2.3 THE COST OF LIQUIDITY ISSUE—THE COST OF TRADING IMMEDIATELY

The issue of transactions costs has been discussed throughout the SUE literature. For example, Ball and Brown (1968) conclude that “the market acts without bias at least to within transactions costs” (p. 174). Bernard and Thomas (1989) devote two sections of their paper to the potential implications of transactions costs to the existence and magnitude of the drift. They present evidence both consistent with and inconsistent with the drift being bounded by transactions costs. Bhushan (1994) supports the idea that the drift is bounded by trading costs by relating its magnitude to two firm-specific proxies for transactions costs. But most empirical studies assume that investors can trade at closing prices without an allowance for the cost of demanding liquidity—the cost of implementing quick trades. Investors seeking to maximize the profits generated by using a SUE-based trading strategy, however, are unlikely to wait until the end of the day to trade. Using intra-day quote data, we examine the importance of assuming investors wait until the close to initiate SUE-based positions by assuming investors initiate their positions at the quotes that were actually available when news of the earnings surprise first became public. We also examine whether SUE-based strategies that involve trading at closing prices are robust to the introduction of liquidity costs.

3. *Description of the Variables and Sample*

We begin by constructing a sample of firm-quarter observations with earnings, size, price, and return data from *Compustat*, *I/B/E/S* and *CRSP*. For each observation in this sample, we construct both time series (seasonal random walk–SRW) and analyst forecast errors. Our sample begins in 1993 and ends after the first quarter of 2002. The sample begins in 1993 due to the availability of TAQ data and of an exchange traded fund (ETF) that we use for intraday return adjustment (described below).

For a firm-quarter observation to qualify for the initial sample, we require the following data: earnings per share, earnings per share lagged four quarters, relevant adjustment factors, the earnings announcement date and the firm’s subsequent earnings announcement date from *Compustat*; earnings per share, at least one analyst earnings forecast, and an earnings announcement date from *I/B/E/S*; and stock returns, firm size, and a stock price of at least \$1.00 from *CRSP*. For an earnings forecast to qualify it must be made within 90 days of the earnings announcement. To ensure that we have lined up *Compustat* and *I/B/E/S* data properly and to ensure we have a close earnings announcement date approximation, we require that *Compustat* and *I/B/E/S* earnings announcement dates agree to within two calendar days. These selection criteria lead to a sample of 86,807 observations.

From this initial sample, we retain as good- (bad-) news observations only those firm quarters whose SRW *and* analyst forecast errors would have placed them in the top (bottom) forecast deciles for the previous calendar quarter (see Foster, Olsen, and Shevlin 1984). Using ticker symbols, we then match the extreme-decile observations with the NYSE TAQ database to obtain the bid and ask prices that exist immediately following the earnings announcement. We require that observations have an initial bid price of at least \$1.00 and an initial bid-ask spread of less than 50% of the mid-point of the bid and ask prices.⁶ We describe the details of the variables below.

⁶Note that all constraints are based on information known to investors.

3.1 ESTIMATING EARNINGS SURPRISE (SUE):

Consistent with most prior studies, we define the earnings surprise as actual earnings per share minus expected earnings per share divided by stock price. Most prior SUE research uses only time series earnings forecasts, while some recent studies use analyst forecasts. Because Livnat and Mendenhall (2006) show that the drift is larger when considering both measures of surprise, we use both time series and analysts' forecasts.

Among those papers using time series forecasts, the most common is the seasonal random walk model: the forecast equals actual earnings per share for the same quarter of the prior calendar year.⁷ We use this forecast to construct one measure of earnings surprise:

$$SUE_{i,q}^{Time\ Series} = \frac{E_{i,q} - E_{i,q-4}}{P_{i,q}}, \quad (1)$$

where $E_{i,q}$ is actual quarterly earnings per share for firm i for quarter q , $E_{i,q-4}$ is actual reported earnings per share for the same quarter of the prior year, and, $P_{i,q}$ is share price twenty days prior to the quarter q earnings announcement.

Our second measure of earnings surprise is similar to the first but is based on analyst earnings forecasts taken from *I/B/E/S*. Specifically, the earnings surprise is estimated as actual earnings per share from *I/B/E/S* minus the average of all outstanding forecasts on the *I/B/E/S* Detail file (among those less than 90 days old) divided by share price twenty days prior to the earnings announcement.⁸

$$SUE_{i,q}^{Analyst} = \frac{E_{i,q} - \hat{E}_{i,q}^{Analyst}}{P_{i,q}} \quad (2)$$

⁷ Foster, Olsen, and Shevlin (1984) find that the SRW model performs as well as more sophisticated models (e.g., ARIMA models) in predicting the drift.

⁸ Livnat and Mendenhall (2006) show that the magnitude of the drift is virtually insensitive to taking actual earnings from *I/B/E/S* or *Compustat*. We confirm that finding (not reported) on our initial large sample.

We classify firms into 10 portfolios based on each measure of SUE. In order to ensure an implementable trading rule, we follow Foster et al (1984) and assign firms to SUE deciles based on the previous calendar quarter's SUE cutoffs.⁹ We then perform the analysis on those observations most likely to be of interest to investors—those in the most extreme surprise decile of each measure of surprise. Next, we have two sets of assistants collect the exact date and time of each of these earnings surprises from the *Factiva* database. We difference the two datasets and reconcile the discrepancies to arrive at our final dataset.

3.2 ESTIMATING ABNORMAL RETURNS

We measure abnormal returns from the time we assume investors initiate their position until the time we assume they terminate it. In some cases, we follow other SUE studies and assume investors initiate (and terminate) their positions at closing prices. But in other instances we assume investors initiate their positions as soon as possible after the earnings announcement becomes public. In these scenarios, the abnormal return is generated over two consecutive periods: from the time investors initiate their position to the close of that day and from the close on initiation day to the close of the termination day. The holding period raw return for the stock is the compound value of the raw returns over the two periods. The stock's expected return is also a compound value of returns for the inter and intraday periods. For the interday period, the expected return is the compound return on an equally-weighted portfolio consisting of all firms in the size (market capitalization of equity) decile of which the stock is a member at the beginning of the calendar year.¹⁰ For the intraday period, the expected return is the contemporaneous return on the Standard and Poor's Depository Receipt (SPDR), commonly referred to as the S&P 500 Spider.

⁹ While Bernard and Thomas (1990) report that the magnitude of the drift is insensitive to this research design choice, we choose this approach because it is implementable.

¹⁰ The sample consists of both exchange-traded (NYSE and Amex) stocks as well as Nasdaq stocks. Market capitalization deciles and expected returns are determined separately for each group. That is, exchange-traded (Nasdaq) stocks are assigned to deciles based on the size distribution of NYSE-Amex (Nasdaq) firms. In calculating abnormal returns, exchange-traded (Nasdaq) sample firms' raw returns are compared to the returns of an equally-weighted index of NYSE-Amex (Nasdaq) firms of the same size decile.

$$FRETQTR_{i,q} = (1 + FRET_{i,q}^{intraday})(1 + FRET_{i,q}^{interday}) - 1 \quad (3)$$

$$ERETQTR_{i,q} = (1 + MRET_{i,q}^{intraday})(1 + DRET_{i,q}^{interday}) - 1 \quad (4)$$

$$ARETQTR_{i,q} = FRETQTR_{i,q} - ERETQTR_{i,q} \quad (5)$$

$FRETQTR_{i,q}$ is the compound raw return on stock i for quarter q from position initiation through position termination, where $FRET_{i,q}^{intraday}$ is the return on the stock from position initiation to initiation day close and $FRET_{i,q}^{interday}$ is the return on the stock from the initiation day close to the termination day close. Similarly, $ERETQTR_{i,q}$ is the expected compound raw return for stock i for quarter q from position initiation through position termination, where $MRET_{i,q}^{intraday}$ is the return on the S&P 500 Spider from position initiation to initiation day close; $DRET_{i,q}^{interday}$ is the return on an equal-weighted portfolio of matched size-decile firms from the initiation day close to the termination day close. $ARETQTR_{i,q}$ is, therefore, our measure of the abnormal return on the stock from position initiation through position termination for stock i in quarter q .¹¹ This paper focuses on return differences between good- and bad-news subsamples. All inferences regarding this return difference are unaltered when using raw returns instead of adjusted returns.

3.3 MICROSTRUCTURE VARIABLES

To evaluate whether an investor could have profited from implementing a SUE strategy, we identify the prices at which an investor could have initiated positions for each sample observation. Since the firm quote rule governs quotes in U.S. equity markets, we obtain intraday quote data from the New York Stock Exchange's Trade and Quote (TAQ) database and construct the National Best Bid and Offer (NBBO) for our sample.¹² We construct the NBBO by determining the highest valid bid and the lowest valid offer at each

¹¹ We reconcile all stock holding period returns calculated using this procedure with returns calculated using TAQ prices along with dividends from *CRSP*. For each subsample, returns calculated using the two methods are the same to within a few basis points.

¹²The firm quote rule mandates that specialists or market makers execute marketable orders for at least the quoted size (which can be for no fewer than 100 shares) at prices that are no worse than their quoted prices. Market makers and specialists are only exempted from this obligation if there is an order ahead or if they are in the process of changing quotes when an order arrives. See Stoll and Schenzler (2006).

moment throughout the trading day.¹³ These are the prices at which at least one liquidity demanding investor could have respectively sold or purchased *at least* 100 shares of the underlying stock. During our sample period, continuous trading for stocks listed on Nasdaq begins each day at 9:30 a.m. as long as the NBBO is not locked (National Best Bid (NBB) equal to the National Best Offer (NBO)) or crossed (NBB exceeding the NBO) and ends each day at 4:00 p.m..¹⁴ Continuous trading for stocks listed on the New York Stock Exchange (NYSE) and the American Stock Exchange (AMEX) typically begins after the opening call auction (conducted some time after 9:30 a.m.) and ends at 4:00 p.m., usually with a closing call auction.¹⁵ In addition to constructing the existing quotes at the time of the earnings announcement, we also estimate average non-event bid-ask spreads. For each stock, we estimate the non-event (or normal) bid-ask spread as the time-weighted spread during the twenty trading days beginning five weeks prior to the date of the announcement.¹⁶

3.4 HOLDING PERIOD AND TRADING ASSUMPTIONS

Based on the variables above, we report results for seven different trading assumptions explained more fully in Figure 1. Scenarios 1 and 2 represent the closing-price assumptions that frequently appear in the literature. Scenario 3 is similar to scenarios 1 and 2, but it uses the *actual* first closing price following the earnings announcement instead of relying on the *Compustat* report date. In Scenario 4, we assume that investors initiate their positions at the midpoint of the bid-ask spread prevailing at the time that the earning surprise is first reported by *Factiva*. While Scenario 4 clearly represents an upper bound on the profitability of actually trading on earnings surprises, the last three scenarios we consider make more conservative assumptions.

¹³Following Bessembinder (2003), quotes are omitted if either the bid or the ask price is non-positive, if the quotes are associated with trading halt or designated order imbalances, or if the quotes are not firm.

¹⁴For those Nasdaq-listed stocks with a locked or crossed NBBO at 9:30 a.m, trading begins once the NBO exceeds the NBB. See Cao, Ghysels, and Hatheway (2000).

¹⁵See Madhavan and Panchapagesan (2000) and Cushing and Madhavan (2000) respectively for more information on the NYSE's opening and closing auctions.

¹⁶See McNish and Wood (1992) for more information on time-weighted bid-ask spreads.

In Scenario 5, we assume that investors decide to trade as soon as possible following the announcement and are willing to pay the cost of demanding immediate liquidity. As Lee, Mucklow, and Ready (1993) show (and we confirm) bid-ask spreads increase in size around the time of earnings announcements, especially in those cases representing large surprises. So, in this scenario, investors decide that the benefits of acting quickly more than offset the inflated costs associated with demanding liquidity at this time. The cumulative abnormal returns associated with post-earnings announcement drift are concave and, if they reverse at all, do not reverse for at least a year following the earnings surprise. Investors hoping to profit from the drift, therefore, have considerable latitude in when and how they terminate their positions. Since investors need not be in a rush to terminate their positions, they need not pay for demanding liquidity, especially in more actively traded stocks. We therefore assume in this scenario that investors terminate their positions at the closing price following the subsequent earnings announcement without a specific payment for demanding liquidity. Scenario 6 is similar to Scenario 5, but assumes that investors incur one-half of the normal (non-event) bid-ask spread when they terminate their positions. This scenario may be more reasonable for less-actively traded stocks. Finally, Scenario 7 assumes that investors wait from the time of the announcement to the first following close and pay one-half normal bid-ask spread for initiating their position and one-half bid-ask spread for terminating their position. So, compared to Scenario 6, in Scenario 7 investors give up the gains to acting immediately in exchange for paying a lower cost of demanding liquidity when initiating their positions.

4. *Empirical Results*

4.1 PRECISE ANNOUNCEMENT TIMES

We begin by examining the precise date and time of earnings announcements for the good- and bad-news subsamples. Table 1 shows that using the closing price of any specific day relative to the *Compustat* earnings report date does not accurately capture the time when investors could trade on earnings information.

For example, Table 1 shows that for over half the sample (59.5% of good-news and 54.5% of bad-news observations) investors received earnings news in time to act by 10:00 a.m. on the *Compustat* report date. But by the close of trading on that day, six trading hours later, investors still could not have acted for about one third of the sample (29.5% for good-news and 37.2% for bad-news observations). Studies that assume investors trade at the closing price on the *Compustat* report date assume investors are clairvoyant one third of the time. But studies that assume investors initiate their positions at the close of the day following the *Compustat* report date assume that over 99% of the time investors wait more than six trading hours before taking a position and over half the time they wait an additional trading day. Results presented later examine the sensitivity of different trading-time assumptions on the estimated profitability of acting on earnings surprises.

4.2 FORECAST ERRORS AND SUBSEQUENT RETURNS–LARGE SAMPLE

Table 2 is a five-by-five grid that divides our initial large sample of firm-quarter observations by analyst forecast error from left to right and by SRW forecast error from top to bottom. Despite a different sample period, different data requirements, and a different definition of abnormal returns, our results are consistent with those of Livnat and Mendenhall (2006). Specifically, comparing the spread along the first five entries in the bottom row to the first five entries in the right-most column, analyst forecast errors provide a greater post-earnings announcement drift (5.1% from top to bottom quintile) than do SRW forecast errors (3.2%). Also consistent with Livnat and Mendenhall (2006) combining the two errors gives a larger drift than either error individually. That is, comparing the abnormal return for observations that fall into the top quintile for both errors (5.4%) to those in the bottom quintile for both errors (-2.3%) gives a post-earnings announcement drift magnitude of 7.7%, which compares to 6.9% for Livnat and Mendenhall (2006). These results indicate that our results are roughly comparable to those of prior research.¹⁷

¹⁷ Livnat and Mendenhall (2006) obtain more symmetric abnormal returns of 3.5% for similarly defined extreme good-news events and -3.4% for the bad-news subsample. When we trim 0.5% from each end of the sample on the basis of abnormal returns, we obtain 3.8% for good-news events and -3.2% for bad-news events. Until the robustness section, we base all results on size-adjusted returns that are neither winsorized nor trimmed. We later discuss the effects of outliers on our results.

The remaining empirical tests focus on a subset of the firm-quarter observations as described above. Specifically, we focus on those observations of the most interest to investors—observations whose analyst forecast error and SRW forecast error would have placed them into the top decile of all surprises in the previous calendar quarter.

4.3 IMPACT OF TIMING ON RETURNS

In Table 3 we present results based on initiating positions at various closing prices and initiating positions at the midpoint of the bid-ask spread following the release of the earnings surprise. Panel A provides results for the good-news sub-sample—those observations with extremely positive analyst and SRW forecast errors; Panel B provides results for the bad-news sub-sample—those with extremely negative analyst and SRW forecast errors; and Panel C provides the returns for a hedged portfolio that is long the good-news stocks and short the bad-news stocks. The first two columns provide results that can be obtained without determining the exact announcement date and time. The first column assumes that investors initiate their positions at the closing price on the *Compustat* report date and the second column assumes initiation at the closing price one day after the *Compustat* report date. As discussed above, each of these dates has been used in the SUE literature.

Note that assuming initiation at the *Compustat* date closing price leads to an abnormal return of 8.80% for the good-news sub-sample and -4.35% for the bad-news sub-sample for a hedged return of 13.15%.¹⁸ Delaying initiation by one day reduces the magnitudes of returns for the good- and bad-news groups to 7.01% and -2.36%, respectively, for a hedged return of 9.37%. So, delaying by one day reduces returns by just under 2% on both the upside and downside for a reduction in the hedged returns of 3.78%. This clearly suggests that assessing the potential profitability of trading on earnings surprises depends critically on knowing the exact timing of the announcements.

¹⁸ This return is significantly larger than that estimated in most prior SUE studies because we include only those firm-quarter observations that would have ranked in the most extreme positive or negative decile for both of two measures of earnings surprise (time series and analyst) in the prior calendar quarter. As mentioned earlier, returns are calculated using the *CRSP* daily return file and confirmed using TAQ closing prices coupled with dividends from *CRSP*.

The hedge return of 9.37% for extreme deciles is analogous to the 7.7% hedge return obtained when comparing the extreme-quintile cells from Table 2. So, focusing on the intersection of extreme deciles (7.2% of total sample observations) instead of extreme quintiles (18.0% of sample) increases hedge returns by almost 170 basis points per quarter. This finding contradicts a result in Bernard and Thomas (1989) that suggests the drift is bounded by transactions costs. If the drift is bounded by transactions costs, Bernard and Thomas argue that dividing their sample into finer and finer earnings surprise partitions should increase post-earnings announcement returns only up to a point. Consistent with this observation, Bernard and Thomas find that, no matter how fine the earnings surprise partitions, they cannot increase hedge returns beyond about 4%, the result obtained when their sample of earning surprises between 1974 and 1986 is partitioned into deciles (20% of total sample). Citing Stoll and Whaley (1983), Bernard and Thomas state that 4%, for a combined long and short position, is probably within transactions costs for the average firm in their sample. But, like all prior and contemporaneous drift studies, Bernard and Thomas use only time series forecasts. By using two measures of earnings surprise (which should not matter under the transactions costs argument) we are able to increase hedge returns by including only more extreme surprises and we are able to obtain returns that are more than double those of Bernard and Thomas. Since transactions costs for our sample are most likely significantly lower than for theirs (see, e.g., Jones 2002), by focusing on extreme surprises based on two types of forecasts, our data offer preliminary evidence that the drift is not bounded by transactions costs.

The third column of Table 3 assumes that investors initiate their positions at the actual first closing price following the time of the earnings announcement. Calculating these results requires that we obtain exact earnings announcements dates and times from *Factiva*. As documented in Table 1, some announcements occur prior to the *Compustat* report date and some occur after. In the case of the good-news sub-sample, the returns obtained if investors initiate their positions at the actual first close following the announcement is 7.30%. This figure is 150 basis point smaller than that obtained using the close on the *Compustat* report date

and 29 basis points larger than that assuming initiation at the close the day following the *Compustat* report date. For the bad-news sub-sample, investors earn 3.19% (by shorting) if they initiate their positions at the first actual close following the announcement. This figure is 116 basis points smaller than when the *Compustat* report date is used and 83 basis points larger than when using the subsequent day. Netting out the differences, Panel C indicates a hedge return of 10.49% when using the true first closing price following the announcement, compared to 13.15% and 9.37% when using the *Compustat* report date or the following date, respectively. If researchers wish to assume that investors initiate their positions at the first closing price following earnings announcements, these results suggest that using the close of the *Compustat* report date is far too heroic, while assuming initiation the day following the *Compustat* date somewhat understates the magnitude of the drift. But Table 1 shows that assuming investors trade at the *Factiva* closing price is tantamount to assuming that, for the vast majority of observations, investors wait nearly the entire trading day to place their orders.

The final column in Table 3 assumes investors initiate positions immediately following the actual time of the earnings announcement as reported on *Factiva* at the midpoint of the prevailing NBBO. This scenario assumes investors capture all of the profits available from trading on earnings information. As a result, this scenario may represent an unrealistic upper bound on the profits available to investors trading on earnings information.

Note that using the exact time of the earnings announcement significantly increases abnormal returns. For both the good-news and bad-news sub-samples, assuming that investors can initiate their positions immediately following the earnings announcement leads to post-earnings announcement drift about one and one quarter percent larger than when investors initiate their positions at the first actual close following the earnings announcement. For the good- (bad-) news sub-sample, the abnormal return is 8.55% (-4.55%), which is 125 basis points (136 basis points) larger than assuming investors trade at the close. Panel C shows that the hedge return of 13.10% is almost the same as the 13.15% obtained using the *Compustat* report date

closing price. This highlights the fact that assuming investors trade at the closing price on the *Compustat* report date is clearly unrealistic. In the next section, we compare the profitability of scenarios using different timing assumptions and estimates of the bid-ask spread.

4.4 IMPACT OF BID-ASK SPREAD ON INVESTMENT RETURNS

Panel A of Table 4 presents descriptive statistics for relative bid-ask spreads during both non-event (or normal) periods and as they exist immediately following the earnings announcement. Normal bid-ask spreads are defined as the time-weighted spread during the twenty trading days beginning five weeks prior to the date of the announcement divided by the TAQ closing price on the date that the earnings surprise is reported by *Compustat*. The first and third columns show distributional statistics for non-event period bid-ask spreads for the good- and bad-news subsamples, respectively. Note that for the good-news sub-sample the mean and median quoted spreads are 2.69% and 1.97%, respectively. For the bad-news sub-sample the quoted spreads are somewhat larger with a mean (median) of 3.63% (2.79%). The second and fourth columns indicate that for both good- and bad-news sub-samples the bid-ask spreads are larger at the time of the earnings surprise, but the effect is somewhat larger for bad-news firms. Based on the medians, relative quoted spreads increase by 14.7% for good-news firms and by 23.7% for bad-news firms. These results are generally consistent with Lee, Mucklow, and Ready (1993), who find that the relative quoted spreads for 230 actively-traded NYSE-listed securities in 1988 increase by an average of 12.5% in the half hour containing an earnings announcement associated with an absolute return in excess of 2%. They find that quoted spreads quickly return to normal levels, leading the authors to conclude that liquidity providers are sensitive to changes in information-asymmetry risk. Our examination of quoted spreads at the instant earnings surprises are announced suggests Lee, Mucklow and Ready (1993) understate the sensitivity of liquidity providers to adverse-selection risk.

Panel B of Table 4 presents descriptive statistics for trading volume during both non-event (or normal) periods and on the day that investors could first trade based on the earnings surprise. Normal trading

volume is the average daily trading volume over the twenty trading days beginning five weeks before the *Factiva* announcement date. Perhaps surprisingly, 75% of the firms in our sample have a normal average daily trading volume in excess of 19,300 shares. For good-news firms, the median daily trading volume increases from an average of 59,475 shares in the days preceding the announcement of the earnings surprise to 140,300 shares on the day of the earnings surprise. The median bad-news firm experiences a much smaller increase in trading volume: from an average daily volume of 58,410 shares to an event-day trading volume of 90,900 shares. Overall, Panel B of Table 4 suggests that for 75% of the good-news (bad-news) firms in our sample, a marginal investor seeking to trade five percent of the volume on the day of the earnings surprise could have acquired (sold) at least 2,300 shares (1,300 shares) of stock.

In Table 5 we explicitly estimate the impact of the bid-ask spread on returns available to investors by comparing the profitability of Scenarios 5, 6, and 7 to each other and to the last scenario reported in Table 3: the scenario in which we assume investors trade immediately following the earnings announcement and do not pay for demanding immediate liquidity. The first column of Table 5 reports results for this scenario, Scenario 4, and is the last column from Table 3 repeated.

Scenario 5 is the first in which we consider the costs to investors of demanding liquidity. Here we assume that investors attempting to profit from earnings surprises act immediately and pay one-half the (inflated) announcement period bid-ask spread for doing so. Specifically, for good- (bad-) news announcements, investors initiate their positions by buying (selling) at the National Best Offer (National Best Bid). We assume that investors unwind their positions at the closing price following the first subsequent earnings announcement (from *Factiva*). In this scenario investors do not pay for liquidity when terminating their positions. CAR plots based on the drift clearly show that after three months average abnormal returns are slightly positive—or at least not negative (e.g., Rendleman et al 1982, Foster et al 1984, and Bernard and Thomas 1989). So, if trading next week is as good as trading today, why should investors pay a premium for

trading today? It seems that one reasonable assumption is that investors could have avoided paying any part of the bid-ask spread by using strategic limit orders to terminate their positions.

Results in the second column of Table 5 indicate that for both good- and bad-news cases profitability drops by about one-half of the announcement period bid-ask spread as documented in Table 4. (The difference is not exactly one half because of differences in denominators.) Note that abnormal returns remain significantly positive (negative) for the good- (bad-) news subsample. The hedge returns remain economically significant at 8.72% per quarter and highly statistically significant with a traditional (time series) t-statistic of 7.18 (7.09).¹⁹ If investors could trade immediately following earnings announcements at existing quotes and eventually unwind their positions without paying for liquidity, their strategies would have been highly profitable over our sample period.

Scenario 6 is identical to Scenario 5 except that we deduct one half of the stock's non-event quoted bid-ask spread as the cost of terminating the position. Depending on the exact portfolio strategies employed, investors may be willing to pay normal liquidity costs to have some control over the timing of their terminating transactions. For example, investors may in some cases wish to terminate positions quickly to synchronize buy and sell transactions in order to maintain hedged positions. As expected, the results for Scenario 6, which appear in the third column of Table 5, are weaker than the results for Scenario 5. The mean return for the bad-news subsample is, for the first time, not significantly different from zero. But the average return for the good-news subsample is significantly greater than zero and significantly greater than that of the bad-news subsample as indicated by a hedge return of 5.57% with t-statistics greater than 4.00. This scenario most likely overstates the costs of trading and understates the returns available to investors. As throughout this paper, we use firm quotes to construct the bid-ask spread. For investors trading NYSE-listed securities in normal market conditions, the assumption that liquidity demanders trade at the NBBO overstates actual trading costs. For example, Petersen and Sirri (2003) find that market orders routed to the NYSE

¹⁹See Fama and MacBeth (1973) for more information on the time series t-statistic.

between October 28, 1996 and November 22, 1996 when the bid-ask spread was an eighth (a quarter) were executed at prices that were, on average, 1.2 cents (8.4 cents) better than the prevailing quote.

Finally, in Scenario 7 we assume that investors initiate their positions at the first actual closing price following the earnings announcement, terminate their positions at the closing price following the next earnings announcement, and pay one non-event bid-ask spread. This scenario is the most conservative scenario that we consider. It is all but certain that investors can trade by the first close following the earnings announcement. Examination of Table 1 suggests that in over 87% of the cases, investors would have at least six trading hours to complete their transactions. This scenario also assumes that investors demand liquidity for both initiating and terminating their transactions—and that they pay the full quoted bid-ask spread for doing so. Despite these factors, abnormal returns for the good-news subsample are significantly positive and the hedge return is 4.18% with t-statistics above 3.00.

Table 5 represents the most important results of this paper. What is the profitability of the drift under different assumptions regarding the timing of investors' trades and the transactions costs they incur? While readers can assign their own probabilities to the different scenarios presented here, we interpret these results as demonstrating that knowledgeable investors could have very profitably exploited the information in earnings announcements over our sample period. After briefly discussing the robustness of our results, in the next section, we discuss different factors that may lead to our estimate of the profitability of the drift being over or understated.

4.5 ROBUSTNESS

Table 6 presents the hedge returns for the most conservative case from Table 5, Scenario 7, under several alternative assumptions. First, in scenarios not based on closing prices, we assume investors can trade only at firm quotes to which liquidity providers are committed. These quotes are always for a specific number of shares (at least 100) rather than for a specific number of dollars. If an investor takes equal-share, as opposed to equal-dollar, positions in each stock, returns should be price weighted instead of equally

weighted. The second row in Table 6 shows that price weighting the returns reduces the most conservative estimate of available returns from 4.18% per quarter to 3.67%. Because of the lower cross-sectional variability of higher priced stocks, however, both the traditional and the time series t-statistics are actually a little higher.

Bhushan (1994) uses price as an inverse proxy for the direct costs of trading, e.g., commissions. Returning to equal weighting, we next omit all stocks with prices less than \$10.00 (see, e.g., Bartov et al 2000). The hedge return is 3.61% per quarter and is very similar to the 3.67% obtained when observations are price weighted. In this case, t-statistics are smaller in part because of the smaller sample size.

Bhushan (1994) uses recent dollar trading volume to proxy for indirect costs of transacting. The next row of Table 6 presents results after deleting stocks with relatively low levels of event period dollar trading volume. Specifically, in the fourth row of results of Table 6 we present the Scenario 7 returns after omitting all observations that rank in the bottom quartile in dollar trading volume on the earnings announcement day.²⁰ The average hedge returns are actually larger than those of the base case. Our results are apparently not attributable to relatively high hedge returns among low-priced or illiquid stocks. This is consistent with the fact that our sample requires observations to have at least one analyst forecast reported to *I/B/E/S* in the 90 days preceding the earnings announcement.²¹ Further evidence that our results are not attributable to small or illiquid stocks is that, with no filters applied, the mean and median NYSE (Nasdaq) stock in our sample has a market capitalization greater than that of the median NYSE (Nasdaq) stock.

Mendenhall (2004) suggests that firms with high idiosyncratic (or arbitrage) risk are difficult to hedge and therefore exhibit larger drifts. We estimate arbitrage risk as the unexplained return standard deviation from a market model regression of stock returns on those of the S&P 500 estimated over 250

²⁰ We replicated this test after omitting the bottom quartile of stocks based on non-event trading volume, days -25 through -6 relative to the announcement date, and the results are nearly identical to those presented.

²¹ Livnat and Mendenhall (2006) report that a similar restriction rules out over 40% of their potential sample observations and that the remaining sample observations consist of firms with significantly larger market values, book values, and share prices than those omitted (see Table 1).

trading days ending ten days prior to the earnings announcement. Consistent with Mendenhall (2004), the correlations between arbitrage risk and returns available for scenarios that do not deduct part or all of the bid-ask spread are significant in the expected directions for both subsamples. But the correlation between arbitrage risk and the returns under Scenario 7 of Table 5 for the good-news (bad-news) subsample is 0.041 (-0.004). The good-news correlation is statistically significant at the five percent level, but neither correlation seems economically large. To ensure that our results are not driven by high-arbitrage risk stocks, in the fifth row of Table 6, we present results after omitting the 25% of our sample observations with the highest levels of arbitrage risk. Note that the returns are similar to those obtained when omitting the bottom quartile on the basis of dollar trading volume.

The final three rows of Table 6 repeat the analyses of rows three through five after winsorizing returns at 1% on each end of the remaining distribution. While extreme return observations have, in some cases, a non-trivial effect on the hedge returns, they are not responsible for the results. In each case, Scenario 7—the most conservative case—hedge returns remain statistically significant at traditional levels and exceed three percent per quarter.

5. Discussion

5.1 WHY MIGHT PROFITABILITY BE OVERSTATED?

In some of our scenarios, investors are assumed to trade immediately upon observing the earnings news. No investor can count on being among the first to trade following every earnings announcement and so this assumption overstates the profitability of the drift. But not all scenarios make this assumption and yet all appear to remain profitable.

We ignore at least two potential costs of trading: price concession and commissions. Price concession is the cost associated with demanding more liquidity than the market is willing to provide. It is therefore much more relevant to large investors such as institutions than to individual investors. Institutional

investors today go to great lengths to break up their orders to avoid both price concession and being identified as potentially informed traders. A skeptic might interpret our results as the returns to an investor who purchases the amount available at the quoted prices, which is at least 100 shares. However, given the positive returns available in Scenario 7, which involves trading at closing prices following the earnings announcement, it is reasonable to assume that an investor could have amassed at least five percent of the trading volume on the earnings announcement date at prices no worse than the closing price on the announcement date. For the median stock in our sample, five percent of the announcement day trading volume is just over 4,500 shares. Further, recall that results presented in the last section indicate that abnormal returns are not concentrated among low-volume stocks. If price concession represents a significant barrier to trading for our sample, low volume stocks should exhibit greater abnormal returns. But the correlation between recent dollar trading volume and returns available in Scenario 7 for the good-news (bad-news) sample is 0.012 (-0.024) and is not significantly different from zero.

Commissions are fees paid to brokers to execute trades. During our sample period commissions for institutions are quoted in cents per share, while commissions for individuals vary greatly in form and magnitude. We believe that the effect of commissions is probably much less than many would suspect. First, any investor who is willing to buy and sell at stated quotes and is willing to commit to at least a few hundred trades per year, should be able to negotiate commissions to zero. Further, in the next section, we discuss why our results may understate the profitability of the drift. We point to one specific assumption that would probably offset any costs of commissions. Finally, Bhushan (1994) uses inverse share price to proxy for percentage commissions. The lack of sensitivity of our results to stock price documented in the last section suggests that commissions do not play a significant role in our findings. Similar to the logic above for trading volume, if commissions represent a significant barrier to those attempting to exploit the information in earnings announcement, low-price stocks should exhibit greater drifts. But the correlation between returns

available under Scenario 7 of Table 5 and the price of good-news (bad-news) stocks just prior to the earnings announcement is -0.010 (-0.019) and is not significant.

5.2 WHY MIGHT PROFITABILITY BE UNDERSTATED?

Like prior research, we examine trading strategies that are based solely on earnings surprise. Although, unlike most prior studies, we use two measures of earnings surprise, we make no attempt to weight the two earnings signals; we simply take all observations in the extreme deciles of both. Further, we ignore all other variables that might, when used with earnings surprise, enhance the drift. Two convenient examples are demonstrated by Chan, Jegadeesh, and Lakonishok (1996) and Collins and Hribar (2000). These papers find that the drift is largely independent of the return momentum and accrual anomalies, respectively. Presumably by combining these signals investors could generate strategies that are more profitable than relying on earnings surprise alone. Further, one can imagine modeling post-earnings announcement returns net of costs as a function of a wide range of variables. A model that balances before-cost drifts with variables related to trading costs might prove more profitable.

Some scenarios discussed above are clearly too optimistic, such as that represented by the returns generated by Scenario 4. In this scenario investors are assumed to trade immediately following the earnings announcement and pay no price for demanding immediate liquidity. But other scenarios are probably overly conservative. Take Scenario 7. In the vast majority of cases, investors wait at least six trading hours to transact and then pay the full price, the quoted bid-ask spread, for demanding to trade quickly. As discussed above, investors probably need not demand liquidity when closing their positions and when they do, they will probably pay less than half of the bid-ask spread as assumed here.

Many studies break samples into earnings-surprise deciles. If sequential earnings forecast errors were unrelated, firm-quarter observations in the top (bottom) surprise decile would have a one in ten chance of being in the top (bottom) decile next quarter. But forecast errors tend to be somewhat positively

autocorrelated.²² Our analysis assumes that investors close their positions every quarter. Different SUE strategies would have different turnover rates and we would expect conscientious SUE investors to optimize along this dimension. We expect that the optimal annual turnover is less than the 400% we assume.²³

Finally, in some scenarios we assume that investors can trade only at firm quotes to which market makers were actually committed immediately following the earnings announcement. In addition to generally ruling out any price improvement, which is common among NYSE stocks, this means that when an earnings announcement occurs after the close of trading for NYSE stocks, we do not assume that investors could initiate positions at the opening call auction. Instead, we assume that investors could initiate positions at the first firm quotes following the opening call. Our logic is that submitting additional orders for the opening call could cause the equilibrium price to shift against the investor. But for small orders (e.g., 100 or 200 shares) the effect on the opening call price would probably be negligible. When we relax this assumption and allow investors to trade at the opening call, this increases profits by approximately \$0.085 per position on average across all observations. This would offset a large fraction of the commissions even for those investors not savvy enough to negotiate the lowest rates.

6. Conclusion

We examine two issues related to the profitability of trading on the information in earnings announcements: timing and liquidity costs. We find that assuming investors initiate positions at the close of the *Compustat* earnings announcement date, as sometimes appears in the literature, significantly overstates the size of the drift—even before liquidity costs. Assuming investors initiate their positions at the close of the day following the *Compustat* earnings report date somewhat understates the gross (before liquidity costs)

²² For SRW see Foster (1977) and for analysts see Mendenhall (1991).

²³ Our focus on the intersection of both SRW and analyst forecast errors leads to very few firms being included in our extreme portfolios in consecutive quarters—only about 2%. So, again, perhaps after transactions costs we have chosen a suboptimal strategy. Further, consistent with other comments in this section, investors might learn that, conditional on having a position in the stock, a less extreme earnings surprise is required to justify maintaining the position for another quarter.

magnitude of the drift relative to trading at the first actual closing price following the time of the earnings announcement. We show further that assuming investors wait to trade until the first actual close following the earnings announcement is extremely conservative.

After documenting an upper bound on the potential profits available from trading at bid-ask midpoints immediately following the time of the earnings announcement, we confirm prior research that indicates bid-ask spreads are larger during the announcement of extreme earnings surprises than at other times. We then present results for three scenarios estimating the timing of trades and the size of liquidity costs from our sample data. In our most conservative scenario, investors initiate their positions at the first market close following the earnings announcement, pay one-half of the stock's normal quoted bid-ask spread when initiating their positions, and again pay one-half the normal bid-ask spread when terminating their positions about three months later. This scenario assumes that: 87% of the time investors wait more than six trading hours to transact; investors pay a conservatively large estimate of liquidity costs when both initiating and terminating their positions; and investors portfolios exhibit 400% annual turnover. Despite these assumptions, drift returns for this scenario are statistically and economically greater than zero. Returns under this conservative scenario are essentially uncorrelated with surrogates for commissions, price concessions, and arbitrage risk. Taken together, our results leave little doubt that an investor could have earned abnormal returns of *at least* 14% per year by acting on the information in earnings surprises.

References

- Ball, Ray, 1978, Anomalies in relationships between securities' yields and yield-surrogates, *Journal of Financial Economics* 6, 103-126.
- Ball, Ray, 1992, The earnings-price anomaly, *Journal of Accounting and Economics* 15, 319-345.
- Ball, Ray, and Eli Bartov, 1996, How naive is the stock market's use of earnings information?, *Journal of Accounting and Economics* 21, 319-337.
- Ball, Ray, and Philip Brown, 1968, An empirical evaluation of accounting numbers, *Journal of Accounting Research* 6, 159-78.
- Bartov, Eli, Suresh Radhakrishnan, and Itzhak Krinsky, 2000, Investor sophistication and patterns in stock returns after earnings announcements, *The Accounting Review* 75, 46-63.
- Battalio, Robert and Richard R. Mendenhall, 2005, Earnings expectations, investor trade size, and anomalous returns around earnings announcements, *Journal of Financial Economics* 77, 289-319.
- Bennett, Paul and Li Wei, 2006, Market Structure, Fragmentation, and Market Quality, *Journal of Financial Markets* 9, 49-78.
- Berkman, Henk, and Cameron Truong, 2006, Event day zero? After-hours earnings announcements, available on SSRN.
- Bernard, Victor L., and Jacob K. Thomas, 1989, Post-earnings-announcement drift, delayed price response or risk premium? *Journal of Accounting Research* 27, 1-36.
- Bernard, Victor L., and Jacob K. Thomas, 1990, Evidence that stock prices do not fully reflect the implications of current earnings for future earnings, *Journal of Accounting and Economics* 13, 305-340.
- Bessembinder, Hendrick, 2003, Quote-based competition and trade execution costs in NYSE-listed stocks, *Journal of Financial Economics* 70, 385-422.
- Bhushan, Ravi, 1994, An informational efficiency perspective on the post-earnings announcement drift, *Journal of Accounting and Economics* 18, 45-65.
- Cao, Charles, Eric Ghysels, and Frank Hatheway, 2000, Price discovery without trading: Evidence from the Nasdaq preopening, *Journal of Finance* 55, 1339-1365.
- Chan, Louis K. C., Narasimhan Jegadeesh, and Josef Lakonishok, 1996, Momentum strategies, *Journal of Finance* 51, 1681-1713.
- Chen, Nai-Fu, Richard Roll, and Stephen Ross, 1986, Economic forces and the stock market, *Journal of Business* 59, 383-403.
- Collins, Daniel W. and Paul Hribar, 2000, Earnings-based and accrual-based market anomalies: one effect or two?, *Journal of Accounting and Economics* 29, 101-123.

Cushing, David and Ananth Madhavan, 2000, Stock Returns And Trading At The Close, *Journal of Financial Markets* 3, 45-67.

Fama, Eugene F., and James D. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy*, 81, 607-636.

Foster, George, 1977, Quarterly accounting data: Time-series properties and predictive-ability results, *The Accounting Review*, 52, 1-21.

Foster, George, Chris Olsen, and Terry Shevlin, 1984, Earnings releases, anomalies, and the behavior of security returns, *The Accounting Review* 59, 574-603.

Francis, Jennifer, Ryan LaFond, Per Olsson, and Katherine Schipper, 2003, Accounting anomalies and information uncertainty, Duke University working paper.

Huang, Roger and Hans Stoll, 1996, Dealer Versus Auction Markets: A Paired Comparison of Execution Costs on NASDAQ and the NYSE, *Journal of Financial Economics* 41, 313-358.

Jones, Charles M., 2002, A century of stock market liquidity and trading costs, Columbia University working paper.

Jones, Charles P., and Robert H. Litzenberger, 1970, Quarterly earnings reports and intermediate stock price trends, *Journal of Finance*, 25, 143-148.

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

Liang, Lihong, 2003, Post-earnings announcement drift and market participants' information processing biases, *Review of Accounting Studies* 8, 321-345.

Litzenberger, Robert H., O. Maurice Joy, and Charles P. Jones, 1971, Ordinal predictions and the selection of common stocks, *Journal of Financial and Quantitative Analysis*, 6, 1059-1068.

Livnat, Joshua, and Richard R. Mendenhall, 2006, Comparing post-earnings announcement drift for surprises calculated from analyst and time-series forecasts, *Journal of Accounting Research*, 44, 177-205.

Madhavan, Ananth, and Venkatesh Panchapagesan, 2000, Price discovery in auction markets: A look inside the black box, *Review of Financial Studies*, 13, 627-658.

Mc Inish, Tom and Robert Wood, 1992, An analysis of intraday patterns in bid/ask spreads for NYSE stocks, *Journal of Finance*, 47, 753-764.

Mendenhall, Richard R., 1991, Evidence on the possible underweighting of earnings information, *Journal of Accounting Research*, 29, 170-179.

Mendenhall, Richard R., 2004, Arbitrage risk and post-earnings-announcement drift, *Journal of Business*, 77, 875-894.

Patell, James, and Mark Wolfson, 1982, Good news, bad news, and the intraday timing of corporate disclosures, *The Accounting Review* 57, 509-527.

Petersen, Mark and Erik Sirri, 2003, Order preferencing and market quality on U.S. equity exchanges, *Review of Financial Studies* 16, 385-415.

Rendleman, Richard J., Charles P. Jones, and Henry A. Latané, 1982, Empirical anomalies based on unexpected earnings and the importance of risk adjustment, *Journal of Financial Economics* 10, 269-287.

Stoll, Hans and Christoph Schenzler, 2006, Trades outside the quotes: Reporting delay, trading option, or trade size?, *Journal of Financial Economics* 79, 615-653.

Stoll, Hans and Robert Whaley, 1983, Transactions costs and the small firm effect, *Journal of Financial Economics* 12, 57-80.

Figure 1

Trading assumptions

Scenario	Initiation	Termination
1.	The closing price on the day that the earnings surprise is reported by <i>Compustat</i> .	The closing price on the day that the subsequent earnings announcement is reported by <i>Compustat</i> .
2.	The closing price on the day after the date that the earnings surprise is reported by <i>Compustat</i> .	The closing price on the day that the subsequent earnings announcement is reported by <i>Compustat</i> .
3.	The closing price after the earnings surprise is reported by <i>Factiva</i> .	The closing price after the subsequent earnings announcement is reported by <i>Factiva</i> .
4.	The midpoint of the bid-ask spread prevailing at the time the earnings surprise is reported by <i>Factiva</i> . If the primary market is closed when the announcement is made, the midpoint of the bid-ask spread when the primary market resumes trading is used.	The closing price after the subsequent earnings announcement is reported by <i>Factiva</i> .
5.	For good- (bad-) news events, the ask (bid) price prevailing when the earnings surprise is first reported by <i>Factiva</i> . If the primary market is closed when the announcement is made, we use the relevant quote prevailing when the primary market resumes trading.	The closing price on the day that the subsequent earnings announcement is reported by <i>Factiva</i> .
6.	For good- (bad-) news events, the ask (bid) price prevailing when the earnings surprise is first reported by <i>Factiva</i> . If the primary market is closed when the announcement is made, we use the relevant quote prevailing when the primary market resumes trading.	For good- (bad-) news events, the closing price after the subsequent earnings announcement is reported by <i>Factiva</i> minus (plus) one-half of the normal bid-ask spread.
7.	For good- (bad-) news events, the closing price after the earnings surprise is reported by <i>Factiva</i> plus (minus) one-half of the normal bid-ask spread.	For good- (bad-) news events, the closing price after the subsequent earnings announcement is reported by <i>Factiva</i> minus (plus) one-half of the normal bid-ask spread.

Notes: Unless otherwise noted, if the announcement date is a non-trading day, we use the closing price on the subsequent trading day. Normal spreads are computed by dividing the average time-weighted bid-ask spread (TWS) in the prevailing in the twenty trading days beginning twenty five trading days before the day that the earnings surprise is reported by *Compustat* by the TAQ closing price on the day the earnings surprise is reported by *Compustat*.

Table 1

Differences in the dates on which earnings announcements appear in *Compustat* and the dates on which investors can trade on the announcements

There are 44 trades initiated two or more business days before the *Compustat* date (24 good news and 20 bad news) and 12 trades initiated two or more days after the *Compustat* date (5 good news and 7 bad news). The cumulative frequency of good news and bad news observations is reported in parentheses.

	Time Interval	Good News (N=2,637)	Bad News (N=3,595)
Trade initiated one business day before the <i>Compustat</i> date	9:30 to 10:00	242 (10.1%)	290 (8.6%)
	10:00 to 11:00	24 (11.0%)	19 (9.2%)
	11:00 to 12:00	15 (11.6%)	9 (9.4%)
	12:00 to 1:00	14 (12.1%)	6 (9.6%)
	1:00 to 2:00	9 (12.4%)	6 (9.7%)
	2:00 to 3:00	6 (12.7%)	10 (10.0%)
	3:00 to 4:00	3 (12.8%)	3 (10.1%)
Trade initiated on the <i>Compustat</i> Date	9:30 to 10:00	1,233 (59.5%)	1,596 (54.5%)
	10:00 to 11:00	89 (62.9%)	100 (57.3%)
	11:00 to 12:00	56 (65.0%)	53 (58.7%)
	12:00 to 1:00	48 (66.9%)	42 (59.9%)
	1:00 to 2:00	34 (68.1%)	43 (61.1%)
	2:00 to 3:00	38 (69.6%)	31 (62.0%)
	3:00 to 4:00	23 (70.5%)	31 (62.8%)
Trade initiated one business day after the <i>Compustat</i> date	9:30 to 10:00	763 (99.4%)	1,310 (99.3%)
	10:00 to 11:00	7 (99.7%)	16 (99.7%)
	11:00 to 12:00	2 (99.7%)	1 (99.7%)
	12:00 to 1:00	0 (99.7%)	2 (99.8%)
	1:00 to 2:00	0 (99.7%)	0 (99.8%)
	2:00 to 3:00	1 (99.8%)	0 (99.8%)
	3:00 to 4:00	1 (99.8%)	0 (99.8%)

Table 2

Mean subsequent-quarter size-adjusted returns by seasonal random walk (SRW) error quintile and analyst forecast error quintile.

Errors equal earnings per share (EPS) minus expected EPS, deflated by share price twenty days prior to the earnings announcement. Expected EPS for the SRW error is EPS for the same fiscal quarter of the prior year. Expected EPS for the analyst forecast error is the *I/B/E/S* mean forecast in the 90-day period before the earnings announcement date. Size-adjusted returns are cumulated from the close on the day following the *Compustat* earnings report date to the close on the *Compustat* earnings report date for the subsequent quarter.

			Analyst Forecast Error Quintile					
			Lowest (20%)	2	3	4	Highest (20%)	All
Time Series Forecast Error Quintile	Lowest (20%)	ret.	-0.023	-0.013	-0.007	0.005	0.010	-0.010
		t-stat	-5.89	-2.30	-0.92	0.83	1.61	-4.14
		N	8,190	2,749	1,282	2,024	3,469	17,714
	2	ret.	-0.023	-0.024	-0.003	0.001	0.018	-0.010
		t-stat	-5.22	-8.04	-0.73	0.26	2.75	-5.59
		N	3,887	5,410	3,203	2,996	2,122	17,616
	3	ret.	-0.007	-0.016	0.002	0.009	0.019	-0.001
		t-stat	-0.96	-5.38	0.89	2.59	1.94	-0.66
		N	1,221	5,228	5,916	3,494	1,215	17,074
	4	ret.	0.003	-0.013	-0.000	0.012	0.024	0.006
		t-stat	0.30	-3.33	-0.05	3.89	5.50	3.40
		N	1,152	3,130	3,874	5,594	3,278	17,026
	Highest (20%)	ret.	-0.009	-0.010	-0.012	0.017	0.054	0.022
		t-stat	-1.28	-1.15	-2.16	3.42	12.55	8.39
		N	2,539	2,293	1,757	3,389	7,395	17,373
	All	ret.	-0.018	-0.016	-0.002	0.010	0.033	0.001
		t-stat	-7.18	-8.63	-0.89	5.27	12.71	1.32
		N	16,989	18,810	16,032	17,497	17,479	86,807

Table 3

The impact of timing discrepancies on abnormal quarterly returns

Scenario 1 involves initiating the position at the closing price on the day that the earnings surprise is reported by *Compustat* and unwinding the position at the closing price on the day that the subsequent announcement is reported by *Compustat*. Scenario 2 involves initiating the position at the closing price on the day after the date that the earnings surprise is reported by *Compustat* and unwinding the position at the closing price on the day that the subsequent announcement is reported by *Compustat*. Scenario 3 involves initiating the position at the closing price after the earnings surprise is reported by *Factiva* and unwinding the position at the closing price after the subsequent earnings announcement is reported by *Factiva*. Scenario 4 involves initiating the position at the midpoint of the bid-ask spread prevailing at the time the earnings surprise is reported by *Factiva* and unwinding the position at the closing price after the subsequent earnings announcement is reported by *Factiva*. If the primary market is closed when the announcement is made, the midpoint of the bid-ask spread when the primary market resumes trading is used.

Panel A: Abnormal quarterly returns for good news announcements: 2,637 buy initiated round-trip trades.

	Scenario 1: Initiate at close of <i>Compustat</i> event day	Scenario 2: Initiate at close of <i>Compustat</i> event day+1	Scenario 3: Initiate at close of <i>Factiva</i> event day	Scenario 4: Initiate at quote midpoint after announcement
Mean	0.0880 (t = 8.52)	0.0701 (t = 7.53)	0.0730 (t = 7.54)	0.0855 (t = 8.21)
75th percentile	0.2104	0.1902	0.2015	0.2230
Median	0.0160	0.0113	0.0105	0.0196
25th Percentile	-0.1407	-0.1397	-0.1474	-0.1454
Std. Dev.	0.5299	0.4778	0.5114	0.5344
% positive	52.6	51.8	51.7	52.9

Panel B: Abnormal quarterly returns for bad news announcements: 3,595 sell initiated round-trip trades.

Mean	-0.0435 (t = -6.96)	-0.0236 (t = -3.68)	-0.0319 (t = -5.06)	-0.0455 (t = -7.17)
75th percentile	0.0983	0.1087	0.1028	0.0967
Median	-0.0729	-0.0582	-0.0627	-0.0783
25th Percentile	-0.2405	-0.2167	-0.2284	-0.2447
Std. Dev.	0.3746	0.3845	0.3785	0.4206
% positive	38.4	39.9	39.6	38.0

Panel C: Hedge Returns: Good News minus Bad News.

Mean	0.1315	0.0937	0.1049	0.1310
t-stat.	10.90	8.29	8.90	10.75
Time series t-stat.	9.58	8.34	8.99	10.73

Table 4

Historical and event-time relative bid-ask spreads and trading volume

For each earnings surprise in our sample, we compute the normal time-weighted bid-ask spread (TWS) by time-weighting the difference between the National Best Offer and the National Best Bid over the twenty trading days beginning five weeks before the day the trade is initiated. We compute 'Normal' spreads by dividing the normal TWS by the TAQ closing price on the *Compustat* date that the position is initiated. 'Actual' spreads are computed by taking the difference between the National Best Offer and the National Best Bid immediately following the *Factiva* announcement if the market is open and dividing the TAQ closing price on the *Compustat* date that the position is initiated. If the market is closed when the announcement is made, we use the first valid quotes when the market reopens. Normal trading volume is the average daily trading volume over the twenty trading days beginning five weeks before the day that the position is initiated. Actual trading volume is the volume on the day that the position is initiated.

Panel A: Relative bid-ask spreads.

	GOOD NEWS n = 2,637		BAD NEWS n = 3,595	
	Normal	Actual	Normal	Actual
Mean	2.69%	3.39%	3.63%	4.81%
75th percentile	3.48%	4.44%	4.79%	6.45%
Median	1.97%	2.26%	2.79%	3.45%
25th Percentile	1.01%	1.02%	1.50%	1.59%
Std. Dev.	2.52%	3.74%	3.12%	4.82%

Panel B: Trading volume (shares).

	GOOD NEWS n = 2,637		BAD NEWS n = 3,595	
	Normal	Actual	Normal	Actual
Mean	188,548	542,246	241,980	535,291
75th percentile	171,200	415,500	164,365	320,500
Median	59,475	140,300	58,410	90,900
25th Percentile	20,220	47,600	19,323	26,400
Std. Dev.	433,478	1,873,397	1,115,219	2,449,537

Table 5

The impact of bid-ask spreads on abnormal quarterly returns

Scenario 4 involves initiating the position at the midpoint of the bid-ask spread prevailing at the time the earnings surprise is reported by *Factiva* and unwinding the position at the closing price after the subsequent earnings announcement is reported by *Factiva*. For good- (bad-) news events, Scenario 5 involves initiating the position at the ask (bid) price prevailing when the earnings surprise is first reported by *Factiva* and unwinding the position at the closing price after the subsequent earnings announcement is reported by *Factiva*. For good- (bad-) news events, Scenario 6 involves initiating the position at the ask (bid) price prevailing when the earnings surprise is first reported by *Factiva* and unwinding the position at the closing price after the subsequent earnings announcement is reported by *Factiva* minus (plus) one-half of the normal bid-ask spread. For good- (bad-) news events, Scenario 7 involves initiating the position at the closing price after the earnings surprise is reported by *Factiva* plus (minus) one-half of the normal bid-ask spread and unwinding the position at the closing price after the subsequent earnings announcement is reported by *Factiva* minus (plus) one-half of the normal bid-ask spread.

Panel A: Abnormal quarterly returns for good news announcements: 2,637 buy initiated round-trip trades.

	Scenario 4: Initiate at announcement - pay no liquidity costs	Scenario 5: Initiate at announcement - pay ½ event spread	Scenario 6: Initiate at announcement - pay ½ event spread & ½ normal spread	Scenario 7: Initiate at close on announcement day - pay normal spread
Mean	0.0855 (t = 8.21)	0.0666 (t = 6.50)	0.0532 (t = 5.19)	0.0461 (t = 4.63)
75th percentile	0.2230	0.1992	0.1846	0.1781
Median	0.0196	0.0085	-0.0107	-0.0156
25th Percentile	-0.1454	-0.1615	-0.1754	-0.1733
Std. Dev.	0.5344	0.5263	0.52634	0.5117
% positive	52.9	50.1	49.9	47.4

Panel B: Abnormal quarterly returns for bad news announcements: 3,595 sell initiated round-trip trades.

Mean	-0.0455 (t = -7.17)	-0.0206 (t = -3.17)	-0.0025 (t = -0.38)	0.0043 (t = 0.68)
75th percentile	0.0967	0.1218	0.1413	0.1394
Median	-0.0783	-0.0570	-0.0396	-0.0282
25th Percentile	-0.2447	-0.2245	-0.2062	-0.1899
Std. Dev.	0.4206	0.3905	0.3913	0.3797
% positive	38.0	41.3	43.9	45.0

Panel C: Hedge Returns: Good News minus Bad News

Mean	0.1310	0.0872	0.0557	0.0418
t-stat.	10.75	7.18	4.58	3.54
Time series t-stat.	10.73	7.09	4.27	3.15

Table 6

Robustness of the quarterly hedge returns generated by paying normal liquidity costs to initiate and terminate positions at closing prices on *Factiva* event dates

For good- (bad-) news events, Scenario 7 involves initiating the position at the closing price after the earnings surprise is reported by *Factiva* plus (minus) one-half of the normal bid-ask spread and unwinding the position at the closing price after the subsequent earnings announcement is reported by *Factiva* minus (plus) one-half of the normal bid-ask spread.

	Mean quarterly hedge return	t-stat	Time series t-stat
Base case - Scenario 7 restated from Table 5	0.0418	3.54	3.15
Weight the returns generated by bad-news events by the bid (relative to other bad news events) and weight the returns generated by good-news events by the offer (relative to other good news events).	0.0367	3.89	3.20
Require the prevailing bid price when earnings is first announced to be greater than \$10.00.	0.0361	3.18	2.48
Omit observations in the bottom quartile on the basis of dollar trading volume on the earnings announcement day.	0.0471	3.43	3.66
Omit observations in the top quartile on the basis of arbitrage risk (idiosyncratic risk based on 250-day market model regression).	0.0454	4.01	4.09
Require the prevailing bid price when earnings is first announced to be greater than \$10.00 <i>and</i> winsorize 1% from each tail of the remaining sample return distribution.	0.0348	3.34	2.58
Omit observations in the bottom quartile on the basis of dollar trading volume on the earnings announcement day <i>and</i> winsorize 1% from each tail of the remaining sample return distribution.	0.0371	3.84	3.06
Omit observations in the top quartile on the basis of arbitrage risk <i>and</i> winsorize 1% from each tail of the remaining sample return distribution.	0.0376	4.53	3.79