KARL-FRANZENS-UNIVERSITÄT GRAZ
UNIVERSITY OF GRAZ
School of Business, Economics and Social Sciences

Earnings Autocorrelation and the Post-Earnings-Announcement Drift – Experimental Evidence

Josef Fink, Stefan Palan, Erik Theissen

Working Paper 2020-03
October 16, 2020

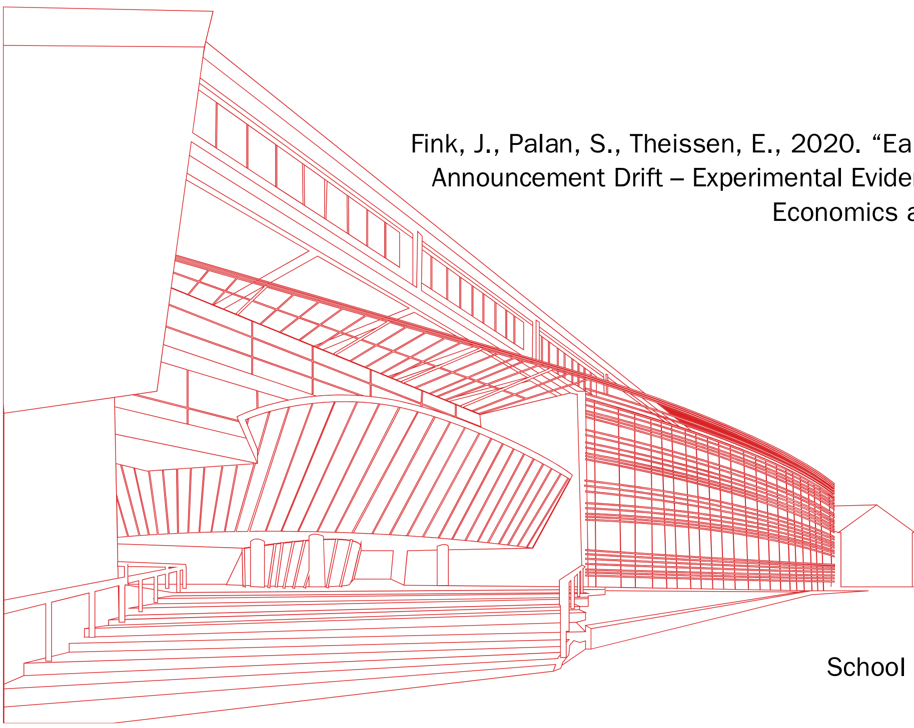

Recommended citation:
Fink, J., Palan, S., Theissen, E., 2020. "Earnings Autocorrelation and the Post-Earnings-Announcement Drift – Experimental Evidence", University of Graz, School of Business, Economics and Social Sciences Working Paper 2020-03.

Working Paper Series
School of Business, Economics and Social Sciences
University of Graz
ISSN 2304-7658
http://bit.ly/graz-working-papers
sowi-wp@uni-graz.at

UNIVERSITY OF GRAZ

Earnings Autocorrelation and the Post-Earnings-Announcement Drift – Experimental Evidence

Josef Fink
University of Graz, Institute of Banking and Finance

Stefan Palan
University of Graz, Institute of Banking and Finance
Corresponding author: +43(316)380-7306, stefan.palan@uni-graz.at

Erik Theissen
University of Mannheim, Finance Area
University of Graz, Institute of Banking and Finance

Working Paper 2020-03
October 16, 2020

Abstract

Post-earnings-announcement drift (PEAD) is one of the most solidly documented asset pricing anomalies. We use the controlled conditions of an experimental lab to investigate whether earnings autocorrelation is the driving cause of this anomaly. We observe PEAD in settings with uncorrelated and correlated earnings surprises, implying that earnings autocorrelation is not a necessary condition for PEAD. It rather is a moderator, as the PEAD is stronger when earnings surprises are serially correlated. We further show that market prices underadjust to fundamental value changes, and that trading strategies can profitably exploit the PEAD. Besides offering new results regarding the PEAD-phenomenon, we thus provide a proof-of-concept for the ability of experiments to generate valuable insights into this asset pricing anomaly.

Keywords: post-earnings-announcement drift, earnings autocorrelation, experimental asset markets

JEL: G12, G14, G40, M41

Working Paper Series
School of Business, Economics and Social Sciences
University of Graz
ISSN 2304-7658
http://bit.ly/graz-working-papers
sowi-wp@uni-graz.at

UNIVERSITY OF GRAZ

Earnings Autocorrelation and the Post-Earnings-Announcement Drift – Experimental Evidence*

Josef Fink[a], Stefan Palan[a,*], Erik Theissen[b,a]

[a]*Department of Banking and Finance, University of Graz, Universitätsstraße 15, 8010 Graz, Austria*
[b]*Finance Area, University of Mannheim, L9, 1-2, 68161 Mannheim, Germany*

Abstract

Post-earnings-announcement drift (PEAD) is one of the most solidly documented asset pricing anomalies. We use the controlled conditions of an experimental lab to investigate whether earnings autocorrelation is the driving cause of this anomaly. We observe PEAD in settings with uncorrelated and correlated earnings surprises, implying that earnings autocorrelation is not a necessary condition for PEAD. It rather is a moderator, as the PEAD is stronger when earnings surprises are serially correlated. We further show that market prices under-adjust to fundamental value changes, and that trading strategies can profitably exploit the PEAD. Besides offering new results regarding the PEAD-phenomenon, we thus provide a proof-of-concept for the ability of experiments to generate valuable insights into this asset pricing anomaly.

Keywords: post-earnings-announcement drift, earnings autocorrelation, experimental asset markets

JEL: G12, G14, G40, M41

*We thank Ryan Riordan and participants of the Austrian Working Group on Banking and Finance workshop 2019, the LFIN seminar at UC Louvain, and seminar participants at the University of Graz for valuable comments. We also thank Adele Theiler for research assistance. Financial support from the Austrian Science Fund FWF (project no. P32124-G27) is gratefully acknowledged. This research project was approved by the IRB of the University of Graz, case ID 39/57/63 ex 2016/17.

*Corresponding author.

Email addresses: josef.fink@edu.uni-graz.at (Josef Fink), stefan.palan@uni-graz.at (Stefan Palan), theissen@uni-mannheim.de (Erik Theissen)

1. Introduction

Share prices react to the surprise component of earnings announcements. Empirical research has, however, documented that the adjustment of prices is not instantaneous. After a positive [negative] earnings surprise, stock prices tend to drift upward [downward] for up to one year. This empirical phenomenon is referred to as the *post-earnings-announcement drift* (PEAD). The existence of PEAD is difficult to reconcile with the efficient markets hypothesis. Consequently, a large number of papers (to be briefly reviewed below) have proposed and tested various explanations for the phenomenon. However, no consensus has emerged so far. The balance of the available evidence is inconsistent with risk-based explanations, such that the PEAD is still considered a puzzle. Eugene Fama has dubbed it "an anomaly above suspicion" and referred to it as "the granddaddy of underreaction events" (Fama, 1998, 304 and 286, respectively). Daniel et al. (2020) go so far as to develop a behavioral three-factor model to explain the cross-section of stock returns and propose a PEAD factor to capture short-term mispricing.

In this paper we use experimental asset markets to study the emergence and the determinants of post-earnings-announcement drift. For our specific purposes, the experimental approach has a number of distinct advantages over empirical studies using field data. The experimenter can directly control, and deliberately manipulate, variables of interest. In particular, we carefully design the earnings process of the firms traded in our experimental markets. We thus know expected earnings with certainty and do not need to estimate them from time-series models or analyst forecasts. As a consequence, we observe earnings surprises without noise or bias. Furthermore, we can manipulate the earnings process and "switch on and off" serial correlation.

As noted above, the existing evidence on the PEAD is inconsistent with risk-based explanations. We therefore concentrate on mispricing as the driving force behind the drift. To this end we deliberately design our experimental markets such that there is no aggregate fundamental value risk. An important consequence of this design feature is that, in equilibrium, prices equal expected values. This has two benefits. First, we know how each earnings announcement affects the fundamental value of an asset. Second, we can analyze how quickly (if at all) prices approach fundamental values after earnings announcements. Another distinctive feature of our experimental design is that there is no aggregate information risk. We can thus rule out the explanation that returns to a PEAD-based trading strategy could be a compensation for information risk.

A further advantage of our laboratory setting is that we control at what time and which information is supplied to the traders in the experimental markets. We can thus ensure that information is symmetrically distributed, and that there is no information leakage. Finally, because we have a complete record of all order submissions, executions and cancellations, we can analyze whether the PEAD in our experimental markets can be profitably exploited after accounting for transaction costs. This is important, because the existing empirical literature on this issue is inconclusive.

We analyze two treatments. In the baseline treatment, it is common knowledge that earnings are serially uncorrelated. Earnings increase or decrease with equal probability, and by a constant amount. Even in this simple setting we observe a statistically significant PEAD. We can, therefore, conclude that earnings autocorrelation is not a necessary condition for a PEAD to arise. The earnings process is so simple that it is also very unlikely that a misspecified prediction model is the cause for the existence of the PEAD. In the correlated earnings treatment, earnings surprises are positively correlated. In this more complex setting we observe a stronger PEAD than in the baseline treatment, implying that earnings autocorrelation is a moderator of the PEAD. Autocorrelation does not cause the drift, but strengthens it. We also observe that more surprising earnings announcements are followed by more pronounced PEAD.

We also study whether there are trading strategies which, by conditioning on earnings news, allow investors to accumulate excess profits. Because of high bid-ask spreads, trading strategies that simply buy shares at the ask price and later sell them at the bid price (or vice versa) are not profitable. However, trading strategies that use limit orders to open and close a position after an earnings announcement earn positive and significant profits. We therefore conclude that the PEAD in our experimental markets can be profitably exploited on an after-transaction cost basis. Finally, we demonstrate that prices indeed underreact to the information content of earnings announcements. The subsequent post-earnings-announcement drift partly, but not fully, corrects the initial underreaction. The price adjustment is more complete in the correlated earnings treatment, a finding consistent with traders devoting more attention to analyzing the implications for asset values of earnings announcements in the more complex setting of this treatment. Finally, we show that the observation of greater mispricing following larger earnings surprises is indeed at least partially driven by underreaction to earnings autocorrelation.

Our findings contribute to the literature on the post-earnings-announcement drift in several important ways. Our result that there is a significant PEAD in a setting without aggregate risk adds to the evidence against risk-based explanations for PEAD. Prior to our work, no consensus had emerged in the literature on whether earnings correlation was a necessary condition for PEAD. We show that it is not, but that earnings autocorrelation affects the strength of the drift. There is also disagreement in the prior literature on the profitability of the PEAD on an after-transaction cost basis. We show that, in our experimental markets, trading strategies that profitably exploit the PEAD exist, but have to be based on limit orders to establish and potentially close a position. Finally, our finding that price changes better reflect changes in fundamental value in a more complex setting supports explanations of the PEAD based on investor (in)attention as proposed by, e.g., DellaVigna and Pollet (2009), Hirshleifer et al. (2009), Hou et al. (2009b) and Hung et al. (2015).

The remainder of the paper is organized as follows. Section 2 provides a brief summary of the relevant literature and develops the hypotheses that we test. Section 3 describes the experimental design. Section 4 presents the results. Section 5 concludes.

2. Literature and Hypotheses

2.1. Literature

Ball and Brown (1968) were the first to document that prices tend to adjust slowly to the information contained in earnings announcements. A large number of empirical studies have subsequently confirmed the existence of a PEAD and have also documented that the magnitude of the drift is systematically related to the strength of the earnings surprise.[1] The observation, made by Ke and Ramalingegowda (2005), that institutional investors exploit the PEAD, furthermore indicates that the phenomenon is also of practical relevance.

Several papers have advocated a risk-based explanation for the occurrence of PEAD (e.g., Ball et al., 1993, Kim and Kim, 2003). However, the changes in market betas observed around earnings announcements are insufficient to explain the PEAD (e.g., Bernard and Thomas, 1989). Similar findings obtain when multi-factor models are considered instead of the CAPM (e.g., Francis et al., 2007, Sadka, 2006, and Chordia et al., 2009). The majority view today thus is that PEAD is a mispricing phenomenon (e.g., Fama, 1998, Richardson et al., 2010, Hung et al., 2015, Daniel et al., 2020). We therefore concentrate on mispricing-based explanations and design our experiments such that confounds stemming from risk-based explanations can be ruled out. To this end we design markets with no aggregate risk, such that the equilibrium risk premium is zero.

Mispricing may arise in the context of earnings announcements when investors underreact to the information contained in the announcements. This underreaction may be driven by investors (1) using a misspecified model to foorecast earnings (e.g., Bernard and Thomas, 1989, Freeman and Tse, 1989, Bernard and Thomas, 1990), (2) being subject to behavioral

[1] See the surveys by Ball (1992), Bernard (1993), and Richardson et al. (2010).

biases (e.g., Frazzini, 2006), or (3) being inattentive (e.g., DellaVigna and Pollet, 2009, Hirshleifer et al., 2009, Hou et al., 2009a, Hung et al., 2015). The most prominent explanation for the occurrence of underreaction is based on the observation, made by Bernard and Thomas (1989, 1990), Rendleman et al. (1987) and others, that quarterly earnings are positively serially correlated. Investors fail to fully account for this autocorrelation and therefore underestimate the implications of the current announcement for future earnings. Consequently, the share price adjusts only partially to the information content of the announcement. Future news releases—even though they could have been anticipated—thus come as surprises to investors and trigger additional share price reactions in the same direction. In a related argument, Freeman and Tse (1989) suggest that investors may not know with certainty whether an earnings change is permanent or transitory. The next earnings announcement then provides information on the permanence of the preceding announcement.

Several studies, such as Ball and Bartov (1996), Rangan and Sloan (1998), Battalio and Mendenhall (2005), Bathke Jr et al. (2006), Bathke et al. (2014), and Chang et al. (2017) find empirical support for the earnings autocorrelation hypothesis. However, whether earnings autocorrelation is a necessary condition for the drift or rather a moderator which affects the strength of the drift but does not cause it remains unclear. One impediment to answering this important question is that, in empirical studies using field data, the existence of autocorrelation in (future) earnings can be estimated and forecasted only from historical earnings data. Such forecasts are subject to model risk and noisy data. In our experimental markets, in contrast, we have full control over the earnings process and thus can "switch on and off" earnings autocorrelation. We can therefore provide a direct test of whether earnings autocorrelation *causes* the drift or whether it only amplifies a drift that would also occur if earnings were serially uncorrelated.

An important question is whether the PEAD can be profitably exploited or whether transaction costs outweigh any profits that could be obtained. The available evidence on this issue is mixed. Chordia et al. (2009), Ng et al. (2008), Zhang and Zhang (2013), and Pavlova and Parhizgari (2011) find that abnormal profits essentially disappear after accounting for transaction costs, while Ke and Ramalingegowda (2005) and Battalio and Mendenhall (2011) find evidence of significant profits even after transaction costs. We use the data from our experimental markets to reexamine the question of whether trading strategies can yield significant profits after accounting for transaction costs.

Ours is the first paper to use financial market experiments to analyze the post-earnings-announcement drift. It is thus also the first experimental paper to study the effect of earnings autocorrelation on the drift. However, we are aware of three papers that have used laboratory experiments to analyze whether and how experimental subjects (mis)interpret the information content of earnings announcements. Maines and Hand (1996), Calegari and Fargher (1997), and Bloomfield et al. (2003) conduct experiments in which subjects forecast future earnings. All three papers agree on the conclusion that subjects underestimate the importance of earnings autocorrelation, possibly because they overweight past information and underweight the contemporaneous information provided by the current earnings announcement (Bloomfield et al., 2003). None of the three papers offers insights into PEAD, however. The experiments of Maines and Hand (1996) and Bloomfield et al. (2003) do not feature a stock market, and Calegari and Fargher (1997) has a single call auction after each earnings announcement, precluding the possibility of observing a drift.

2.2. Hypotheses

The empirical literature reviewed above provides strong evidence that a post-earnings-announcement drift exists. We therefore expect that a PEAD will also arise in our experimental market data.

Hypothesis 1. *Post-earnings-announcement drift can be observed in experimental asset markets.*

The majority view in the literature is that the PEAD is a mispricing phenomenon, frequently assumed to be caused by underreaction to the news contained in a given earnings announcement. The most popular explanation for this underreaction posits that consecutive earnings surprises are correlated, and that investors underestimate the implications of any given announcement for future earnings (Bernard and Thomas, 1989, 1990). What remains open, however, is whether earnings autocorrelation is a necessary condition for PEAD, or whether earnings autocorrelation only amplifies a drift that would also occur in its absence. We test the following hypothesis to shed light on this issue.

Hypothesis 2. *Earnings autocorrelation is a necessary condition for post-earnings-announcement drift.*

The available empirical evidence shows that the magnitude of the PEAD depends on the size of the earnings surprise. In our correlated earnings treatment there are two types of announcements which trigger very different changes in the fundamental value of the asset. We use this design feature to test the following hypothesis.

Hypothesis 3. *Greater earnings surprises are followed by greater drift.*

To determine whether the PEAD is only a statistically significant phenomenon or is also economically relevant, its magnitude needs to be judged against the size of the transaction costs. As noted above, prior evidence from empirical studies using field data is inconclusive in this respect. We therefore test whether the drift in our idealized markets is strong enough to be profitably exploited. Specifically, we test the following hypothesis.

Hypothesis 4. *The observed post-earnings-announcement drift can be profitably exploited even after accounting for transaction costs.*

The PEAD is generally thought to be an underreaction phenomenon. The reasoning goes as follows: The earnings surprise contains information about a change in the fundamental value of the stock. The stock price only partially adjusts to the new value and then continues to drift in the direction of the earnings surprise. While this is a plausible hypothesis, it is difficult to test with field data because the fundamental value of a stock, and the change in this fundamental value revealed by the earnings surprise, are hard to measure precisely. In our experiments, we know the fundamental value of the stocks traded in our laboratory markets, and we also know how this value changes upon an earnings announcement. We can therefore perform a direct test of the following hypothesis.

Hypothesis 5. *Prices initially underadjust to the information content of the earnings announcement and then continue to drift in the direction of the earnings surprise until price equals fundamental value.*

When earnings changes are positively serially correlated there are non-surprising announcements (e.g. an earnings increase following a previous earnings increase) and surprising announcements (e.g. an earnings decrease following an increase). The surprising earnings changes trigger a larger change in the fundamental value of the asset (a decrease after an increase predicts that future earnings changes are also more likely to be negative). If investors fail to fully account for earnings autocorrelation (as hypothesized by, among others, Bernard and Thomas (1989)), they will overestimate the implications for the fundamental value of the asset of non-surprising earnings announcements. This, in turn, will alleviate the tendency for prices to only partially adjust to changes in fundamental value. On the other hand, investors will underestimate the implications for the asset value of surprising announcements. This, in turn, will reinforce the tendency for prices to only partially adjust to changes in fundamental value. We thus have the following hypothesis.

Hypothesis 6. *In the treatment with correlated earnings, prices adjust more fully to the information content of the earnings announcement after non-surprising than after surprising announcements.*

3. Experimental design

Our experiment was run in May and November of 2019 in the experimental research laboratories of four large research and teaching universities, using bachelor, master and

PhD students of Economics and Business as subjects.[2] The experiment consists of two experimental designs, referred to as 'treatments'. We name the two treatments Base and Corr. Each of the 20 sessions we run employs exactly one of our two treatments (10 Base, 10 Corr), and every one of our 238 subjects is thus exposed to one treatment only ('between-subjects' design). The subjects are compensated for their participation by cash payments tied to their performance, paid at the end of the experiment. In each experimental session, eleven to twelve subjects form a cohort and interact over a sequence of four independent periods.[3]

Each session is structured as follows: After the experimenter has checked subjects' IDs and welcomed them to the lab, they are randomly assigned to computers. Following the best practices laid out in Freeman et al. (2018), we report that all subjects then receive the same written instructions, providing information on the trading interface. The experimenter reads the instructions out aloud while subjects follow along to create common knowledge of their contents (i.e., to ensure subjects know the contents, know that all other subjects also learnt the same contents, etc.). Afterwards, the trading mechanism and the most important screens are explained in detail, followed by a trial period to allow subjects to familiarize themselves with the trading interface. The trial period is followed by further instructions (delivered in the same manner as before) describing the earnings announcements and their relevance for the values of the stocks. Finally, subjects answer control questions to ensure their understanding of the instructions before the first trading period commences. After the fourth period of trading has ended, one period is randomly chosen for payout. Subjects then

[2]The combined requirements of using only students of economics, business and related programs (financial mathematics, information technology for business, business education) and running a relatively large number of sessions with 11-12 participants each, forced us to conduct experiments in more than one experimental lab. A positive side effect of this procedural feature is that we can report that our results are robust not only within a certain lab, but also across multiple labs, at different universities, located in two different countries.

[3]In sessions 13 and 14, a higher than expected number of registered subjects did not show up. Therefore, these sessions had only 11 instead of the intended 12 subjects.

complete a post-experiment questionnaire, are paid in private and in cash, and leave.

3.1. Trading environment

In all trading periods, subjects trade shares of two fictitious companies, firm A and firm B. The trading mechanism is a continuous double auction with open order books, implemented in a modified version of GIMS v7.4.11 (Palan, 2015), running on z-Tree v4.1.7 (Fischbacher, 2007). A screenshot of the trading interface is provided in Figure A.1 in the appendix. Each subject is endowed with 900 talers (the experimental currency) and 9 shares of either the stock of firm A (stock A hereafter) or the stock of firm B (stock B)—but not both—at the beginning of every period.[4] Traders can submit any combination of limit and market orders in the markets for stocks A and B.[5] Each order is for one share. The order book is empty at the beginning of a period, and it is anonymous, i.e., the identity of the trader submitting an order is not displayed. Order execution is governed by price and time priority, following the same algorithm as applied at NASDAQ. Unexecuted orders can be cancelled at no cost, the risk-free rate of interest on taler holdings is set to zero, and there are no transaction costs beyond the bid-ask spread that emerges endogenously. Traders can sell short up to a short-selling capacity of 9 A and 9 B shares. Similarly, traders can buy on margin, with negative cash balances of up to 900 talers.

At the beginning of every period, subjects are informed about the earnings per share of both firms. During the trading period, there are four earnings announcements where subjects receive updated information about these earnings per share. After the end of the trading period, there is a fifth announcement. Subjects can trade continuously throughout five 'phases' of equal length (180s) in each period. Phase 0 is the interval of trading start-

[4] In the markets with only 11 subjects, five were endowed with 9 shares of firm A and 900 talers, five were endowed with 9 shares of firm B and 900 talers, and one was endowed with five shares each of firm A and firm B, and with 1000 talers. This ensured an absence of risk at the market-level even in these slightly understaffed sessions.

[5] We will use the terms 'subject' and 'trader' interchangeably.

ing with the beginning of the period and leading up to announcement 1, Phase 1 follows announcement 1, etc., as illustrated in Figure 1.

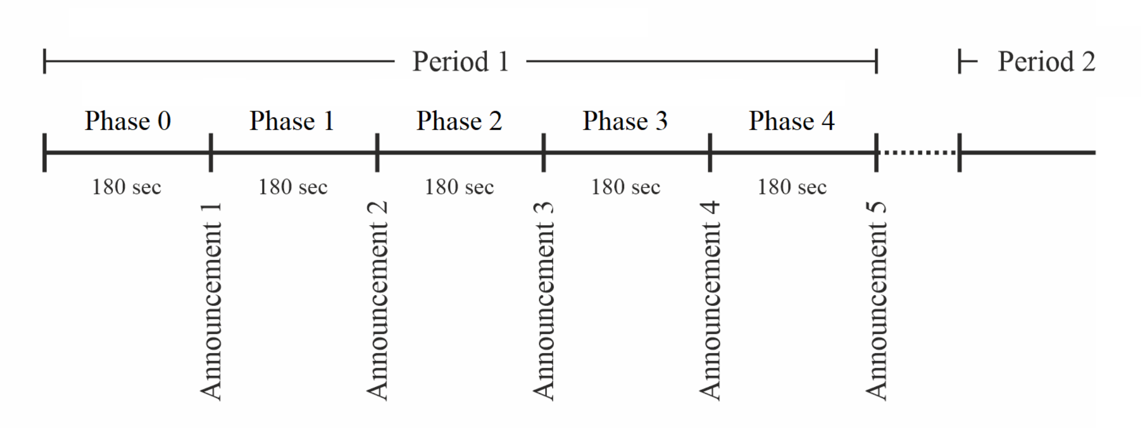

Figure 1: Structure of a trading period. Each of the four 900s trading periods in a session is structured into five 180s phases, separated by four inter-phase earnings announcements, and concluded by one closing earnings announcement.

3.2. Earnings announcements

Each period lasts 900s, with earnings announcements taking place after 180s, 360s, 540s and 720s, and with a final announcement after the period has ended (i.e., after 900s). The shares are bought back by the experimenter after the end of the period for the fundamental value (FV) of 20 times earnings after the fifth announcement, mimicking perpetual discounting at a rate of 5%.

The trading interface displays a countdown to the upcoming earnings announcement (see Figure A.1). At the time of an earnings announcement, the updated earnings of the two firms are shown on screen (and highlighted by a flashing red box) while trading continues uninterruptedly. The earnings per share for both firms are 5 talers at the beginning of each period. In each announcement, the earnings for a firm can increase or decrease by 0.5 talers:

$$\Delta E_\tau^\theta \in \left\{\delta^- = -0.5, \delta^+ = 0.5\right\}, \tag{1}$$

where ΔE_τ^θ is the change, in treatment $\theta \in \{\textsc{Base}, \textsc{Corr}\}$, in earnings published in announcement τ. The constants δ^- and δ^+ hold the two possible earnings changes (in talers).

Our two treatments differ in the dependence structure of successive earnings changes. In treatment Base, positive and negative earnings changes are equally likely:

$$p\left(\Delta E_\tau^{\textsc{Base}} = \delta^-\right) = p\left(\Delta E_\tau^{\textsc{Base}} = \delta^+\right) = 0.5 \;\forall\, \tau, \tag{2}$$

where p is the probability operator. Furthermore, the signs of successive earnings changes are independent. This structure implies that the expected change in earnings is zero and that, therefore, the earnings surprise (i.e., the unexpected component of the announcement) equals the earnings change. Furthermore, the serial correlation of earnings surprises is 0. The earnings process in treatment Base implies that earnings as well as FV follow a recombining binomial tree over the course of a period. This tree is illustrated in Figure 2.

Our second treatment, Corr, models firms whose earnings surprises are autocorrelated. We accomplish this by increasing the probability that successive changes in earnings have the same sign. Specifically, a change in earnings is followed by another change in the same direction with probability 0.75 and is followed by a change in the opposite direction with probability 0.25. Formally,

$$\begin{aligned} p\left(\Delta E_\tau^{\textsc{Corr}} = \delta^- \middle|\, \Delta E_{\tau-1}^{\textsc{Corr}} = \delta^-\right) = p\left(\Delta E_\tau^{\textsc{Corr}} = \delta^+ \middle|\, E_{\tau-1}^{\textsc{Corr}} = \delta^+\right) = 0.75, \\ p\left(\Delta E^{\textsc{Corr}}\tau = \delta^- \middle|\, E_{\tau-1}^{\textsc{Corr}} = \delta^+\right) = p\left(\Delta E_\tau^{\textsc{Corr}} = \delta^+ \middle|\, E_{\tau-1}^{\textsc{Corr}} = \delta^-\right) = 0.25 \end{aligned} \tag{3}$$

The autocorrelation of firms' earnings is thus 0.5. The earnings process in treatment Corr implies that earnings as well as FV follow a non-recombining binomial tree. This tree is illustrated in Figure 3.

In treatment Base each announcement only reveals information about current earnings. In

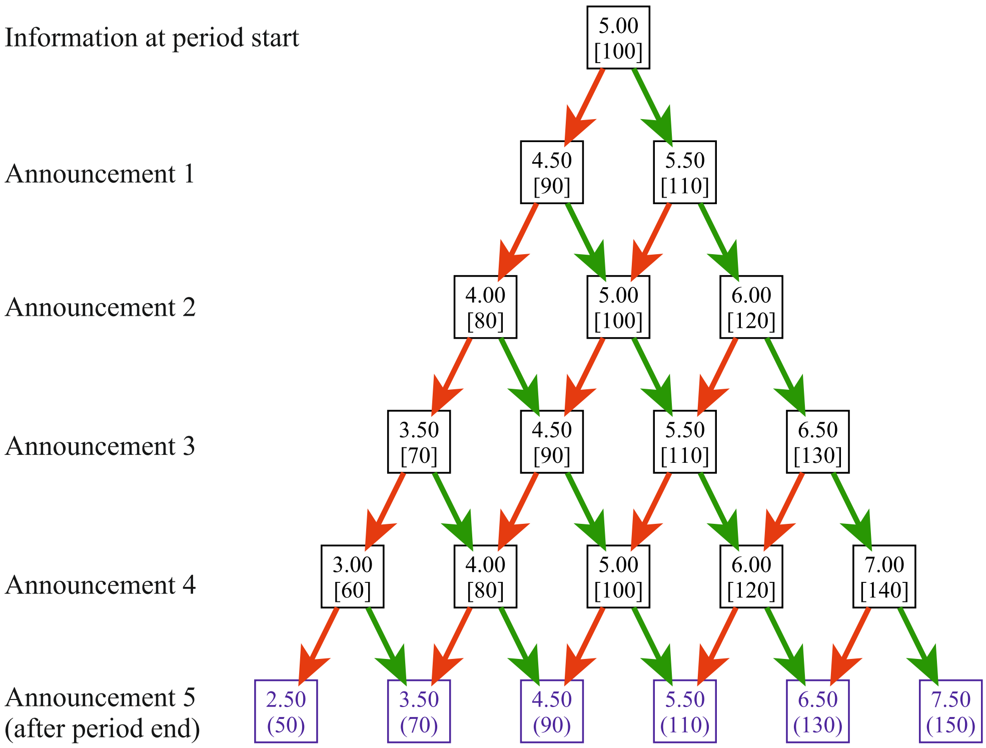

Figure 2: Illustration of the possible earnings trajectories within a period in treatment Base. The first number in each box is the current level of earnings. The amounts in parentheses and square brackets are the respective fundamental values, i.e., the expected future payoffs per share corresponding to the current level of earnings (figures in parentheses were communicated to subjects, figures in square brackets were not). Red (Green) arrows lead to boxes following decreases (increases) in earnings. The blue boxes (bottom row) are only reached after trading for that period has concluded.

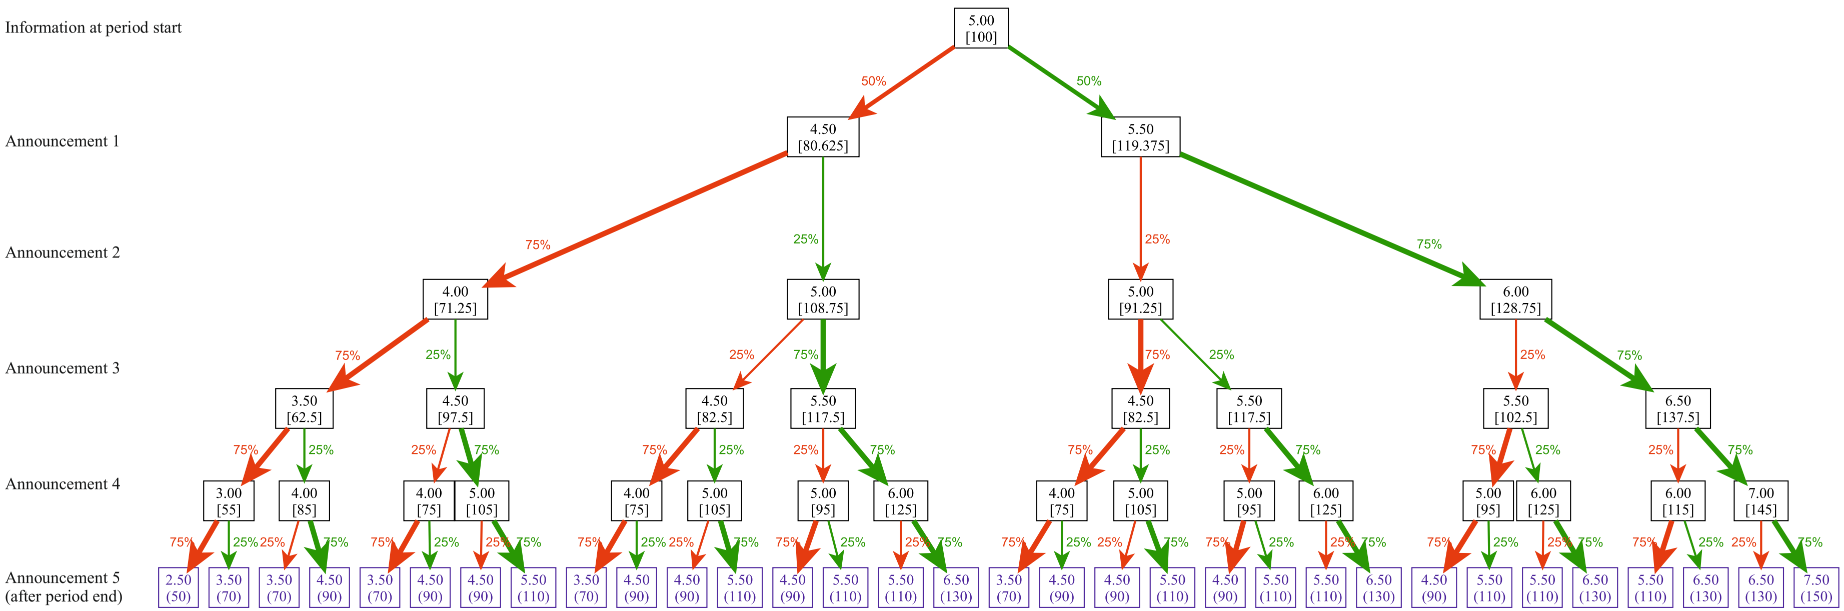

Figure 3: Illustration of the possible earnings trajectories within a period in treatment CORR. The first number in each box is the current level of earnings. The amounts in parentheses and square brackets are the respective fundamental values, i.e., the expected future payoffs per share corresponding to the current level of earnings (figures in parentheses were communicated to subjects, figures in square brackets were not). Red (Green) arrows lead to boxes following decreases (increases) in earnings. Bold arrows indicate a high probability of events unfolding along this path (noted next to the arrows in the same color). The blue boxes (bottom row) are only reached after trading for that period has concluded.

contrast, each announcement in treatment Corr contains information about both current earnings and subsequent earnings announcements. Consequently, surprising announcements in treatment Corr have greater impact on the fundamental value than the announcements in treatment Base, while unsurprising announcements in Corr have a smaller impact on the fundamental value.

In any given announcement, if firm A's earnings are announced to have increased, firm B's earnings are announced to have decreased, and vice versa. The earnings changes always have the same magnitude (0.5 talers) but opposite signs. The earnings changes of the two firms are thus perfectly negatively correlated, and so are the changes in fundamental values (FV). All of this is public information.

The perfectly negative correlation of the changes in fundamental values implies that a portfolio containing an equal number of A and of B shares is risk-free. Remember that each subject is endowed with only A shares *or* only B shares. Risk-averse traders thus have an incentive to trade in order to equate their holdings of A and B shares. Because the aggregate number of A shares equals that of B shares, the aggregate risk is zero. Equilibrium with risk-averse traders then implies that (1) all traders hold risk-free portfolios, and (2) the risk premium is zero. Consequently, in equilibrium, prices should equal expected values.

The negative correlation of the fundamental value changes also implies that the sum of the values of an A share and of a B share is always 200. In a market with N subjects, the aggregate endowment is $\frac{N}{2} \cdot 9$ shares of each type. The total value of the market's share endowment thus is $\frac{N}{2} \cdot 9 \cdot 200 = 900N$. Since the market's total cash endowment is also $900N$, the cash-to-asset ratio equals 1, thus enabling trading while avoiding cash endowment effects.[6]

[6]See Palan (2013) and Noussair and Tucker (2016) and the references therein for evidence on the relationship between the cash/asset ratio and mispricing. Note that in our experiments the short selling capacity (each trader can short up to 9 A and 9 B shares) and the margin buying capacity (each trader can buy on margin up to a negative cash balance of 900 talers) are also symmetric and thus do not distort the cash/asset ratio.

3.3. Subject payment

Subjects' payoffs are based on their wealth $W_{p,T}$, determined at the end of each period. Specifically, one period is chosen for payout using a physical randomization device and each subject's wealth in this period is calculated as:

$$W_i = \sum_{f \in \{A,B\}} 20 E_f n_{i,f} + c_i, \tag{4}$$

where E_f are the final earnings of firm $f \in \{A, B\}$, $n_{i,f}$ is subject i's final balance of shares of firm f's stock, and c_i are the subject's final cash (taler) holdings.[7]

Payoffs are calculated by converting W_i using an exchange rate of 100 talers = €1. Additionally, subjects receive a base compensation of €5 to €8, depending on the rules of the lab a given session was run in.[8] Overall earnings are bounded from below by €0. In total, subjects thus earn an average of 1800 talers (€18) from trading plus the base compensation for an experiment lasting around 2h. Table 1 reports the actual payments from the experiment. The average subject earns about €25.

Table 1: Overview of subject payoff in euros. Information about subject payoffs by treatment, including mean payoff, payoff standard deviation (SD, within-session average), minimum (Min) and maximum (Max) payoff.

Treatment	Mean	SD	Min	Max
Base	24.8	2.6	11.4	30.8
Corr	25.0	5.0	9.1	47.1

4. Results

Before we turn to a test of our hypotheses we provide a short description of the trading activity in our experimental markets.[9] As we document in Table 2, we see substantial trading

[7] We suppress the p and τ indices of W, n, E and c for notational simplicity.

[8] Labs in different cities have differing requirements regarding average subject compensation. These different requirements usually reflect differences in wage and price levels between the cities.

[9] We analyze our data using *R* (R Core Team, 2017) and generate regression tables using *stargazer* (Hlavac, 2018) and *texreg* (Leifeld, 2013). For the remaining tables we use *kableExtra* (Zhu, 2019), for the

activity. The average trader engages in 48.6 actions (submitting a limit order, cancelling an order, or accepting another trader's order) over the course of a period and is involved in 26.3 transactions per period. These numbers correspond to one action per trader every 18.5s and one transaction per trader every 34.2s. The figures in Table 2 also reveal that there is more trading activity in the Base treatment than in the Corr treatment. Both the number of transactions and the number of actions per period are higher by approximately 15% in Base, most likely because of the less complex setting in this treatment.

Table 2: Overview of trading activity. 'Actions' include order submissions, order cancellations, and the acceptance of outstanding orders.

	Trades per period				Actions per period			
Treatment	Mean	SD	Min	Max	Mean	SD	Min	Max
Base	179	74	89	468	630	187	349	1159
Corr	170	57	86	318	579	148	352	870

4.1. Existence of PEAD

In order to analyze the prices in our experimental markets we subdivide each period into intervals of 10s. We refer to these intervals as *windows*. As we will demonstrate in section 4.3, the bid-ask spreads in the experimental markets are large and particularly so shortly after an earnings announcement. We therefore base our main analysis on quote midpoints (but still refer to them as "prices" for ease of exposition) in order to eliminate the effect of bid-ask bounce. The analysis thus—and in line with most prior empirical research on the PEAD—ignores transaction costs. Also in section 4.3 we then analyze the profitability of trading strategies aimed at exploiting the PEAD. This analysis is based on executable bid and ask prices and thus explicitly accounts for the execution costs traders incur.

Figure 4 plots cumulative taler price changes relative to the quote midpoint at the time of the earnings announcement. The upper [lower] line in each diagram tracks price changes

figures *ggplot2* (Wickham, 2016).

after positive [negative] earnings surprises. The dashed horizontal lines refer to the price levels reached by the end of the announcement window, i.e., the window starting at the moment of the earnings announcement. These prices capture the initial price reaction to the earnings announcement.

Panel (a) of Figure 4 shows results pooled across both treatments. There is clear evidence of PEAD. Prices of stocks with positive earnings news jump by 3.96 talers in the announcement window, but then continue to drift upwards by another 3.80 talers in the 170s following the announcement window. The drift is even more pronounced after negative earnings surprises. In this case prices drop by 2.10 talers in the announcement window and then drift downwards by another 7.88 talers over the remainder of the phase. These figures imply that a strategy, initiated at the end of the announcement window, that goes long stocks with positive earnings surprises and short stocks with negative surprises earns a profit (before transaction costs) of 11.71 taler.

Table 3 reports the results of t-tests.[10] The figures in the first column indicate that the profits of both a 'long' strategy that invests in stocks with positive announcements and a 'short' strategy that sells the stocks with negative announcements earn a statistically significant profit, with t-statistics of 6.35 and 13.61, respectively. The profit of the combined 'long-short strategy' is also highly significant, with a t-statistic of 14.48. These results clearly support Hypothesis 1.

Result Hypothesis 1. *There is clear evidence of post-earnings-announcement drift in our experimental asset market data.*

[10]Following Benjamin et al. (2018), we highlight significance at the 0.005 level alongside the more conventional 0.05 and 0.01 levels.

Figure 4: Taler gains and losses relative to the quote midpoint at the time of the announcement. Average price changes relative to the quote midpoint at the time of the announcement, using the closing quote midpoint for each 10s window following the announcement. Panel (a) plots results for the pooled data from all treatments; panel (b) reports results separately for treatments BASE (left) and CORR (right); panel (c) plots only CORR data and reports results separately for unsurprising (left) and surprising (right) earnings news, using only data from Phases 2 through 4. The blue (orange) line plots the cumulative price changes following positive (negative) earnings news. The dashed horizontal lines of the same colors indicate the price levels at the end of the 10s announcement window.

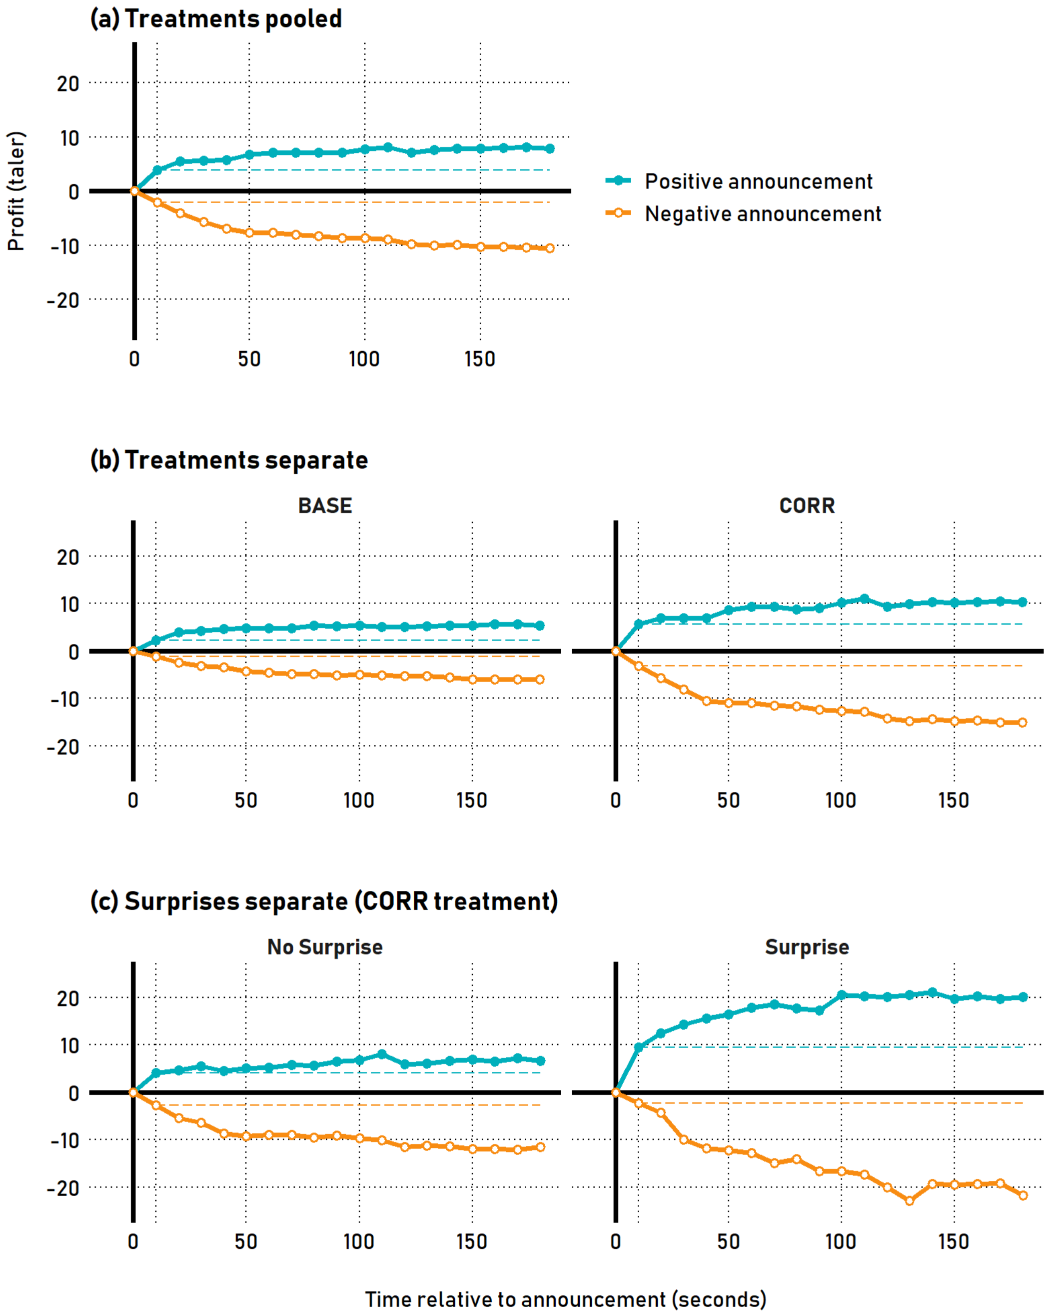

Table 3: Post-earnings-announcement drift. Mean taler gains based on quote midpoints from the end of the post-announcement-window until the end of the phase. The data for treatment Corr separated by the surprise variable exclude the phase following the first announcement. We report t-statistics in parentheses.

	Treatments				CORR Treatment[a]		
	All[b]	BASE	CORR	Δ	No Surprise	Surprise	Δ
Long	3.80***	3.17***	4.42***	1.25	2.75*	10.22***	7.47**
	(6.35)	(6.08)	(4.12)	(1.05)	(2.25)	(4.36)	(2.82)
Short	-7.88***	-4.68***	-11.36***	-6.69***	-8.71***	-18.42***	-9.71**
	(-13.61)	(-9.41)	(-11.37)	(-5.99)	(-8.54)	(-5.55)	(-2.80)
Long-Short	11.71***	7.66***	16.20***	8.53***	11.94***	26.09***	14.15***
	(14.48)	(10.08)	(11.72)	(5.41)	(6.78)	(7.56)	(3.65)

Note:
*p<0.05, **p<0.01, ***p<0.005 (two-tailed t-test)
Values based on all phases where both the first and the last 10s windows following an announcement (windows 0 and 17) have both a bid and an ask; thus permitting us to calculate quote midpoints.
[a] CORR Treatment data separated by the surprise variable excludes the first announcement.
[b] Only phases with valid midpoints to calculate return considered (includes 95% , 87.5% , and 85% of Long, Short, and Long-Short phases, respectively, for 'All' column).

4.2. PEAD and earnings autocorrelation

To test whether earnings autocorrelation is a necessary condition for a PEAD to arise we next analyse the experimental markets separately for each of the two treatments. We present the results in Panel (b) of Figure 4. There is clear evidence of PEAD in both treatments. However, the drift is much more pronounced in the Corr treatment than in the Base treatment. The results shown in columns 2 and 3 of Table 3 confirm this finding. After a positive earnings surprise and excluding the announcement window, prices drift upwards by 3.17 talers in treatment Base and by 4.42 talers in treatment Corr. Both values are highly significant. Similarly, prices drift downwards after negative earnings surprises by 4.68 talers in treatment Base and by 11.36 talers in treatment Corr.[11] Again both values are highly significant. The long-short strategy yields a return of 7.66 talers in treatment Base and of 16.20 talers in treatment Corr. The differences between the two treatments are significant

[11] We test whether the immediate price reactions after positive and negative earnings surprises differed significantly. We find significant differences in the Base treatment (Welch two-sample $t(214.08) = 1.0850$, $p = 0.0312$) and somewhat stronger differences in the Corr treatment (Welch two-sample $t(289.97) = 2.5758$, $p = 0.0022$).

for the short strategy and for the combined long-short strategy.

We run regressions to gain deeper insights into the price dynamics, using the window-to-window taler quote midpoint price changes as our dependent variable. We exclude the announcement window from the analysis, because it captures the initial price reaction to the earnings announcement rather than the post-announcement drift. All price changes after negative earnings surprises are multiplied by (-1). This allows us to pool the data from windows following positive and negative announcements. Our independent variables include a dummy that identifies observations from the CORR treatment, and count variables for the period within a session, the phase within a period, and the window within a phase.[12] The period and phase count variables capture changes in the strength of the PEAD across the periods of a session and the phases following the four earnings announcements within a period. We use their coefficients to judge whether the PEAD diminishes as the experimental subjects gain more experience. We use the coefficient of the window count variable to analyze the dynamics of the drift following an earnings announcement. Because the pattern depicted in Figure 4 suggests that the strength of the drift decreases at a decreasing rate, we also include the square of the window number to allow for a non-linear change in the strength of the drift. Finally, we include a dummy variable that identifies phases following a positive earnings surprise. We can thus test whether the drift is different following positive and following negative surprises.

We report the results in the first column of Table 4. The constant is positive and significantly different from zero, implying that there is a drift in the direction of the earnings announcement in the BASE treatment.[13] The drift is weaker after positive announcements, as

[12]Remember that each session consists of four periods and each period is sub-divided into five phases of 180s each. Consequently, each phase consists of 18 windows of 10s each. The regression analysis excludes the first phase of each period (because it is not preceded by an earnings announcement) and the first window of each phase (because it captures the initial price reaction to an announcement rather than post-announcement drift).

[13]Note that we re-base the Period, Phase and Window variables (i.e. we set the respective count variable

is evidenced by the negative coefficient for the 'Positive earnings change' dummy. However, the sum of the two coefficients is still positive and significant, implying that there is also a significant PEAD after positive announcements in the BASE treatment.

The positive and significant coefficient on the CORR treatment dummy implies that the drift is stronger in the presence of earnings autocorrelation. The pattern of the 'Window0' and 'Window0^2' coefficients suggests that the drift indeed decreases over time, but at a decreasing rate. The coefficients on the count variables for the period and the phase are small and not significant, implying that the drift is not lower in later phases of a period or in later periods of a session. We therefore conclude that the occurrence of PEAD is a persistent phenomenon in our experimental markets.[14]

The results presented in this sub-section support our earlier finding that there is a significant and persistent PEAD in our experimental markets. They furthermore clearly reject hypothesis 2. Earnings autocorrelation is not a necessary condition for the occurrence of PEAD. Rather, it is an accelerant that strengthens PEAD compared to a situation without earnings autocorrelation.

Result Hypothesis 2. *Earnings autocorrelation is not a necessary condition for post-earnings-announcement drift. However, the drift is significantly more pronounced in the presence of earnings autocorrelation.*

The CORR treatment is characterized by positively autocorrelated earnings changes. This pattern in earnings announcements implies that an earnings change with the same sign as the previous change is more likely (and thus less surprising) and has a smaller impact on the fundamental value of the stock than a change in the opposite direction. We will

for the first Period, Phase and Window to 0 rather than to 1) in order to make the constant interpretable. With these variable definitions the constant in our regression captures the drift in the first phase, of the first period, immediately after the announcement window, after a negative announcement, in the BASE treatment. We identify the re-based variables by adding a "0" to the variable name, e.g. 'Period0'.

[14] Note that the regression R^2 is low. Thus, while the large number of observations in our sample allows us to document clear evidence regarding drivers and dynamics of the return drift following earnings announcements, the patterns that we detect explain only a small fraction of the overall return variability.

refer to the two types of announcements as "unsurprising" and "surprising" announcements, respectively. Note that the first announcement in treatment Corr is neither unambiguously surprising nor unsurprising.[15] We therefore exclude the first announcement from all analyses that distinguish between surprising and unsurprising announcements.

We plot results for surprising and unsurprising announcements in Panel (c) of Figure 4. The figure provides clear evidence that the PEAD is more pronounced after surprising announcements. This is in line with the empirical literature that has documented stronger drift after larger earnings surprises (e.g., Bernard and Thomas, 1989). Our previous finding that the drift is stronger after negative than after positive announcements continues to hold. The results shown in the last three columns of Table 3 confirm these conclusions. The drift is statistically significant after both surprising and unsurprising announcements, but is much more pronounced after surprising announcements. The differences are significant for the long strategy and for the combined long-short strategy. We present the results of an augmented regression model in Column 2 of Table 4. We replace the dummy for the Corr treatment by three separate dummy variables which identify phases in treatment Corr with unsurprising and surprising announcements and phases following the first announcement which, as outlined above, is neither ambiguously surprising nor unsurprising. The results confirm our previous findings. There is a PEAD after both surprising and non-surprising announcements, yet it is significantly more pronounced after surprising announcements. These findings support Hypothesis 3.

Result Hypothesis 3. *Greater earnings surprises are followed by more substantial post-earnings-announcement drift.*

The results documented thus far have important implications. They stem from a well-controlled, precisely engineered laboratory experimental setting. In this setting, several

[15]In the first announcement, earnings are as likely to increase as they are to decrease, and the absolute size of the impact on the fundamental value of the announcement is independent of the direction of the announcement.

of the factors that have been conjectured to cause the post-earnings-announcement drift, such as informational asymmetries or changes in the riskiness of a stock that are related to earnings news, can be ruled out. We can thus focus on the question whether (as proposed by Bernard and Thomas, 1989, and many others) earnings autocorrelation causes the drift. Our results do not support this hypothesis. Rather, they suggest that earnings autocorrelation, while strengthening the PEAD, does not *cause* it.

Table 4: Regression analysis of window-to-window changes in taler closing quote midpoints. OLS regressions of returns over consecutive post-announcement windows. The dependent variable is the absolute change in taler closing midpoints per window. Returns are signed based on direction of previous earnings change (i.e., the signs of returns following negative announcements are reversed). 'Period0' is the period number within the session, rebased to the range 0...3 (instead of 1...4). 'Phase0' is the phase number within the period, rebased to 0...3 (instead of 1...4; thus excluding the phase preceding the first announcement, and designating the first post-announcement phase as 0). 'Correlated' is a dummy variable for treatment Corr. 'Correlated (No surprise)' is a dummy variable for an earnings change carrying the same sign as the earnings change in the previous announcement in the Corr treatment. 'Correlated (Surprise)' is a dummy for an earnings change carrying the opposite sign as in the previous announcement in the Corr treatment. 'Correlated (First announcement)' is a dummy for returns stemming from the phase following the first announcement in the Corr treatment (which is neither unambiguously surprising nor unsurprising). 'Window0' is the consecutive ID number of the time window, starting with the window following the announcement window (thus excluding the window directly after the announcement), rebased to 0...16 (instead of 1...17).

	Model 1	Model 2
Constant	1.181^{***}	1.137^{***}
	(0.141)	(0.141)
Correlated	0.227^{***}	
	(0.061)	
Correlated (No surprise)		0.054
		(0.062)
Correlated (Surprise)		0.563^{***}
		(0.137)
Correlated (First announcement)		0.378^{***}
		(0.041)
Positive earnings change	-0.199^{***}	-0.198^{***}
	(0.066)	(0.066)
Period0	-0.056	-0.059
	(0.034)	(0.035)
Phase0	-0.007	0.027
	(0.022)	(0.019)
Window0	-0.188^{***}	-0.189^{***}
	(0.029)	(0.029)
Window0^2	0.008^{***}	0.008^{***}
	(0.002)	(0.002)
R^2	0.004	0.005
Adj. R^2	0.004	0.004
Num. obs.	10389	10389
RMSE	5.578	5.576

$^{***}p < 0.005$, $^{**}p < 0.01$, $^{*}p < 0.05$. Standard errors, clustered at the Session level, in parentheses.

4.3. Trading strategies

The results in the previous sections clearly document the existence of PEAD in our experimental markets. They furthermore show that trading strategies can profitably exploit the PEAD on paper, i.e., when transaction costs are ignored. We now turn to the question of whether traders could employ strategies to profitably exploit the PEAD in practice, i.e., when transaction costs are taken into account. As we laid out in section 2.1, the empirical evidence from field data on this question is ambiguous.

The source of the transaction costs in our markets is the bid-ask spread. Figure 5 shows the evolution of the bid and ask prices over time, averaged over all periods of all sessions of the Base and Corr treatments. Note that in the figure we center bid and ask prices on the fundamental value of the stock, not on the quote midpoint.[16]

Four patterns are apparent from the figure. First, spreads are generally wide, making trading using market orders expensive. Average spreads in the 10s-windows range from less than 5 talers to more than 30 talers. These values are substantial, in particular when compared to the change in the fundamental value caused by an earnings announcement.[17] Second, spreads tend to decline over the course of a period. They are largest in Phase 0 and lower in later phases of a period. The largest spreads (observed immediately after the start of a period) are explained by the fact that the order book is empty at the beginning of Phase 0. Orders submitted early, far away from the fundamental value, can thus establish a wide spread that soon narrows as more orders are submitted. Third, spreads tend to widen after an earnings announcement and then decline during the remainder of the phase. What contributes most to the widening of the spread following an announcement is an increase

[16]The ask prices tend to be farther away from the fundamental value than the bid prices. This pattern is indicative of slight overpricing. We provide an analysis of price efficiency in section 4.4.

[17]This change is 10 talers in treatment Base and up to 28.125 talers after surprising announcements in treatment Corr. The spreads in our experimental markets are thus wider than those in typical financial markets outside of the lab, but finding wide spreads when earnings are unexpectedly small or large is not specific to our experimental setting. Ng et al. (2008) report average quoted spreads of 6.48% and 5.78%, respectively, for their decile portfolios with the lowest and highest unexpected earnings.

in market order submissions (i.e., limit order executions), a reduction in new limit order submissions, and, finally, an increase in limit order cancellations .[18] Furthermore—and unsurprisingly—the spread widens asymmetrically. The existing limit orders in the book constitute free trading options. Following the announcement, traders focus on that side of the book whose limit orders the announcement has moved closer to, or even into, the money. Thus, there are more transactions and cancellations of limit asks [bids] following positive [negative] earnings changes. The reduction in new limit order submissions, in contrast, is roughly symmetrical, suggesting that traders are too busy trying to execute or cancel stale limit orders to submit new ones. Fourth, spreads are consistently wider in treatment CORR than in treatment BASE. This is in line with the greater variability of both fundamental values (standard deviation 25.2 vs. 12.5 talers) and of trading prices (standard deviation 23.7 vs. 11.9 talers) in treatment CORR as compared to treatment BASE.

The regression results shown in Table 5 confirm the conclusions from the visual inspection of the data. The dependent variable is the bid-ask spread at the end of each 10s window. As independent variables we include different dummy variables for treatment CORR in Models 1 and 2. In Model 1 we include a dummy variable identifying any observations from treatment CORR, while in Model 2 we include separate dummy variables identifying observations from phases following non-surprising announcements in treatment CORR, observations from phases following surprising announcements in treatment CORR, and observations from the phase following the first announcement in treatment CORR. We furthermore include a dummy variable identifying observations from phases following positive earnings changes, and count variables for the period number, the phase within a period, and the window within a phase. As in Table 4, we also include the squared window number to allow for nonlinear

[18]There are about 6 times as many transactions as cancellations in the announcement window. Furthermore, transactions outnumber spread-widening cancellations (i.e., cancellation of the best bid or ask) approximately 30:1. These factors, which work to reduce the number of limit orders in the book, are accompanied by a reduction in new limit order submissions by more than half compared to the average level of limit order submissions over all windows.

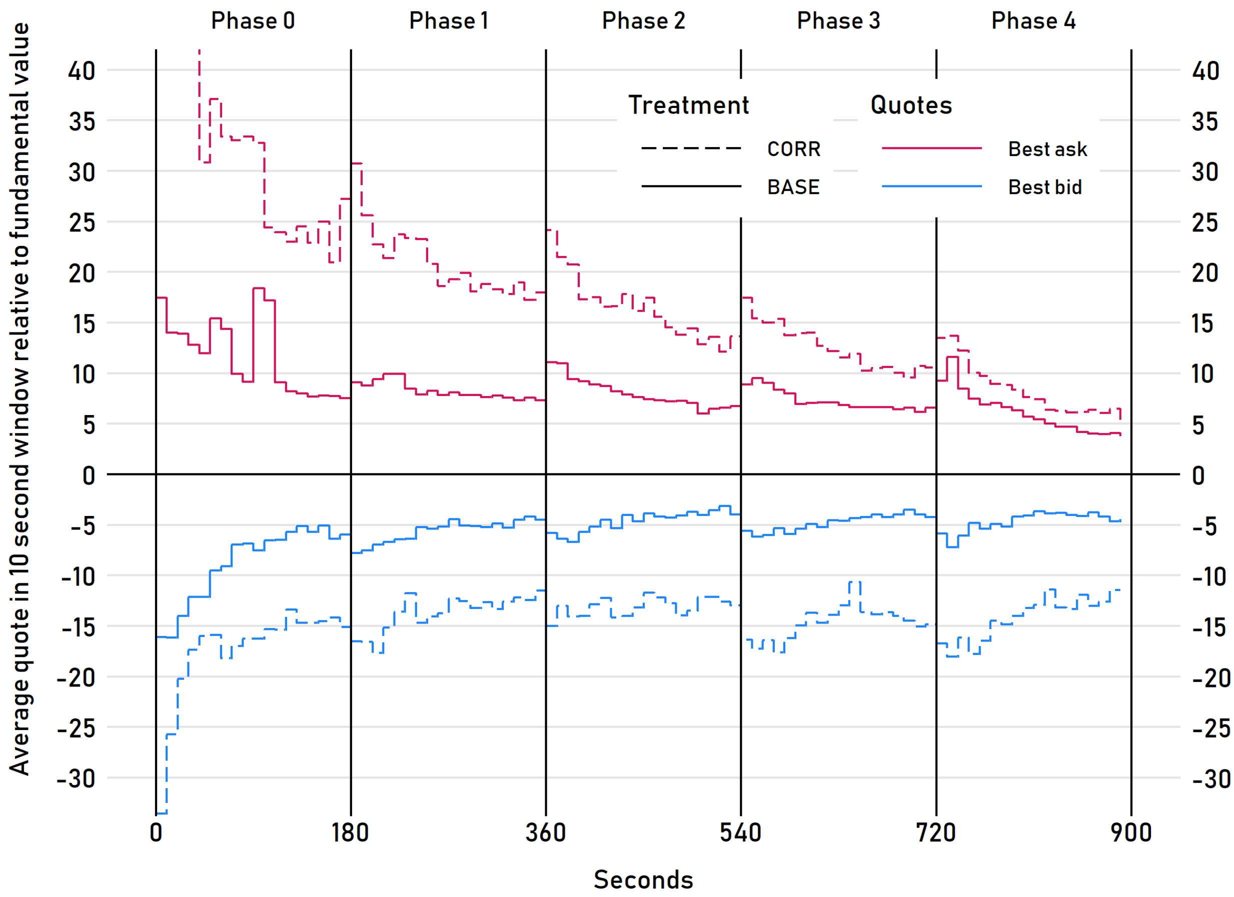

Figure 5: Development of spreads over time. Average difference between best ask quote and FV (red), and between best bid quote and FV (blue), in talers, over the trading period. Solid lines plot treatment Base data, dashed lines treatment Corr data.

dynamics.

The results show that spreads start out close to 30 talers at the beginning of the first period of treatment Base markets. Spreads are consistently larger in treatment Corr, even if the difference is only significant for observations following surprising (and first) announcements. There are no significant differences between the spreads following positive and following negative earnings announcements. Importantly, the initially high spreads decline over the periods of a session, the phases of a period, and the windows of a phase.

To analyze whether the PEAD can be profitably exploited we analyze three trading strategies and discuss them in turn. The first and simplest trading strategy consists of buying a share at the best ask price immediately after a positive earnings announcement and selling the share at the best bid price at the end of the phase (and doing the reverse after a negative announcement).[19]

Strategy 1. *Following an announcement with positive earnings news, buy a share at the best ask price. Then, 10s before the next announcement, sell it at the best bid price. Following an announcement with negative earnings news, do the reverse.*

The results of applying this strategy to our data are presented in the top panel of Table 6. We report separate results for treatment Base and for both unsurprising and surprising announcements in treatment Corr.[20] The six rows of the panel are based on different assumptions regarding the timing of the initial trade. The first line assumes that the opening trade is made immediately after the announcement, the second line assumes that it is made 2s after the announcement, and so on.

The results indicate that a strategy based on market orders is clearly unprofitable. Because of the wide bid-ask spreads we documented before, the strategy yields significantly

[19] Note that—in all analyses—we implicitly assume that the opening trades necessary to implement the proposed strategies do not affect the prices at which the position can be unwound at the end of the phase.

[20] As noted before, we exclude the results for the first phase in treatment Corr (the trading phase following the first announcement) from the analysis. A separate analysis of this phase (results not shown) reveals that profits are negative.

Table 5: Spread regressions. OLS regressions of taler spreads at the close of consecutive 10s windows starting with the announcement window. 'Correlated' is a dummy variable for treatment Corr. 'Correlated (No surprise)' is a dummy variable for an earnings change carrying the same sign as the earnings change in the previous announcement in the Corr treatment. 'Correlated (Surprise)' is a dummy variable for an earnings change carrying the opposite sign as in the previous announcement in the Corr treatment. 'Correlated (First announcement)' is a dummy variable for the phase following the first announcement in the Corr treatment (which is neither unambiguously surprising nor unsurprising). 'Period0' is the period number within the session, rebased to the range 0...3 (instead of 1...4). 'Phase0' is the phase number within the period, rebased to 0...3 (instead of 1...4; thus excluding the phase preceding the first announcement, and designating the first post-announcement phase with 0). 'Window' is the consecutive ID number of the time window (0...17), starting with the "announcement window", i.e., the 10s-window starting at the time of the announcement.

	(1) Spread	(2) Spread
Constant	28.179^{***}	27.140^{***}
	(3.263)	(2.987)
Correlated	16.163	
	(8.538)	
Correlated (No surprise)		14.949
		(9.287)
Correlated (Surprise)		15.075^{**}
		(5.759)
Correlated (First announcement)		19.687^{*}
		(9.171)
Positive earnings change	1.832	1.829
	(1.469)	(1.464)
Period0	-4.921^{***}	-4.922^{***}
	(1.615)	(1.600)
Phase0	-2.415^{***}	-1.714^{***}
	(0.704)	(0.506)
Window	-1.077^{***}	-1.080^{***}
	(0.193)	(0.192)
Window^2	0.033^{***}	0.033^{***}
	(0.008)	(0.008)
R^2	0.155	0.157
Adj. R^2	0.155	0.157
Num. obs.	11144	11144
RMSE	24.657	24.630

$^{***}p < 0.005$, $^{**}p < 0.01$, $^{*}p < 0.05$. Standard errors, clustered at the session level, in parentheses.

negative profits in most cases. Furthermore, the timing of the initial trade is important. The longer the delay between the announcement and the initial trade, the less profitable the strategy becomes. The latter finding is caused by the facts that (1) prices drift in the direction of the announcement, thus diminishing the potential profit available from trading in the direction of the announcement, and (2) spreads widen in the 10s following the announcement, thus increasing the transaction costs.

However, subjects in our experiments do not have to rely on market orders. Rather, they can open a position using a limit order and then simply hold on to their position until the end of the period where they receive [pay] the fundamental value for each stock that they have bought [sold short]. Submission of a limit order requires the specification of a price limit. We propose the following simple procedure. We assume that, after an earnings announcement, the trader submits a limit order with a price limit equaling the current quote midpoint. We then use the experimental data to check whether this limit order would have executed. We categorize it as executed if, at any time during the remainder of the phase, we observe a transaction at a price equal to or better than (i.e., lower [higher] in the case of a buy [sell] limit order) the price limit of the limit order.[21] If the limit order is not executed, the profit is recorded as zero. If the limit order is executed, we assume that the resulting position is held until the end of the period where the trader receives [pays] the fundamental value for each stock that (s)he owns [is short].[22] We thus have the following strategy:

Strategy 2. *Following an announcement with positive earnings news, submit a limit buy order, priced at the quote midpoint. If this order gets executed, hold the position until*

[21] Assume, for example, that following a positive announcement, the best bid price is 100 and the best ask price is 110. The trader then submits a limit buy order with a price limit of 105. We categorize this order as executed (at the price of 105) if, in our experimental data, a transaction was concluded, or an ask limit order submitted, at a price equal to or below 105 at any time during the remainder of the phase following the announcement.

[22] We use the FV of the current phase as the relevant closing price, even though in practice subjects can redeem their shares at the fundamental value only at the end of the period, after the final announcement. The FV of the current phase, however, is the expected value of this final redemption value. The returns shown in Table 6 are thus correct in expectation, but are not risk-free.

the end of the period and receive the fundamental value for each share held. Following an announcement with negative earnings news, do the reverse.

The middle panel of Table 6 presents the profits that can be obtained by following this strategy. The returns reported in the table include the cases in which the initial limit order is not executed and the profit is thus zero. The profits shown in Table 6 are therefore not conditional on execution. The proportion of initial limit orders that get executed is approximately 75.4% (when the limit order is placed 10s after the announcement). The profits are positive and significant in all cases we consider; in treatments BASE and CORR and, in the latter, both following surprising and unsurprising earnings announcements. Furthermore, we find that profits are higher the sooner after the earnings announcement the position is opened. They are also higher (for the 10s delay) in treatment CORR than in treatment BASE (Welch two-sample $t(379.73) = 1.9196$, $p = 0.0557$)[23] and, within treatment CORR, following surprising announcements than following unsurprising announcements, though not significantly so (Welch two-sample $t(88.183) = 0.2885$, $p = 0.7736$).

Subjects can implement Strategy 2 in our experiment. However, outside of the lab, there is no good analogue to the redemption of the shares at their respective fundamental values at the end of a period. Rather, traders in markets outside of the lab have to actively close their positions in order to realize their profits (or losses). We therefore devise a third strategy that can be implemented also in markets outside of the lab. We assume, as in Strategy 2, that the trader submits a limit order with a price limit equal to the current midpoint to open a position. If this limit order is not executed within the first 120s of the trading phase, it is cancelled and the profit is recorded as zero. If the limit order is executed, the trader submits a second limit order to close the position 120s after the start of the phase, again with a price limit equal to the then current midpoint.[24] If this limit order is executed during the next 50s, we calculate and record the profit or loss. If it is not executed, we assume that the

[23]The comparison excludes the first post-announcement phase in treatment CORR.

[24]We have also analyzed a version of this strategy where the price limit was set to one taler above the

trader cancels the limit order 170s after the announcement (i.e., 10s before the subsequent announcement) and closes the position at the current best bid or ask price using a market order.

Strategy 3. *Following an announcement with positive earnings news, submit a limit buy order, priced at the quote midpoint. If this order does not get executed in the first 120s of the phase, cancel it and do nothing else. If it gets executed, submit a limit sell order, priced at the quote midpoint, after 120s have elapsed since the announcement. If this limit order gets executed within the next 50s, do nothing else. If it does not get executed until 170s have elapsed since the announcement, cancel it and submit a market sell order instead. Following an announcement with negative earnings news, do the reverse.*

The third panel of Table 6 presents the profits that can be obtained by following Strategy 3. The proportion of initial limit orders that get executed is approximately 64.6% (when the limit order is placed 10s after the announcement).[25] The profits are lower than those for Strategy 2, but they are positive and significant in all cases considered. As before, the profits are higher in treatment Corr than in treatment Base (Welch two-sample $t(326.19) = 2.1026$, $p = 0.0363$),[23] and, within treatment Corr, higher following surprising than following unsurprising announcements. However, the latter difference is not significant (Welch two-sample $t(103.22) = 1.0467$, $p = 0.2977$). The profits decrease with increases in the delay between the announcement and the time at which the position is opened. The trading profits also are of an economically relevant magnitude, ranging from 1.13% to 6.13%.

Overall our results imply that the PEAD can be exploited profitably even after accounting for transaction costs. However, in order to earn a profit, traders have to "manage" transaction costs by using limit orders instead of market orders, a finding that echoes empirical results by Li (2016).

current best bid for a buy limit order and to one taler below the best ask in case of a sell limit order. The results are consistent with those presented below. In fact, profits are even slightly higher in this alternative version of the strategy.

[25] This percentage is lower than the corresponding percentage reported for Strategy 2 above. The reason is that, for Strategy 3, we require the limit order to be executed within 120s, while for Strategy 2 we only require it to be executed before the end of the phase, i.e., within 180s.

Result Hypothesis 4. *There are trading strategies that can profitably exploit the observed PEAD even after accounting for transaction costs.*

Table 6: Trading strategy returns per phase. 'Delay' is the number of seconds between the earnings announcement and the time at which the position is opened.

	BASE		CORR			
	-		No Surprise		Surprise	
Delay	Mean	SE	Mean	SE	Mean	SE
Strategy 1:	Open at best bid/ask, close at best bid/ask.					
0	-5.18***	(0.51)	-19.54***	(3.47)	-4.59	(5.27)
2	-5.52***	(0.49)	-20.22***	(3.46)	-5.60	(5.20)
4	-6.57***	(0.50)	-20.98***	(3.43)	-6.52	(5.09)
6	-7.22***	(0.48)	-21.88***	(3.39)	-10.27*	(5.10)
8	-7.83***	(0.54)	-23.18***	(3.41)	-13.67*	(5.16)
10	-8.21***	(0.55)	-24.52***	(3.49)	-16.50***	(5.11)
Strategy 2:	Open using limit order at quote midpoint, close by holding and earning FV.					
0	6.71***	(0.45)	9.48***	(0.91)	11.73***	(2.15)
2	6.39***	(0.45)	9.08***	(0.94)	9.82***	(2.71)
4	6.28***	(0.46)	9.00***	(0.94)	10.32***	(2.13)
6	5.95***	(0.44)	8.36***	(0.97)	8.73***	(2.09)
8	5.65***	(0.45)	7.64***	(0.92)	8.83***	(1.93)
10	5.54***	(0.44)	7.16***	(0.90)	7.74***	(1.77)
Strategy 3:	Open using limit order at quote midpoint, close by midpoint limit order after 120s or market order at 170s.					
0	2.20***	(0.27)	3.97***	(0.74)	6.13***	(1.60)
2	1.79***	(0.26)	3.52***	(0.72)	4.44*	(2.15)
4	1.68***	(0.25)	3.62***	(0.72)	5.78***	(1.57)
6	1.46***	(0.25)	3.06***	(0.70)	4.39***	(1.49)
8	1.23***	(0.26)	2.53***	(0.70)	4.40***	(1.36)
10	1.13***	(0.26)	2.18***	(0.71)	3.61***	(1.17)

Note:
*p<0.05, **p<0.01, ***p<0.005 (two-tailed t-test)

4.4. Price adjustment and fundamental value

An important advantage of the experimental setting over empirical market data is that we observe (and control) the fundamental values of the stocks traded in our experimental markets, and that we know how earnings announcements change these fundamental values. We can therefore explicitly test whether prices underreact to the news contained in an earnings announcement, whether and to which extent the post-earnings-announcement drift corrects the initial underreaction, and how long this adjustment takes if it occurs.

Before we report our results regarding price adjustment, we analyze the pricing efficiency of our experimental markets more generally. We follow the approach of Powell (2016) and define two measures of mispricing as follows.

$$GD \equiv \frac{1}{T-\eta} \cdot \int_{\eta}^{T} ln\left(\frac{P_{k,t}}{FV_k}\right) \cdot dt \tag{5}$$

and

$$GAD \equiv \frac{1}{T-\eta} \cdot \int_{\eta}^{T} \left| ln\left(\frac{P_{k,t}}{FV_k}\right) \right| \cdot dt \tag{6}$$

$P_{k,t}$ is the quote midpoint ("price") at time t in phase k, FV_k is the fundamental value in this phase, η is the duration of the announcement window (10s) and T is the length of each trading phase (180s). GD thus measures the time-weighted average geometric deviation of market prices from the fundamental value in the 170s following the announcement window. It is a signed measure that indicates by what percentage average prices exceed or fall short of the fundamental value. GAD measures the time-weighted average *absolute* geometric deviation. It can be interpreted as the (log) percentage by which average prices differ from the fundamental value, irrespective of the sign of the difference.

Table 7 shows the pricing efficiency results for our markets, first pooled over both treatments and then separately for treatments BASE and CORR. In each case we report measures of overall mispricing, of mispricing in the starting phase of each period (Phase 0, i.e., the phase leading up to the first earnings announcement), and of mispricing in phases following positive and negative earnings surprises. For treatment CORR the table furthermore

Table 7: Mispricing. Measures of relative (GD) and absolute (GAD) mispricing relative to FV. 'Starting phase' is Phase 0, the phase prior to the first earnings announcement. 'First announcement' is Phase 1, the phase following the first announcement, in treatment Corr.

Treatment	**Phases**	**GD**	**GAD**	**Observations**
All	All	1.44%	8.00%	800
	Starting phase	2.69%	9.05%	160
	Positive earnings change	-3.77%	6.09%	320
	Negative earnings change	6.03%	9.40%	320
BASE	All	1.31%	6.26%	400
	Starting phase	0.51%	6.86%	80
	Positive earnings change	-2.71%	5.04%	160
	Negative earnings change	5.73%	7.19%	160
CORR	All	1.58%	9.74%	400
	Starting phase	4.86%	11.23%	80
	First announcement	3.82%	9.89%	80
	Positive earnings change	-4.82%	7.13%	160
	First announcement	-2.65%	7.30%	40
	Surprise	-4.66%	7.50%	29
	No Surprise	-5.83%	6.95%	91
	Negative earnings change	6.33%	11.61%	160
	First announcement	10.29%	12.49%	40
	Surprise	7.58%	10.03%	29
	No Surprise	4.18%	11.72%	91

provides separate results for phases following surprising and unsurprising announcements.

When we pool all observations we find that prices on average exceed the fundamental value by 1.44%. We thus observe slight overpricing. This tendency is already visible in the starting phase leading up to the first earnings announcement, and it is slightly more pronounced in treatment Corr than in treatment Base (average GD of 1.58% as compared to 1.31%). Most importantly, the mispricing is negative after positive earnings surprises and positive after negative earnings surprises. Thus, while prices drift in the direction of the earnings announcement, they tend to stay below [above] the fundamental value after positive [negative] announcements. These findings imply that prices underreact to the news contained in the earnings announcement.

The absolute mispricing as measured by GAD is, by definition, larger than the signed mispricing GD because, in the former, incidences of positive and negative mispricing do not cancel out. GAD is greater after negative than after positive earnings surprises, a finding that is due to the general tendency towards overpricing documented above. Mechanically, when market prices lie above FV, a positive shock to FV is partially mitigated by the pre-existing overpricing, while a negative shock to FV is exacerbated by such overpricing. Interestingly, the absolute mispricing in treatment CORR does not differ markedly when preceded by surprising versus when preceded by unsurprising announcements.

We use regressions to gain a better understanding of the dynamics of absolute mispricing (GAD). As independent variables we include separate dummy variables identifying observations from phases following non-surprising announcements in treatment CORR, observations from phases following surprising announcements in treatment CORR, and observations from the phase following the first announcement in treatment CORR. We furthermore include count variables for the period within a session, the phase within a period, and the window within a phase. As before we also include the square of the window number to account for non-linearity. We estimate a model including data from all earnings announcements (model 1) and separate models for data from phases following positive and negative announcements, respectively (models 2 and 3).

We present the results in Table 8. The value of the intercept in model 1, 13.44%, measures the average GAD at the beginning of the first period in treatment BASE markets. Mispricing is larger in treatment CORR, yet with no clear differences between the phases following the first announcement, those following unsurprising announcements and those following surprising announcements. The negative coefficients on the count variables for the period and the phase indicate that mispricing tends to decrease over the course of the experiment. This tendency is, of course, a direct reflection of the existence of a post-earnings announcement drift. Within a phase, the mispricing decreases at a decreasing rate, as is evidenced by the negative coefficient on the 'Window' count variable and the positive coefficient on its square. Finally, a comparison of the results for models 2 and 3 confirms our earlier finding of greater mispricing following negative earnings surprises.

Table 8: Regression analysis of absolute mispricing. OLS regressions of mispricing at the close of consecutive 10s-windows starting at the time of the announcement. Model (1) reports mispricing pooled across announcement types, while models (2) and (3) report mispricing following positive and negative announcements, respectively. The dependent variable is the geometric absolute deviation (GAD) in %, calculated using time-weighted midpoints for each window. 'Correlated (No surprise)' is a dummy variable for phases following an earnings change carrying the same sign as in the preceding announcement in treatment Corr. 'Correlated (Surprise)' is a dummy for an earnings change carrying the opposite sign as in the preceding announcement in treatment Corr. 'Correlated (First announcement)' is a dummy for observations from the phase following the first announcement. 'Period0' is the period number within the session, rebased to the range 0...3 (instead of 1...4). 'Phase0' is the phase number within the period, rebased to 0...3 (instead of 1...4; in the case of the 'Pooled' regression, the phase prior to the first announcement is included, such that the range of the phase number is 0...4 in this case). 'Window' is the consecutive ID number of the time window (0...17), starting with the "announcement window", i.e., the 10s-window starting at the time of the announcement.

	(1) Pooled	(2) Positive announcements	(3) Negative announcements
Constant	13.435***	10.377***	15.263***
	(1.374)	(1.229)	(1.700)
Correlated (No surprise)	3.743*	2.193	5.396**
	(1.785)	(1.661)	(1.945)
Correlated (Surprise)	3.819***	3.470*	3.930***
	(0.984)	(1.475)	(1.331)
Correlated (First announcement)	3.270*	1.689	3.674*
	(1.435)	(1.779)	(1.812)
Period0	−1.092***	−1.082***	−0.443
	(0.334)	(0.331)	(0.381)
Phase0	−0.658***	−0.539***	−1.353***
	(0.225)	(0.176)	(0.342)
Window	−0.829***	−0.624***	−1.050***
	(0.099)	(0.076)	(0.138)
Window^2	0.031***	0.026***	0.039***
	(0.004)	(0.003)	(0.006)
R^2	0.099	0.097	0.159
Adj. R^2	0.098	0.096	0.158
Num. obs.	14026	5730	5608
RMSE	8.759	6.438	8.186

$^{***}p < 0.005$, $^{**}p < 0.01$, $^{*}p < 0.05$. Standard errors, clustered at the session level, in parentheses.

The analysis so far has focused on the *level* of mispricing, i.e., the distance between the prices in our markets and the fundamental value. We will now turn to a more detailed analysis of the changes in prices in response to the changes in fundamental values caused by the earnings announcements.

We proceed as follows. We base our analysis on quote midpoints but continue to refer to them as "prices". We normalize the price prior to an earnings announcement to 0 and the pre-announcement price plus the change in the fundamental value to 100% [-100%] in the case of a positive [negative] announcement. A post-announcement price of ($\pm$)100% implies that the price change equals exactly the change in the fundamental value.[26] An price change of less than [more than] ($\pm$)100% indicates underreaction [overreaction].

Panel (a) of Figure 6 presents the results for the pooled observations from both treatments. The figure shows clearly that the initial price reaction to the earnings announcement is smaller than the implied change in the fundamental value (and more so after negative than after positive earnings surprises). Throughout the remainder of the trading phase the price drifts further in the direction of the announcement but fails to fully adjust. Following positive earnings surprises, prices reflect 33% of ΔFV within the 10s announcement window and reach around 46% within 20s of the announcement. However, by the end of the phase they only reflect 65% of the change in fundamental value. Following negative earnings announcements, prices reflect only 18% of ΔFV within the post-announcement window and it takes another 30s to surpass 50%. Yet by the end of the phase following negative earnings surprises, prices reflect about 87% of ΔFV.

The finding that prices adjust more fully to changes in fundamental value after negative than after positive earnings surprises seems to contradict our earlier finding of larger mispricing following negative earnings surprises. Two factors explain this apparent contradiction.

[26]Note that this is only a statement regarding the price *change*; it does not imply that the price *level* equals the fundamental value. Therefore, the analysis is unaffected by the slight overpricing in the experimental markets documented above.

Figure 6: Adjustment of stock price as a percentage of the change in fundamental value induced by an announcement. Panel (a) plots results for the pooled data from all treatments; panel (b) reports results separately for treatments BASE (left) and CORR (right); panel (c) plots only CORR data and reports results separately for unsurprising (left) and surprising (right) earnings news. The blue, upward trending [orange, downward trending] line plots price adjustment following positive [negative] earnings news. The bold, black, horizontal lines indicate full adjustment of prices to the change in FV induced by the earnings announcement. The dotted horizontal line at 0 indicates the price level at the moment of the earnings announcement. The dashed horizontal line indicates FV prior to the earnings announcement.

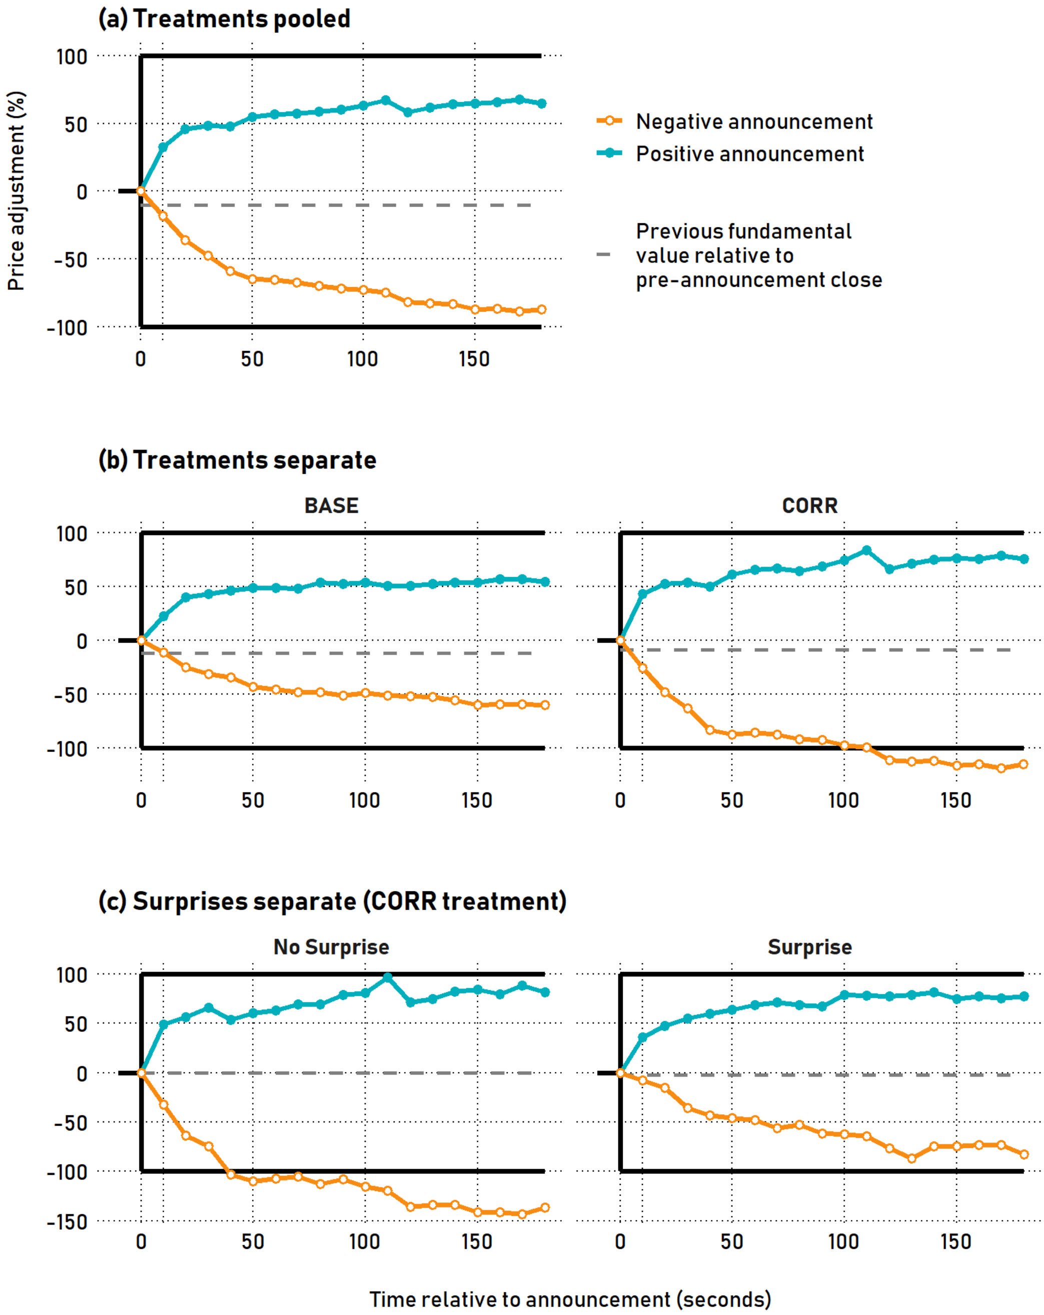

First, as shown above, prices on average slightly exceed the fundamental values of the stocks. This implies that, even after a price change that fully reflects the change in fundamental value following a negative earnings surprise, prices can still be above fundamental values, contributing to mispricing. The dashed gray lines in Figure 6 represent the fundamental values of the stocks prior to the earnings announcement. The lines help to show that the adjustment of prices following positive and negative announcements is roughly symmetric relative to this pre-announcement fundamental value. Second, the mispricing reported in Table 7 is an average over the entire trading phase following the announcement window, while the 90% price adjustment mentioned above reflects the price level at the end of the phase.

Panel (b) of Figure 6 shows separate results for the two treatments. The adjustment of prices to the news contained in the announcement is faster and more complete in treatment Corr than in treatment Base.[27] This is particularly true after negative earnings surprises. Here, the price change fully reflects the change in fundamental value implied by the announcement in treatment Corr after around 110s. We conjecture that the more complete price adjustment in Corr may be driven by subjects devoting greater attention to announcements in this treatment. Subjects know that there can be two types of announcements—surprising and non-surprising—with very different implications for the asset value, and therefore evaluate the announcements more thoroughly.

The findings of this section can be summarized as follows.

Result Hypothesis 5. *Prices initially underadjust to the information content of the earnings announcement and then continue to drift in the direction of the earnings surprise. The price adjustment is more complete in treatment* Corr.

The final question we wish to analyze is whether the experimental subjects can correctly assess the effect of earnings autocorrelation on asset values. In section 4.2 we documented

[27] As will become clear from Panel (c), price changes are faster and more complete than in treatment Base following both unsurprising and surprising announcements in treatment Corr.

that the PEAD is more pronounced after surprising than after non-surprising earnings announcements. At the same time, Panel (c) of Figure 6 documents that prices adjust faster and—particularly in the case of negative announcements—more completely to changes in fundamental value after non-surprising announcements. In fact, we even observe overshooting after negative non-surprising announcements (i.e., a price change in excess of the change in fundamental value).

This pattern is consistent with underestimation of the effect of earnings autocorrelation. Recall that, in treatment Corr, non-surprising announcements lead to a smaller absolute change in the fundamental value than do announcements in treatment Base, simply because the announcement could be anticipated because of the serial correlation. Subjects who underestimate the effect of earnings autocorrelation thus *overestimate* the effect of a non-surprising announcement on the fundamental value. Consequently, we expect to find a stronger adjustment of prices to the change in the fundamental value and a less pronounced PEAD. By the same argument, subjects who underestimate the effect of earnings autocorrelation thus *underestimate* the effect of a surprising announcement on the fundamental value. Consequently, we expect to find a weaker adjustment of prices to the change in the fundamental value and a more pronounced PEAD after surprising announcements. This is precisely the pattern we documented above. Our results thus confirm hypothesis 6 and can be summarized as follows.

Result Hypothesis 6. *In the presence of earnings autocorrelation, prices adjust more fully to the information content of non-surprising than of surprising announcements. This pattern is consistent with underestimation of the implications for asset values of earnings autocorrelation.*

5. Conclusion

The post-earnings-announcement drift (PEAD) is one of the—if not the—most solidly documented market anomalies in the literature. Empirical investigations into its causes are complicated by the fact that many relevant variables cannot be directly observed and therefore need to be estimated. These difficulties are absent under the controlled conditions of the experimental laboratory. In this paper we report results of a series of experimental markets designed to analyze the importance of earnings autocorrelation for the emergence and strength of the post-earnings-announcement drift. To focus on this aspect we carefully design the experiments to rule out several other potential causes of the drift, such as risk-based explanations and informational asymmetries.

Our results provide clear evidence of PEAD in stock returns, both without and with positive autocorrelation in earnings news. We thus show that earnings autocorrelation is not a necessary condition for PEAD. We do find, however, that autocorrelation increases the strength of the drift. We further show that the drift is driven by prices adjusting slowly and incompletely to changes in fundamental values. Finally, we demonstrate that the PEAD in our markets can be profitably exploited, thus underlining the economic significance of the phenomenon.

Our results imply that the post-earnings-announcement drift is not solely a field phenomenon, but can be replicated in the simplified and controlled environment of the experimental lab. This is an important finding in its own right because it demonstrates that the PEAD is not driven by the idiosyncrasies of the current institutional environment of securities trading. Exchange regulations, brokers, analysts, news services, etc., are all incidental to the phenomenon. We believe that the opportunities offered by the experimental method should thus be further exploited in future research. The results of our study suggest that pursuing explanations based on investor attention may be a particularly promising strategy.

References

References

Ball, R., 1992. The Earnings-Price Anomaly. Journal of Accounting and Economics 15, 319–345.

Ball, R., Bartov, E., 1996. How Naive is the Stock Market's Use of Earnings Information? Journal of Accounting and Economics 21, 319–337.

Ball, R., Brown, P., 1968. An Empirical Evaluation of Accounting Income Numbers. Journal of Accounting Research 6, 159–178.

Ball, R., Kothari, S., Watts, R., 1993. The Economic Determinants of the Relation between Earnings Changes and Stock Returns. The Accounting Review 68, 622–638.

Bathke, A., Morton, R., Notbohm, M., Zhang, T., 2014. Objective Estimation versus Subjective Perceptions of Earnings Patterns and Post-Earnings-Announcement Drift. Accounting and Finance 54, 305–334.

Bathke Jr, A.W., Lorek, K.S., Willinger, G.L., 2006. The security market's reaction to firms' quarterly earnings evidencing varying degrees of autocorrelation. Advances in Accounting 22, 29–43.

Battalio, R.H., Mendenhall, R.R., 2005. Earnings expectations, investor trade size, and anomalous returns around earnings announcements. Journal of Financial Economics 77, 289–319.

Battalio, R.H., Mendenhall, R.R., 2011. Post-earnings announcement drift: Bounds on profitability for the marginal investor. Financial Review 46, 513–539.

Benjamin, D.J., Berger, J.O., Johannesson, M., Nosek, B.A., Wagenmakers, E.J., Berk, R., Bollen, K.A., Brembs, B., Brown, L., Camerer, C., Cesarini, D., Chambers, C.D., Clyde, M., Cook, T.D., de Boeck, P., Dienes, Z., Dreber, A., Easwaran, K., Efferson, C., Fehr, E., Fidler, F., Field, A.P., Forster, M., George, E.I., Gonzalez, R., Goodman, S., Green, E., Green, D.P., Greenwald, A.G., Hadfield, J.D., Hedges, L.V., Held, L., Hua Ho, T., Hoijtink, H., Hruschka, D.J., Imai, K., Imbens, G., Ioannidis, J.P.A., Jeon, M., Jones, J.H., Kirchler, M., Laibson, D., List, J., Little, R., Lupia, A., Machery, E., Maxwell, S.E., McCarthy, M., Moore, D.A., Morgan, S.L., Munafó, M., Nakagawa, S., Nyhan, B., Parker, T.H., Pericchi, L., Perugini, M., Rouder, J., Rousseau, J., Savalei, V., Schönbrodt, F.D., Sellke, T., Sinclair, B., Tingley, D., van Zandt, T., Vazire, S., Watts, D.J., Winship, C., Wolpert, R.L., Xie, Y., Young, C., Zinman, J., Johnson, V.E., 2018. Redefine statistical significance. Nature Human Behaviour 2, 6–10. URL: `https://doi.org/10.1038/s41562-017-0189-z`, doi:`10.1038/s41562-017-0189-z`.

Bernard, V., 1993. Stock Price Reactions to Earnings Announcements: A Summary of Recent Anomalous Evidence and Possible Explanations. Russell Sage Foundation. volume 1. chapter Stock Price Reactions to Earnings Announcements: A Summary of Recent Anomalous Evidence and Possible Explanations. pp. 303–340.

Bernard, V., Thomas, J., 1989. Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium? Journal of Accounting Research 27, 1–36.

Bernard, V., Thomas, J., 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.

Bloomfield, R., Libby, R., Nelson, M., 2003. Do Investors Overrely on Old Elements of the Earnings Time Series? Contemporary Accounting Research 20, 1–31.

Calegari, M., Fargher, N.L., 1997. Evidence that Prices Do Not Fully Reflect the Implications of Current Earnings for Future Earnings: An Experimental Markets Approach. Contemporary Accounting Research 14, 397–433.

Chang, T., Hartzmark, S., Solomon, D., Soltes, E., 2017. Being Surprised by the Unsurprising: Earnings Seasonality and Stock Returns. The Review of Financial Studies 30, 281–323.

Chordia, T., Goyal, A., Sadka, G., Sadka, R., Shivakumar, L., 2009. Liquidity and the Post-Earnings-Announcement Drift. Financial Analysts Journal 65, 18–32.

Daniel, K., Hirshleifer, D., Sun, L., 2020. Short- and Long-Horizon Behavioral Factors. Review of Financial Studies 33, 1673–1736.

DellaVigna, S., Pollet, J.M., 2009. Investor inattention and friday earnings announcements. The Journal of Finance 64, 709–749.

Fama, E., 1998. Market Efficiency, Long-Term Returns, and Behavioral Finance. Journal of Financial Economics 49, 283–306.

Fischbacher, U., 2007. z-tree: Zurich toolbox for ready-made economic experiments. Experimental Economics 10, 171–178. doi:10.1007/s10683-006-9159-4.

Francis, J., Lafond, R., Olsson, P., Schipper, K., 2007. Information Uncertainty and Post-Earnings-Announcement-Drift. Journal of Business Finance and Accounting 34, 403–433.

Frazzini, A., 2006. The Disposition Effect and Underreaction to News. Journal of Finance 61, 2017–2046.

Freeman, D.J., Kimbrough, E.O., Petersen, G.M., Tong, H.T., 2018. Instructions. Journal of the Economic Science Association 4, 165–179.

Freeman, R.N., Tse, S., 1989. The multiperiod information content of accounting earnings: Confirmations and contradictions of previous earnings reports. Journal of Accounting Research 27, 49–79.

Hirshleifer, D.A., Lim, S.S., Teoh, S.H., 2009. Driven to Distraction: Extraneous Events and Underreaction to Earnings News. Journal of Finance 64, 2289–2325.

Hlavac, M., 2018. stargazer: Well-Formatted Regression and Summary Statistics Tables. Central European Labour Studies Institute (CELSI). Bratislava, Slovakia. URL: `https://CRAN.R-project.org/package=stargazer`. r package version 5.2.2.

Hou, K., Peng, L., Xiong, W., 2009a. A Tale of Two Anomalies: The Implications of Investor Attention for Price and Earnings Momentum. Working Paper .

Hou, K., Xiong, W., Peng, L., 2009b. A tale of two anomalies: The implications of investor attention for price and earnings momentum. Available at SSRN 976394 .

Hung, M., Li, X., Wang, S., 2015. Post-earnings-announcement drift in global markets: Evidence from an information shock. The Review of Financial Studies 28, 1242–1283.

Ke, B., Ramalingegowda, S., 2005. Do institutional investors exploit the post-earnings announcement drift? Journal of Accounting and Economics 39, 25–53.

Kim, D., Kim, M., 2003. A Multifactor Explanation of Post-Earnings Announcement Drift. Journal of Financial and Quantitative Analysis 38, 383–398.

Leifeld, P., 2013. texreg: Conversion of statistical model output in R to LaTeX and HTML tables. Journal of Statistical Software 55, 1–24.

Li, J., 2016. Slow price adjustment to public news in after-hours trading. The Journal of Trading 11, 16–31.

Maines, L., Hand, J., 1996. Individuals' Perceptions and Misperceptions of Time Series Properties of Quarterly Earnings. The Accounting Review 71, 317–336.

Ng, J., Rusticus, T.O., Verdi, R.S., 2008. Implications of Transaction Costs for the Post-Earnings Announcement Drift. Journal of Accounting Research 46, 661–696.

Noussair, C.N., Tucker, S., 2016. Cash inflows and bubbles in asset markets with constant fundamental values. Economic Inquiry 54. doi:10.1111/ecin.12320.

Palan, S., 2013. A review of bubbles and crashes in experimental asset markets. Journal of Economic Surveys 27, 570–588.

Palan, S., 2015. GIMS - software for asset market experiments. Journal of Behavioral and Experimental Finance 5, 1–14. doi:10.1016/j.jbef.2015.02.001.

Pavlova, I., Parhizgari, A., 2011. In search of momentum profits: are they illusory? Applied Financial Economics 21, 1617–1639.

Powell, O., 2016. Numeraire independence and the measurement of mispricing in experimental asset markets. Journal of Behavioral and Experimental Finance 9, 56–62.

R Core Team, 2017. R: A language and environment for statistical computing. URL: https://www.R-project.org/.

Rangan, S., Sloan, R., 1998. Implications of the Integral Approach to Quarterly Reporting for the Post-Earnings-Announcement Drift. The Accounting Review 73, 353–371.

Rendleman, R., Jones, C., Latane, H., 1987. Further Insight into the Standardized Unexpected Earnings Anomaly: Size and Serial Correlations Effect. The Financial Review 22, 131–144.

Richardson, S., Tuna, I., Wysocki, P., 2010. Accounting anomalies and fundamental analysis: A review of recent research advances. Journal of Accounting and Economics 50, 410–454.

Sadka, R., 2006. Momentum and Post-Earnings Announcement Drift Anomalies: The Role of Liquidity Risk. Journal of Financial Economics 80, 309–349.

Wickham, H., 2016. ggplot2: Elegant Graphics for Data Analysis. Springer-Verlag New York. URL: https://ggplot2.tidyverse.org.

Zhang, G., Zhang, S., 2013. Information efficiency of the us credit default swap market: Evidence from earnings surprises. Journal of financial stability 9, 720–730.

Zhu, H., 2019. kableExtra: Construct Complex Table with 'kable' and Pipe Syntax. URL: https://CRAN.R-project.org/package=kableExtra. r package version 1.1.0.

Appendix A. Additional figures and tables

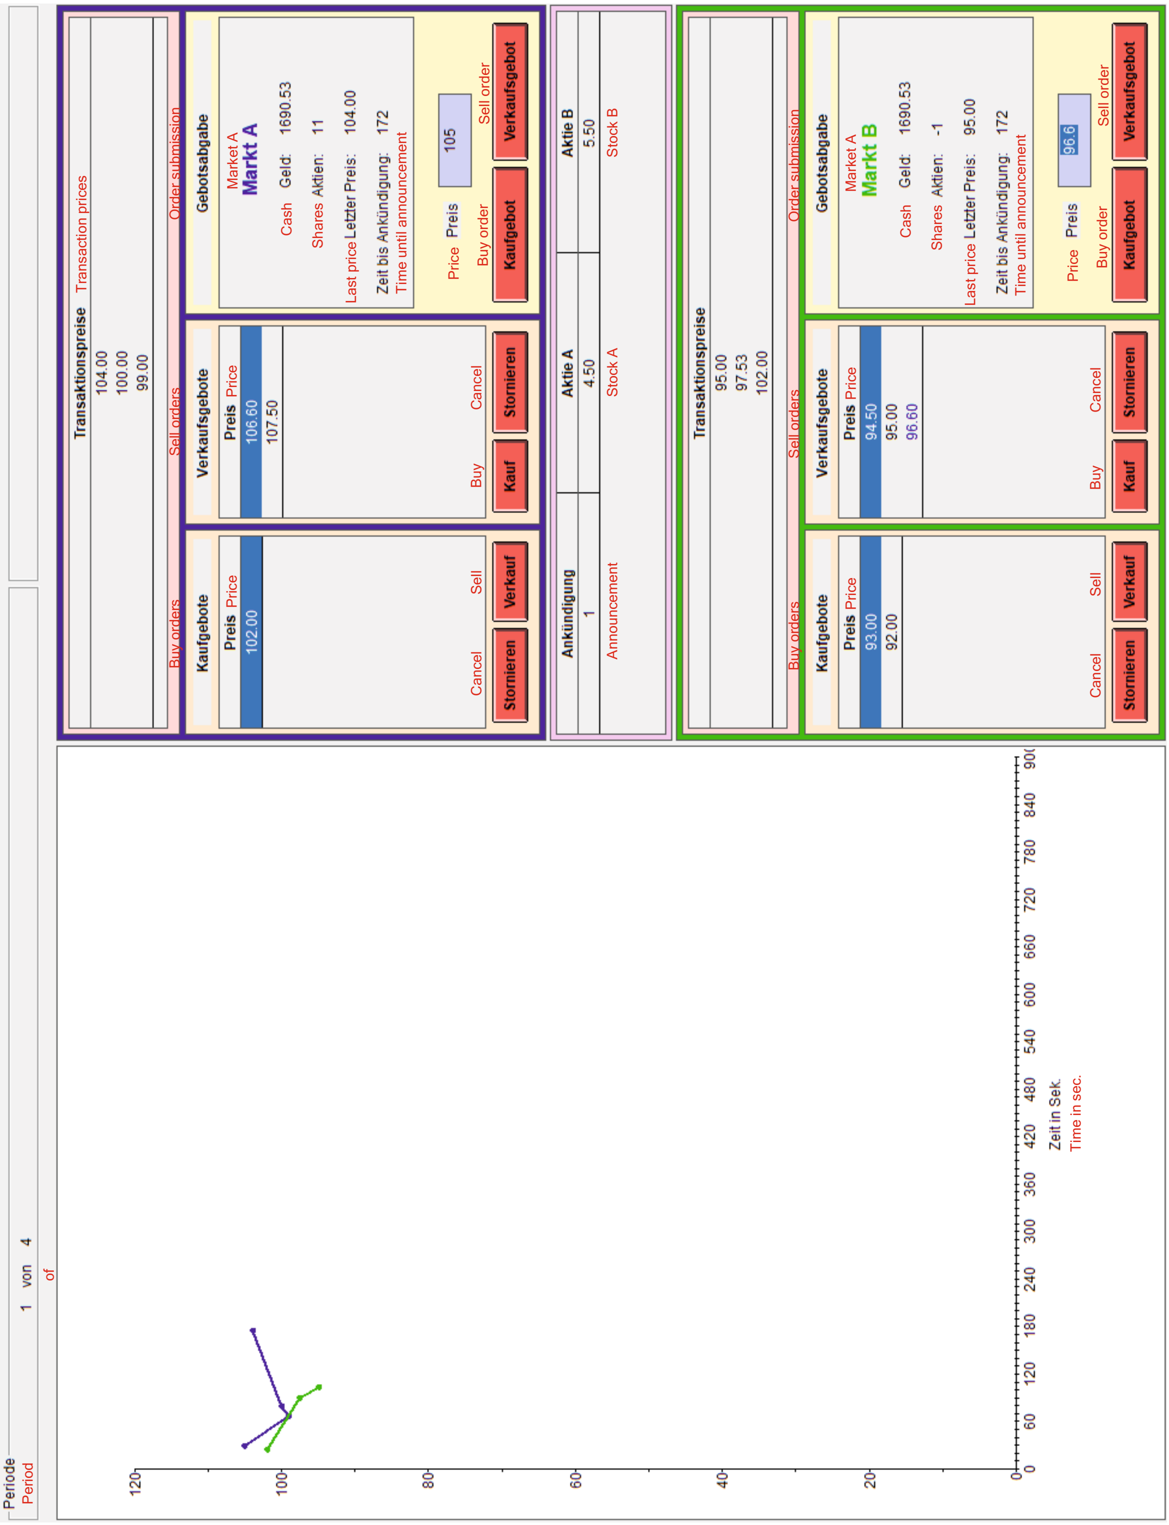

Figure A.1: Screenshot of the trading interface. Translations (in red) were not present in the experiment.

Table A.1: Post-earnings-announcement drift. Mean log returns in quote midpoints from the end of the post-announcement-window until the end of the phase. The data for treatment CORR separated by the surprise variable exclude the phase following the first announcement. We report t-statistics in parentheses.

	Treatments				**CORR Treatment[a]**		
	All[b]	**BASE**	**CORR**	**Δ**	**No Surprise**	**Surprise**	**Δ**
Long	3.85***	3.19***	4.50***	1.30	2.21*	11.52***	9.31***
	(7.35)	(6.68)	(4.85)	(1.25)	(2.11)	(4.49)	(3.36)
Short	-8.62***	-4.75***	-12.84***	-8.08***	-11.61***	-16.30***	-4.69
	(-14.04)	(-8.97)	(-12.45)	(-6.97)	(-8.97)	(-5.88)	(-1.53)
Long-Short	12.42***	7.81***	17.53***	9.72***	14.27***	25.52***	11.24***
	(15.47)	(10.44)	(13.06)	(6.32)	(7.91)	(8.12)	(3.10)

Note:
*p<0.05, **p<0.01, ***p<0.005 (two-tailed t-test)
Values based on all phases where both the first and the last 10s windows following an announcement (windows 0 and 17) have both a bid and an ask; thus permitting us to calculate quote midpoints.
[a] CORR Treatment data separated by the surprise variable excludes the first announcement.
[b] Only phases with valid midpoints to calculate return considered (includes 95% , 87.5% , and 85% of Long, Short, and Long-Short phases, respectively, for 'All' column).

Table A.2: Regression analysis of window-to-window log returns of closing quote midpoints. OLS regressions of returns over consecutive post-announcement windows. The dependent variable is the log-return in % using closing midpoints per window. Returns are signed based on direction of previous earnings change (i.e., the signs of returns following negative announcements are reversed). 'Period0' is the period number within the session, rebased to the range 0...3 (instead of 1...4). 'Phase0' is the phase number within the period, rebased to 0...3 (instead of 1...4; thus excluding the phase preceding the first announcement, and designating the first post-announcement phase with 0). 'Correlated' is a dummy variable for treatment CORR. 'Correlated (No surprise)' is a dummy variable for an earnings change carrying the same sign as the earnings change in the previous announcement in the CORR treatment. 'Correlated (Surprise)' is a dummy for an earnings change carrying the opposite sign as in the previous announcement in the CORR treatment. 'Correlated (First announcement)' is a dummy for returns stemming from the phase following the first announcement in the CORR treatment (which is neither unambiguously surprising nor unsurprising). 'Window0' is the consecutive ID number of the time window, starting with the window following the announcement window (thus excluding the window directly after the announcement), rebased to 0...16 (instead of 1...17).

	Model 1	Model 2
Constant	1.190^{***}	1.158^{***}
	(0.123)	(0.124)
Correlated	0.245^{***}	
	(0.055)	
Correlated (No surprise)		0.101
		(0.063)
Correlated (Surprise)		0.542^{***}
		(0.142)
Correlated (First announcement)		0.355^{***}
		(0.043)
Positive earnings change	-0.228^{***}	-0.227^{***}
	(0.069)	(0.068)
Period0	−0.055	−0.058
	(0.033)	(0.033)
Phase0	0.006	0.032
	(0.021)	(0.017)
Window0	-0.188^{***}	-0.189^{***}
	(0.029)	(0.029)
Window0^2	0.008^{***}	0.008^{***}
	(0.002)	(0.002)
R^2	0.006	0.006
Adj. R^2	0.005	0.005
Num. obs.	10389	10389
RMSE	5.033	5.032

$^{***}p < 0.005$, $^{**}p < 0.01$, $^{*}p < 0.05$. Standard errors, clustered at the session level, in parentheses.