JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 59, No. 6, Sep. 2024, pp. 2799–2837

doi:10.1017/S0022109023000881

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

Josef Fink
McKinsey and Company Vienna
josef_fink@mckinsey.com

Stefan Palan
University of Graz Institute of Banking and Finance,
stefan.palan@uni-graz.at (corresponding author)

Erik Theissen
University of Mannheim Finance area
theissen@uni-mannheim.de

Abstract

Post-earnings-announcement drift (PEAD) is one of the most solidly documented asset pricing anomalies. We use the controlled conditions of the 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, confirming that earnings autocorrelation is not a necessary condition for PEAD. Instead, it acts as an accelerator: PEAD is stronger when earnings surprises are correlated. We further show that market prices underadjust to fundamental value changes, and that trading strategies can profitably exploit the PEAD.

I. Introduction

Over the past 4 decades, researchers have uncovered a large number of seemingly anomalous patterns in financial market returns. Their existence raises the question of whether they are related to (possibly unidentified) sources of risk and are thus consistent with financial market equilibrium, or whether they constitute

We thank Jennifer Conrad and Elena Asparouhova—respectively, the editor and referee of the current version—as well as two anonymous reviewers of a previous submission, Ray Ball, Te Bao, Bruno Biais, Peter Bossaerts, Jürgen Brandner, Peiran Jiao, Alexander Kempf, Michael Kirchler, Olaf Korn, Christian Leuz, Roland Mestel, Kerstin Mitterbacher, Stefan Nagel, Thomas Post, Ryan Riordan, Andrea Schertler, Thomas Stöckl, Martin Weber, Yilong Xu, and the participants of the 2019 Austrian Working Group on Banking and Finance workshop, the 2020 virtual finance workshop at Radboud University, the 2021 Experimental Finance conference, and seminar participants at the UC Louvain, the University of Cologne, the University of Graz, the University of Maastricht, and the University of Mannheim for valuable comments. We also thank Adele Theiler for research assistance. This work was supported by the Austrian Science Fund FWF (project no. P32124-G27). The research project was approved by the IRB of the University of Graz (case ID 39/57/63). The authors have no conflicts of interest to disclose.

systematic deviations from equilibrium. One category of such patterns have subsequently become accepted as proxying for systematic risk factors (e.g., the size and BM effects, see Fama and French (1993)). A second category was found to constitute "mistakes" that, once discovered, disappeared as the market corrected and accounted for them (e.g., McLean and Pontiff (2016)). It is the third category that truly captured the attention of researchers: despite having been discovered, they neither disappeared, nor yielded to a ready risk-based explanation. Unfortunately, distinguishing between risk-based and behavioral explanations is onerous because fundamental asset values are unobservable and because measuring investor expectations and identifying the information investors base their decisions on is difficult.

One group of anomalies in the third category are characterized by underreaction to public information. Arguably, the most prominent members of this group are momentum and the post-earnings-announcement drift (PEAD). We focus on the latter. PEAD is the phenomenon that share prices react slowly to the surprise component of earnings announcements. The price initially underreacts and then drifts in the direction of the surprise for extended periods of time. A large number of articles (reviewed in Section II) have tested various explanations for the phenomenon. However, no consensus has emerged so far. Because of its persistence and the lack of a rational explanation, PEAD has been characterized as "an anomaly above suspicion" and "the granddaddy of underreaction events" by Eugene Fama (Fama (1998), pp. 304 and 286).

Ours is the first article to use experimental asset markets to study the emergence and the determinants of PEAD. We design our markets such that fundamental values are known. The experimental approach is a valuable complement to empirical studies using field data. It allows us to directly control and deliberately vary the variables of interest. In particular, we carefully design the earnings process of the firms traded in our markets. We also clearly communicate the characteristics of this earnings process to traders so they need not estimate expected earnings from time-series models or analyst forecasts. As a consequence, we observe earnings surprises and traders' reactions to them without noise or bias. This allows us to show that PEAD can be reliably reproduced in experimental asset markets. As we lay out below, we focus on the link between earnings autocorrelation and PEAD. We show that earnings autocorrelation amplifies the PEAD but does not cause it. We further find that underlying the drift is a slow and incomplete adjustment of prices to fundamental values, which in our setting are unambiguously defined. Finally, we show that trading strategies can profitably exploit the drift in our markets, even after accounting for transaction costs.

The existing evidence on PEAD is inconsistent with rational, risk-based explanations. We thus focus on mispricing as the driver behind the drift. We control the timing and content of the information that is available to traders to ensure that information is symmetrically distributed and rule out leakage. Our design furthermore eliminates aggregate fundamental value risk such that, in equilibrium, prices equal expected values. This has two benefits. First, we know how each earnings announcement affects the fundamental value. Second, we can analyze how quickly prices approach fundamental values after announcements.

One of the most widely accepted drivers of the PEAD is that "prices are affected by investors who fail to recognize fully the implications of current earnings for future earnings" (Bernard and Thomas (1989), p. 2). Obviously, these implications are relatively easy to assess when earnings news are serially uncorrelated, but are less readily evaluated when earnings are serially correlated. Empirical data typically shows that corporate earnings are in fact serially correlated (e.g., Ball and Bartov (1996)). However, determining which information investors base their decisions on is difficult, and in particular, researchers cannot readily measure investors' expectations regarding future earnings autocorrelation.

We isolate the impact of earnings autocorrelation on the emergence of a PEAD by analyzing two experimental settings, one without and one with earnings autocorrelation. In our baseline setting, we create common knowledge that earnings are serially uncorrelated. Earnings increase or decrease with equal probability, and by a constant amount. In this setting, the current earnings surprise contains *no* information about future earnings surprises. We observe statistically significant PEAD even in this simple setting. We thus confirm that earnings autocorrelation is not a necessary condition for PEAD.

In our correlated earnings setting, earnings surprises are positively correlated and we create common knowledge of the precise nature of this correlation. Here, we observe a stronger PEAD than in the baseline setting, implying that earnings autocorrelation is an accelerator of the PEAD. Autocorrelation thus does not cause the drift, but strengthens it. We also observe that more surprising earnings announcements are followed by more pronounced PEAD, echoing results from markets outside of the lab.

The ability to observe fundamental values in our experimental data allows us to show that the drift manifests as a slow and incomplete adjustment of prices to values. The price adjustment is more complete in the correlated earnings setting, a finding consistent with traders devoting more attention to analyzing the implications for asset values of earnings announcements in the more complex environment of this setting. We further show that prices generally underreact to the information content of earnings announcements. The subsequent PEAD partly, but not fully, corrects the initial underreaction. Finally, we show that the greater mispricing following larger earnings surprises is indeed at least partially driven by underreaction to earnings autocorrelation.

We also study whether there are trading strategies which, by conditioning on earnings news, allow traders to earn excess returns. Because of economically significant 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 unprofitable in our markets. However, trading strategies that use limit orders to open and close positions after an earnings announcement earn positive and significant excess returns. We therefore conclude that the PEAD in our experimental markets can be profitably exploited even after accounting for transaction costs.

Our findings contribute to the literature on the PEAD in several ways. First, our result that there is significant PEAD in a setting without aggregate risk adds to the evidence that the PEAD is not a compensation for risk. Second, our article sheds light on the role of earnings autocorrelation for the emergence and strength of the PEAD. Our results show that earnings autocorrelation is not a necessary condition for PEAD, but that it affects the strength of the drift. Third, there is disagreement in the prior literature on the profitability of the PEAD on an after-transaction cost basis. We show that, in our controlled 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 observation 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, for example, Hirshleifer, Lim, and Teoh (2009), Hung, Li, and Wang (2015), and Chakrabarty, Moulton, and Wang (2022).

While we are the first to study PEAD using an experiment, we build on prior experimental literature. Our participants receive information on earnings changes in settings without and with earnings autocorrelation and should use this information to update their beliefs about the fundamental value before trading in the market. Maines and Hand (1996), Calegari and Fargher (1997), and Bloomfield, Libby, and Nelson (2003) conduct experiments in which participants forecast future earnings. All three articles conclude that participants underestimate the importance of earnings autocorrelation, possibly because they overweight past information and underweight the information provided by the current announcement (Bloomfield et al. (2003)). However, neither Maines and Hand (1996) nor Bloomfield et al. (2003) study a stock market and Calegari and Fargher (1997) conduct only a single call auction after each earnings announcement, precluding the possibility of observing a drift.

The task our traders are expected to solve is reminiscent of individual decision-making experiments in which participants have to forecast time series or, more generally, process new information.[1] Asparouhova, Hertzel, and Lemmon (2009), for example, show that short runs in a discrete random walk lead to predictions of reversal, while long runs trigger predictions of continuation. Frydman and Nave (2017) show their participants performance surprises of a firm and elicit their willingness to pay for a share of the firm. They find participants to have extrapolative beliefs. Andries, Bianchi, Huynh, and Pouget (2022) ask participants to forecast an index while observing an indicator variable that may or may not predict the index. When participants believe that the variable has no explanatory power, they tend to exhibit extrapolative beliefs and their investments tend to underreact to their forecasts. When they believe that the indicator variable is informative, they form rational beliefs and form investment decisions accordingly. Our setting, while sharing similarities with these individual decision-making experiments, differs in important ways. In contrast to all studies referenced above, our participants know the process that generates the earnings changes. This process is also very simple, particularly in our baseline setting. Furthermore (and in contrast to Andries et al. (2022)), all signals that our participants receive are informative and this is common knowledge. To the extent that the results of Andries et al. (2022) apply to our setting, we would therefore expect that our experimental participants form rational beliefs.

Our article also contributes to the broader literature on underreaction to information in financial markets. Market experiments have shown that participants, under certain conditions, underreact to the information at hand. In the market experiments of Weber and Welfens (2007), Caginalp, Porter, and Hao (2011),

[1]For a more comprehensive survey of the experimental forecasting literature than we can provide here, see Leitner and Leopold-Wildburger (2011) and Assenza, Bao, Hommes, and Massaro (2014).

https://doi.org/10.1017/S0022109023000881 Published online by Cambridge University Press

and Janssen, Li, Qiu, and Weitzel (2020), the expected value of a single asset changes halfway through a trading period. Prices in these markets do not fully adjust to the news. While the setting analyzed in these articles resembles ours in some respects, there are important differences. First, these articles study only a single news event, precluding the analysis of serially correlated announcements. Second, there is aggregate risk in their markets, implying that equilibrium prices depend on traders' risk preferences. Third, the markets in Weber and Welfens (2007) and Janssen et al. (2020) are subject to important limits to arbitrage. Specifically, traders cannot buy on margin or sell short, and each trader can have only one buy and one sell limit order outstanding at a time.[2]

Overall, we believe that bringing financial market anomalies into the controlled confines of the experimental lab can be a powerful complement to empirical approaches, particularly, if the suspected driver of the anomaly under question is a behavioral pattern that does not lend itself to ready investigation using existing financial market data.

The remainder of the article is organized as follows: Section II provides a brief summary of the relevant literature on PEAD and develops our hypotheses. Section III describes the experimental design. Section IV presents the results. Section V concludes.

II. Literature and Hypotheses

In this section, we briefly review the literature on PEAD and derive our hypotheses. Ball and Brown (1968) were the first to document that prices adjust slowly to the information contained in earnings announcements. A great number of empirical studies have subsequently confirmed the existence of PEAD.[3] While it can last across multiple earnings announcements (e.g., Bernard and Thomas (1989), Doyle, Lundholm, and Soliman (2006)), the bulk of the drift is observed between the initial announcement and the next.[4] Overall, the empirical literature provides strong evidence that PEAD exists and is economically significant. We therefore expect that a PEAD will also arise in experimental markets.

Hypothesis 1. Post-earnings-announcement drift occurs in experimental asset markets.

Several articles have advocated risk-based explanations as drivers for PEAD (e.g., Ball, Kothari, and Watts (1993), Kim and Kim (2003)), yet the changes in market betas (or, more generally, in factor exposure) observed around earnings

[2]The description of the experimental design in Caginalp et al. (2011) leaves open whether similar restrictions were in place in their experiments.

[3]See the surveys by Ball (1992), Bernard (1993), Richardson, Tuna, and Wysocki (2010), and Fink (2021). Note that there is evidence that the PEAD may have weakened recently (e.g., Richardson et al. (2010), Chordia, Subrahmanyam, and Tong (2014), and Martineau (2022)). However, this latter finding is not uncontested. Meursault, Liang, Routledge, and Scanlon (2023), for example, find strong PEAD using a text-based definition of earnings surprises in the years 2008–2019.

[4]In light of this result, we follow Hung et al. (2015), Daniel, Hirshleifer, and Sun (2020), and Martineau (2022), among others, and focus on the period between two earnings announcements in our own analyses.

announcements are insufficient to explain the drift (e.g., Bernard and Thomas (1989), Sadka (2006), Francis, Lafond, Olsson, and Schipper (2007), and Chordia, Goyal, Sadka, Sadka, and Shivakumar (2009)). The majority view today, thus, is that PEAD constitutes mispricing (e.g., Richardson et al. (2010), Hung et al. (2015), and Daniel et al. (2020)). We therefore also concentrate on mispricing-based explanations. To this end, we design our experiment such that there is no aggregate risk in the market. Equilibrium prices then equal expected payoffs. Consequently, if a PEAD arises in our markets, we can rule out that it is a compensation for risk.

Mispricing arises in the context of earnings announcements when investors underreact to the announced information. This underreaction may be caused by i) a misspecified model to forecast earnings (e.g., Bernard and Thomas (1989), Freeman and Tse (1989), and Bernard and Thomas (1990)), ii) behavioral biases (e.g., Frazzini (2006)), and/or iii) inattention (e.g., DellaVigna and Pollet (2009), Hirshleifer et al. (2009), Hou, Xiong, and Peng (2009), and Hung et al. (2015)). One of the most prominent explanations combines all of these and is based on the observation that quarterly earnings are positively serially correlated (see, e.g., Rendleman, Jones, and Latane (1987), Bernard and Thomas (1989), p. 9, and Kothari, Lewellen, and Warner (2006)). According to this explanation, 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.

Several studies find empirical support for the earnings autocorrelation hypothesis (e.g., Ball and Bartov (1996), Rangan and Sloan (1998), Battalio and Mendenhall (2005), Bathke Jr, Lorek, and Willinger (2006), Bathke, Morton, Notbohm, and Zhang (2014), and Bhattacharya, Olsson, and Park (2020)). However, whether earnings autocorrelation is the root cause of the drift, or rather an accelerator that affects its strength, is not entirely clear. One impediment to answering this question is that, in 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 control the earnings process and thus can “switch on and off” earnings autocorrelation. We can therefore directly test the extent to which earnings autocorrelation causes the drift or rather only amplifies a drift that would also occur if earnings were serially uncorrelated. We thus formally test the following hypothesis:

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

The literature has documented a systematic, positive relation between the magnitude of the earnings surprise and the strength of the post-announcement drift (e.g., Bernard and Thomas (1989)). Accordingly, we design our correlated earnings setting to have two types of announcements, which trigger very different changes in the fundamental value of the asset. These two types allow us to test the following hypothesis:

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

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

Hypothesis 4. The observed post-earnings-announcement drift can be exploited to earn excess returns even after accounting for transaction costs.

As discussed above, the consensus view today is that PEAD is a misvaluation phenomenon caused by underreaction to the information content of earnings announcements. The reasoning goes as follows: The announcement 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 and its changes revealed by the earnings announcement are unobservable. Most recently, Jiang, Li, and Wang (2021) have documented that the phenomenon of underreaction to news, followed by a significant drift in the direction implied by the news, is not confined to earnings announcements but rather extends to other firm news. Their findings highlight the importance of additional research into the PEAD and related phenomena.

In the experiments we conduct, we know the fundamental value of the stocks, and we know how this value changes upon an earnings announcement. We can therefore directly test the following hypothesis:

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

As noted above, a prominent explanation of the PEAD states that investors fail to fully account for the information content of earnings announcements when earnings are serially correlated. We revisit the question of the (mis)interpretation of autocorrelated earnings news. When earnings changes are positively serially correlated, there are unsurprising announcements (an earnings increase following a previous earnings increase, $++$, or a decrease following a decrease, $--$) and surprising announcements (a decrease following an increase, $+-$, or an increase following a decrease, $-+$). The surprising earnings changes trigger a greater change in the fundamental value of the asset (a decrease after an increase predicts that future earnings changes are more likely to be negative, whereas prior to the announcement, they had been more likely to be positive). If investors fail to fully account for earnings autocorrelation, they will overestimate the implications for the asset's fundamental value of *unsurprising* 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, thus, reinforcing the tendency for prices to only partially adjust to changes in fundamental value. We thus test the following hypothesis.

Hypothesis 6. In the presence of earnings autocorrelation, prices adjust more fully to the information content of the earnings announcement after unsurprising than after surprising announcements.

Note that our hypotheses serve a dual purpose. While ours is an approach that is well-established in other areas of finance, it has not hitherto been applied to studying PEAD. We thus feel the necessity to document the ability of experimental markets to replicate well-known characteristics of PEAD from the field. In this sense, testing Hypotheses 1, 3, and 4 does little to advance our state of knowledge about PEAD. We nevertheless test them to establish the suitability of our research method to study the PEAD and related phenomena. By contrast, studying Hypotheses 2, 5, and 6 generates new insights into the PEAD phenomenon, which previous studies were not designed to offer with the same clarity, or at all. Testing these latter three hypotheses thus does help us better understand the puzzle that is PEAD.

III. Experimental Design

Our experiment was run from May 2019 to June 2020 in the WULABS at the Vienna University of Economics and Business, in the mLab at the University of Mannheim, in the Experimental lab of the WiSo Faculty at the University of Hamburg, and in the MaxJungLab at the University of Graz, using bachelor, master, and PhD students of economics and business as participants.[5,6] The experiment consists of two experimental designs, referred to as "treatments," which we name Base and Corr. Each of the 20 sessions we run employs exactly one of the two (10 Base and 10 Corr), and every one of our 238 participants is thus exposed to one treatment only ("between-participants" design). The participants are compensated by cash payments tied to their performance, paid at the end of the experiment. In each experimental session, 11 to 12 participants form a cohort and interact over a sequence of four periods. All asset and cash balances are reset between periods.[7]

[5]We follow the example of Maines and Hand (1996) and Calegari and Fargher (1997) and use only students of economics, business, and related programs (financial mathematics, information technology for business, business education) for this study, requiring us to conduct our experiments in more than one lab. We spread the two treatments evenly across labs, always conducting the same (±1) number of sessions of each treatment in any given lab. While we cannot test for lab fixed effects due to insufficient sample size per lab, this procedure ensures that between-treatment comparisons remain unaffected by such lab effects.

[6]In 2 of our 20 sessions, insufficient participant numbers required us to use 11 instead of 12 participants.

[7]Note that some articles that study thin, complex markets find that such markets may fail to reliably converge to the competitive equilibrium. Bossaerts and Plott (2002), for example, report incomplete convergence to competitive equilibrium outcomes when they study the CAPM in the lab. However, their markets are considerably more complex than ours, in that traders trade three securities, face aggregate risk at the market level, and have to use a trading interface that allows for multi-unit trading, yet does not

In each session, after the experimenter has checked their IDs and welcomed them to the lab, the participants are randomly assigned to computers. Following the best practices laid out in Freeman, Kimbrough, Petersen, and Tong (2018), we report that all participants then receive the same written instructions, providing information on the trading interface.[8] The experimenter reads the instructions out aloud while participants follow along to create common knowledge of their contents (i.e., to ensure participants know the contents, know that all other participants also learnt the same contents, etc.). Afterward, the trading mechanism and the most important screens are explained in detail, followed by a trial period to allow participants 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, participants answer control questions to ensure their understanding of the instructions before the first trading period commences.[9] After the fourth trading period has ended, one period is randomly chosen for payout. Participants then complete a post-experiment questionnaire, are paid in private, in cash, and leave.

A. Trading Environment

In all trading periods, participants simultaneously trade shares of two fictitious companies, firms A and B, in 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)). Figure B.1 in Appendix B shows the trading interface. Each participant starts with 900 talers (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.[10] Traders can submit any combination of limit and market orders in the markets for stocks A and B.[11] Each order is for 1 share. The order book is empty at the beginning of a period, and it is anonymous, that is, the identity of the trader submitting an order is not displayed

employ an order book. In contrast, our own traders trade only two securities, face no aggregate risk, and use a simpler trading interface with a fixed single unit order size and an open order book. Despite the greater complexity of their markets, Bossaerts and Plott (2002) had an average of only 9.7 traders (with a minimum of 5), while we had an average of 11.9 (with a minimum of 11). Furthermore, the same authors assert that markets with 8–12 participants that involve “at most two goods/securities” are acknowledged to generate convergence toward competitive equilibrium in the laboratory (Bossaerts and Plott (2004), p. 136). Similarly, Plott and Sunder (1982), Smith (1982), and Friedman, Harrison, and Salmon (1984) highlight that efficient outcomes in limit order markets require only few traders, and Palan (2013) concludes from his review of the issue that experimental results in these less complex markets are unaffected by market size.

[8]See Section C of the Supplementary Material for a copy of the experimental instructions.

[9]If participants had questions, the experimenter pointed them to the section of the instructions explaining the topics in question. The experiment commenced once all participants had correctly answered all questions. See Appendix A for evidence on participant comprehension of the instructions.

[10]In the markets with only 11 participants, 5 were endowed with 9 shares of firm A and 900 talers, 5 were endowed with 9 shares of firm B and 900 talers, and 1 was endowed with 5 shares each of firms A and B and with 1,000 talers. This ensured an absence of risk at the market-level and a constant cash/asset ratio even in these slightly understaffed sessions.

[11]We will use the terms “participant” and “trader” interchangeably.

FIGURE 1

Structure of a Trading Period

Figure 1 shows each of the 4 900-s trading periods in a session, structured into 5 180-s phases, separated by 4 inter-phase earnings announcements, and concluded by 1 closing earnings announcement.

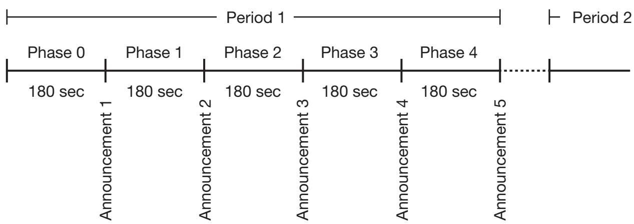

(traders' own orders are flagged, though). Order execution is governed by price and time priority, following the algorithm of NASDAQ (2019). Traders can cancel their own unexecuted orders without cost and at any time. Unless canceled, unexecuted orders remain in the order book until the end of the period. Taler holdings pay no interest and there are no transaction costs beyond the (endogenous) bid–ask spread. Traders can sell short up to 9 A and 9 B shares. Similarly, traders can buy on margin for up to 900 talers.

At the beginning of every period, participants are informed about both firms' earnings per share. During the period, there are 4 announcements where participants receive updated earnings information. After the end of the trading period, there is a fifth announcement. Participants trade continuously, without interruption or reset, throughout 5 "phases" of equal length (180 s) in each period. Phase 0 extends from the start of the period to the first earnings announcement, Phase 1 follows the first announcement, etc. (see Figure 1).

B. Earnings Announcements

The 4 trading periods last 900 s each, with announcements after 180 s, 360 s, 540 s, and 720 s. A final announcement after the end of the period (i.e., after 900 s) ensures that share values remain stochastic even in the last phase of each period. The shares are bought back by the experimenter after the end of the period for the fundamental value (F) of 20 times earnings after the fifth announcement (mimicking perpetual discounting at a rate of 5%).

The trading interface counts down to the upcoming earnings announcement (see Figure B.1 in Appendix B). At the announcement, the firms' updated earnings are shown on screen (highlighted by a flashing red box) while trading continues. Both firms' initial earnings per share are 5 talers. In each announcement, the earnings for a firm can increase or decrease by 0.5 talers:

$$\text{(1)} \qquad \Delta E_\tau^\theta \in \{\delta^- = -0.5,\ \delta^+ = 0.5\} \forall \tau,$$

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

FIGURE 2

Illustration of the Possible Earnings Trajectories Within a Period in Treatment BASE

The first number in each box in Figure 2 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 (numbers in parentheses were communicated to participants and numbers 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.

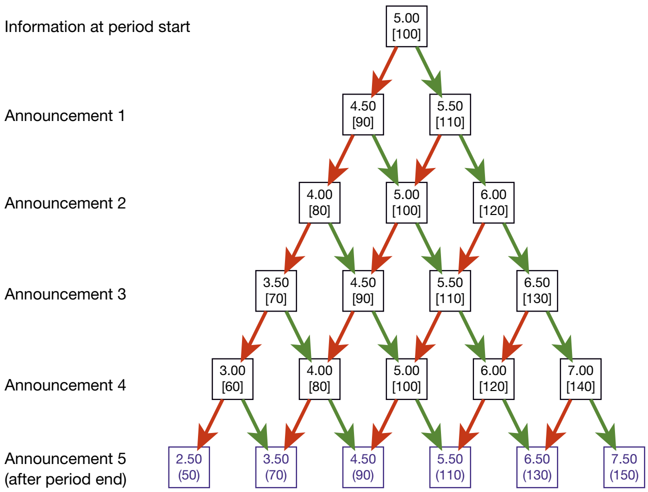

Our two treatments differ in the dependence structure of successive earnings changes. In treatment BASE, positive and negative earnings changes are equally likely:

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

where p is the probability operator. Successive earnings changes are independent (Markov chain of order 1). The expected earnings change is thus 0 and the earnings surprise (i.e., the unexpected component of the announcement) equals the earnings change. Earnings surprises are serially uncorrelated. Earnings and F in treatment BASE thus follow a recombining binomial tree over a period (see Figure 2). The figure lists, for each announcement, the possible earnings E and the corresponding fundamental values F. Participants in the experiment saw the same tree as in Figure 2 except for the F values in square brackets.

In our second treatment, CORR, earnings surprises are autocorrelated. Specifically, a change in earnings is followed by another change in the same direction $(++/--)$ with probability 0.75 and by a change in the opposite direction $(+-/-+)$ with probability 0.25 (Markov chain of order 2). Formally,

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

The autocorrelation of both firms' earnings is thus 0.5.[12,13] The earnings process in treatment CORR implies that earnings as well as F follow a non-recombining binomial tree. This tree is illustrated in Figure 3. Again the figure lists, for each announcement, the possible earnings E and the corresponding fundamental values F. Participants in the experiment saw the same tree as in Figure 3, except for the F values in square brackets.

Each announcement in treatment BASE only contains information about current earnings, while each announcement in treatment CORR contains information about current earnings and, implicitly, about the distribution of subsequent earnings announcements. Surprising announcements in treatment CORR thus have greater impact on F than the announcements in treatment BASE, whereas unsurprising announcements in CORR have smaller impact. While fundamental value changes by ± 10 talers in every announcement in treatment BASE (Figure 2), the change varies between ± 5 and ± 28.125 talers in treatment CORR (Figure 3).

As noted previously, we focus on mispricing-based explanations for PEAD and therefore design experimental markets without aggregate risk. We do so by establishing perfectly negative correlation between the earnings changes of firms A and B. If the earnings of firm A are announced to have increased, those of firm B are announced to have decreased, and vice versa. Earnings changes thus always have the same magnitude (0.5 talers) but opposing signs. The perfect, negative correlation of earnings changes also implies perfectly negative correlation of fundamental value changes. All of this is public information. Any portfolio consisting of equal numbers of A and of B shares thus is risk-free and the fundamental values of 1 A and 1 B share always add up to 200, a fact all participants were explicitly made aware of in the instructions. The total number of A shares equals that of B shares, such that there is no aggregate risk. Equilibrium with risk averse agents then implies that each trader holds a balanced portfolio, and the equilibrium risk premium is 0.[14] Consequently, in equilibrium, prices should equal expected values. At the same time, because each participant is endowed with *only* A or *only* B shares, participants have an incentive to trade in order to equate their holdings of A and B shares. Our design is thus not subject to

[12]The dynamics of earnings announcements found in empirical data is more complex than the simple structure we assume in our experiments. In particular, earnings have been found to be positively correlated in adjacent quarters but negatively correlated 4 quarters apart (e.g., Bernard and Thomas (1990)).

[13]See Section B.1 of the Supplementary Material for evidence that the random draws conformed to this distribution.

[14]A reviewer proposed an alternative design that would have retained traders' endowments and the two-asset structure with perfectly negatively correlated earnings changes, but would have permitted trading in only one asset while rendering the other non-tradeable. This design results in clear predictions for portfolio adjustments among risk-averse traders, because traders endowed with only the tradeable asset should sell all of their holdings, whereas traders endowed only with non-tradeable assets should buy assets to equate the number of tradeable and non-tradeable assets in their portfolios and thus eliminate price risk. In addition, with only one traded asset, the scope for arbitrage opportunities is greatly reduced. On the other hand, the alternative design yields fewer observations of trading activity. Furthermore, the existence of non-tradeable assets in equity portfolios does not correspond to participants' real-life experience and may therefore negatively affect the external validity of the experimental results. See Section A of the Supplementary Material for a more detailed discussion of the alternative design.

FIGURE 3

Illustration of the Possible Earnings Trajectories Within a Period in Treatment Corr

In Figure 3, 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 (numbers in parentheses were communicated to participants and numbers 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.

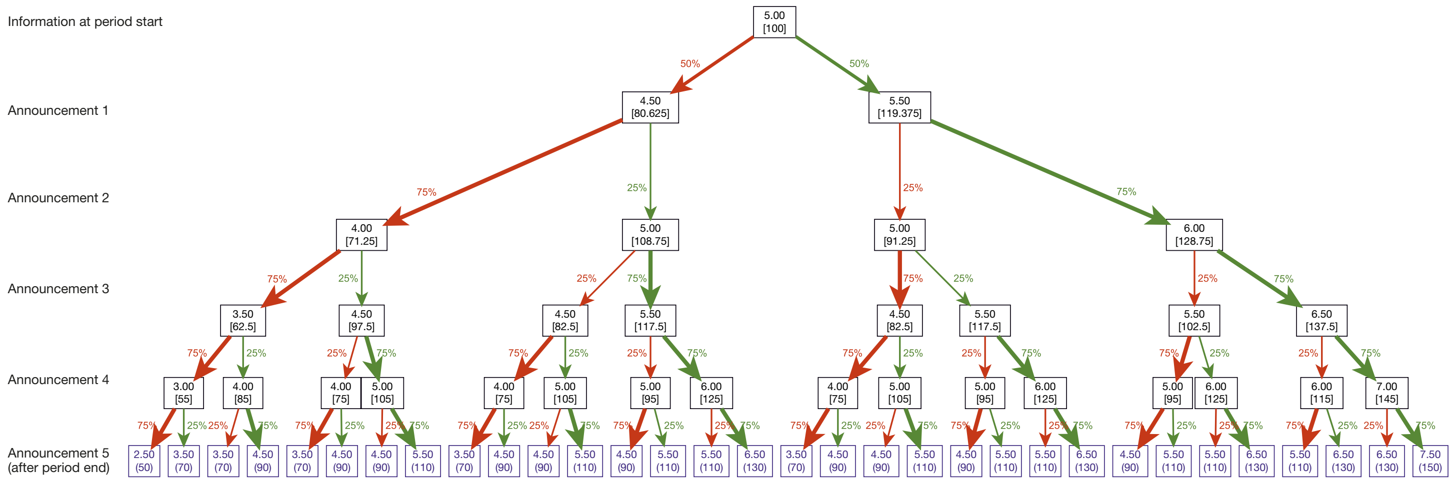

https://doi.org/10.1017/S0022109023000881 Published online by Cambridge University Press

TABLE 1
Overview of Participant Payoff in Euros

Table 1 presents information about participant payoffs by treatment, including mean payoff, payoff standard deviation (Std. Dev., within-session average), minimum (Min), and maximum (Max) payoff.

Treatment	Mean	Std. Dev.	Min	Max
BASE	24.8	2.6	11.4	30.8
CORR	25.0	5.0	9.1	47.1

the *no trade* criticism frequently leveled at other experimental asset market studies.[15]

The aggregate endowment of N participants is $\frac{N}{2} \times 9$ shares of each firm, for a total market value of $\frac{N}{2} \times 9 \times 200 = 900N$. Since the market's total cash endowment is also $900N$, the cash/asset ratio equals 1, thus enabling trading while avoiding cash endowment effects.[16]

C. Participant Payment

Participants' payoffs are based on their end-of-period wealth, $W_{p,T}$, in a period that is randomly chosen using a physical randomization device. They are 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 participant i's final balance of shares of firm f's stock, and c_i are the participant's final cash (taler) holdings. (We suppress the p and τ indices of W, n, E, and c for notational simplicity.)

Payoffs are calculated by converting W_i using an exchange rate of 100 talers = € 1. Participants further receive a base compensation of € 5 to € 8, depending on the rules of the lab a given session was run in, which reflects differences in wage and price levels between cities. Total earnings are bounded from below by € 0, but no participant went bankrupt. In total, participants earn an average of 1,800 talers (€ 18) from trading plus the base compensation for an experiment lasting approximately 2 h. Table 1 reports the actual payments from the experiment. The average participant earns around € 25.

IV. Results

We start with a short description of the trading activity in our experimental markets.[17] As Table 2 documents, we see substantial activity. The average trader

[15]See Milgrom and Stokey (1982) and Tirole (1982), the two seminal articles on *no trade* in capital markets, and Kleinlercher and Stöckl (2021), who study deviations from *no trade* in experimental markets.

[16]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 and the margin buying capacity are also symmetric, and thus do not distort the cash/asset ratio.

[17]We analyze our data using *R* (R Core Team (2017)), reading it in using package *ztree* (Kirchkamp (2019)) and generating regression tables using *stargazer* (Hlavac (2018)) and *texreg* (Leifeld (2013)). For the remaining tables, we use *kableExtra* (Zhu (2019)), for the figures *ggplot2* (Wickham (2016)).

TABLE 2
Overview of Trading Activity

"Actions" in Table 2 include order submissions, order cancelations, and the acceptance of outstanding orders.

	Actions per Period				Trades per Period			
Treatment	Mean	Std. Dev.	Min	Max	Mean	Std. Dev.	Min	Max
Base	628	186	349	1,156	179	74	89	468
Corr	578	148	352	870	170	57	86	318

engages in 50.8 actions (submitting a limit order, canceling an order, or accepting another trader's order) and is involved in 29.3 transactions per period, corresponding to 1 action every 17.7 s and 1 transaction every 30.7 s. The figures in Table 2 also suggest that there is more trading activity in the Base than in the Corr treatment. Both the number of actions and the number of transactions per period are higher by 5%–10% in Base, yet the differences are not significant (actions: Welch two-sample $t(74.234) = 1.3262$, $p = 0.1888$; transactions: $t(73.282) = 0.6153$, $p = 0.5403$).

The traders in our experiments also react as expected to the incentives they face. While the literature shows that investors typically hold non-optimal portfolios (see, e.g., Ackert, Church, and Qi (2016) and the references therein), the majority of our traders nevertheless reduce their exposure to price risk by reducing the absolute difference between their firm A and firm B shareholdings. By the end of Phase 0, 63.9% of all traders have reduced this difference, a proportion that increases to 71.0% by the end of Phase 4. In fact, 2.4% of all traders fully equalize their firm A and firm B share holdings by the end of Phase 0 and 7.1% do so by the end of Phase 4.[18] At the treatment level, the proportions of traders who have reduced their exposure by the end of Phase 0 [Phase 4] are 64.6 [70.6]% (Base) and 63.1 [71.5]% (Corr). The proportions of traders who completely equalized their shareholdings by the end of Phase 0 [Phase 4] are 2.3 [8.1]% (Base) and 2.5 [6.3]% (Corr).

A. Existence of PEAD

To analyze the price dynamics in our markets, we subdivide each period into 10-s *windows*. We base our analysis on quote midpoints (but refer to them as "prices" for simplicity) to eliminate noise resulting from bid–ask bounce. The analysis thus (and in line with most prior empirical research on the PEAD) ignores transaction costs. However, we account for transaction costs in the form of bid–ask spreads in Section IV.C, where we analyze the profitability of trading strategies aiming to exploit the PEAD. This analysis is based on executable bid and ask prices and thus explicitly accounts for the implementation cost. We further note that the analyses in Sections IV.A–IV.C mimic analyses using field data in that they consider *price changes* following earnings announcements (excluding the immediate

[18]Nevertheless, common knowledge of rationality would imply that all (risk-averse) participants only move to balance their portfolios and then cease trading. In line with the findings from all other empirical and experimental asset market research that we are aware of, the traders in our experiment do not follow this strict theoretical prediction, likely because common knowledge of rationality is absent.

price reaction to the announcement). In Section IV.D, we relate the price change to the *change in fundamental value* caused by the earnings announcement, utilizing the advantages afforded us by our experimental approach by going beyond what can be analyzed with field data.

Figure 4 plots cumulative taler price changes relative to the quote midpoint at the time of an announcement. The upper [lower] line in each diagram tracks changes after positive [negative] earnings surprises. The dashed horizontal lines represent the price levels reached by the end of the announcement window (i.e., the window starting at the time of the announcement). These horizontal lines thus capture the initial price reaction to the announcement.

FIGURE 4

Price Changes in Talers Relative to the Quote Midpoint at the Time of the Announcement

Figure 4 shows the average price changes relative to the quote midpoint at the time of the announcement, using the closing quote midpoint for each 10-s window following the announcement. Graph A plots results for the pooled data from all treatments; Graph B reports results separately for treatments BASE (left) and CORR (right); Graph C plots only CORR data and reports results separately for unsurprising (left) and surprising (right) earnings news, using only data from Phases 2 to 4. The blue, upward-trending [orange, downward-trending] lines plot 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 10-s announcement window.

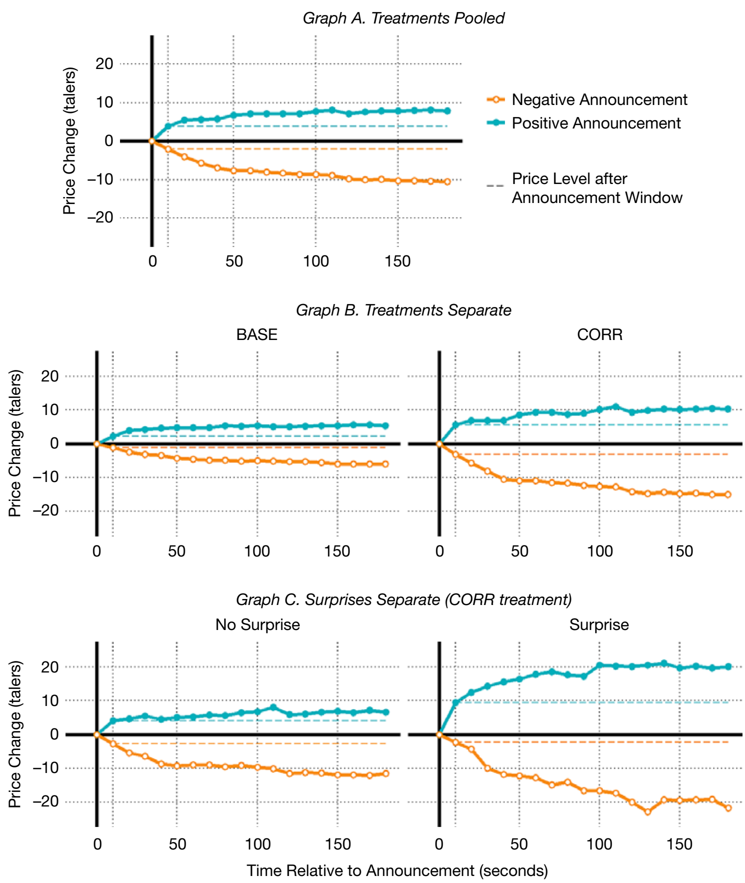

https://doi.org/10.1017/S0022109023000881 Published online by Cambridge University Press

TABLE 3

Post-Earnings-Announcement Drift

In Table 3, we show the mean taler gains based on quote midpoints from the end of the announcement window until the end of the phase. We report *t*-statistics in parentheses. Values based on all phases where both the first and the last 10-s windows following an announcement (windows 0 and 17) have both a bid and an ask, thus permitting us to calculate quote midpoints (includes 95%, 87.5%, and 85% of long, short, and long–short phases, respectively). *, **, and *** indicate significance at the 5%, 1%, and 0.5% levels, respectively.

	Treatments				CORR Treatment[a]		
	All	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)

[a] CORR treatment data separated by the surprise variable excludes the phase following the first announcement (Phase 1).

Graph A of Figure 4 pools results across 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 upward by another 3.80 talers in the 170 s following the announcement window. The drift is even more pronounced following negative earnings surprises (prices drop by 2.10 talers in the announcement window and then drift downward by another 7.88 talers over the remainder of the phase).[19,20] This implies that a strategy, initiated at the end of the announcement window, that buys stocks with positive earnings surprises and shorts stocks with negative surprises, earns 11.71 talers (before transaction costs).

Table 3 reports the results of *t*-tests.[21,22] The figures in the first column indicate that the profits of both a "long" strategy that buys the stocks with positive announcements and a "short" strategy that sells the stocks with negative announcements are positive and statistically significant. The profit of the combined "long–short

[19]Consistent with, for example, Louhichi (2008), Nam, Wang, and Zhang (2008), and the seminal review by Karpoff (1987), traders react more actively to positive news than to negative news. Our results thus lend support to findings of greater drift following negative than following positive earnings news, as reported by Booth, Kallunki, and Martikainen (1997), Truong ((2010), (2011)), and Sun (2015) (while contradicting Doyle et al. (2006), who report the opposite).

[20]We checked whether short selling restrictions drive our results. In treatment BASE [CORR], there are 23 out of 1,888 [9 out of 1,920] participant × phase observations where, at the end of the announcement window, a trader has reached the shorting capacity of 9 shares of the stock with negative earnings news and is unable to sell. These phases do not differ noticeably from others in terms of PEAD. Yet participants are quite willing to enter into short positions. In treatment BASE, 92 out of 118 participants (78.0%), and in treatment CORR, 99 out of 120 participants (82.5%) held a short position at some point during the experiment.

[21]In light of Benjamin et al. (2018), we highlight significance at the 0.005 level alongside the more conventional 0.05 and 0.01 levels throughout the article.

[22]As a robustness check, we repeat the analysis reported in Table 3 using window lengths of 5 s and 15 s instead of 10 s. We report the (qualitatively unchanged) results in Section B.2 of the Supplementary Material.

strategy" is also highly significant, providing clear evidence of PEAD. Our results thus support Hypothesis 1.[23]

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

B. PEAD and Earnings Autocorrelation

To test whether earnings autocorrelation drives PEAD, we next analyze each treatment separately. Graph B of Figure 4 presents the results. There is clear evidence of PEAD in both treatments, yet the drift is much more pronounced in the CORR treatment, as confirmed by columns 2 and 3 of Table 3. After positive earnings news (and excluding the announcement window) prices drift upward by 3.17 talers in treatment BASE and by 4.42 talers in treatment CORR. Both values are highly significant. Similarly, after negative earnings news prices drift downward by 4.68 talers in treatment BASE and by 11.36 talers in treatment CORR. 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 for the short and the combined long–short strategies.

Table 4 reports regressions studying the dynamics of window-to-window taler price changes. We exclude the announcement window because it captures the initial price reaction to the announcement rather than PEAD. We furthermore multiply all price changes after negative earnings surprises by (-1) so we can pool the data from 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. We re-base the count variables for the first period, phase, and window from 1 to 0 in order to make the constant interpretable. The period and phase variables capture changes in the strength of the PEAD across the periods of a session and the phases within a period so we can tell whether the PEAD diminishes with trader experience. The coefficient on the window variable lets us assess the dynamics of the drift following an announcement. Since Figure 4 suggests that the strength of the drift decreases non-linearly, we also include the square of the window number to allow for such a functional form. Finally, we include a dummy variable identifying phases following a positive earnings surprise to test whether the drift differs following positive and following negative surprises.

We start by discussing the first column of Table 4.[24] 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 treatment BASE. The constant is positive and significantly different from 0, implying that there is drift in

[23]We provide further evidence supporting Hypothesis 1 in the regression analysis discussed in Section IV.B.

[24]Alternative specifications with period × phase × window fixed effects or using log returns instead of taler changes and using PEAD over the entire phase instead of per window yield similar results (reported in Table SM.6 in Section B.3 of the Supplementary Material).

TABLE 4

Regression Analysis of Window-to-Window Changes in Taler Closing Quote Midpoints

Table 4 presents 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). 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 treatment Corr. CORRELATED_SURPRISE is a dummy for an earnings change carrying the opposite sign as in the previous announcement in treatment Corr. CORRELATED_FIRST_ANNOUNCEMENT is a dummy for returns stemming from the phase following the first announcement in treatment Corr (which is neither unambiguously surprising nor unsurprising). POSITIVE_EARNINGS_CHANGE is a dummy variable for a positive earnings change. 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). 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). *, **, and *** indicate significance at the 5%, 1%, and 0.5% levels, respectively. Standard errors, clustered at the session level, in parentheses.

	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
No. of obs.	10,389	10,389
RMSE	5.578	5.576

the direction of the earnings announcement. This drift is weaker after positive announcements, as 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 of 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 decreases over time, but at a decreasing rate. The coefficients for the period and the phase are small and not significant, implying little variation in the drift over time.[25]

[25]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 we detect explain only a small fraction of the overall return variability.

Scrutinizing our data, we find that arbitrage opportunities may have interfered with price adjustment in our markets, thus contributing to the PEAD we observe. To rule this out, we run another 10 sessions under a design mirroring that in Corr, but with an algorithmic arbitrageur that eliminates arbitrage opportunities. Our findings remain virtually unchanged in these markets (see Appendix D). To rule out noise or spurious results caused by a lack of participant experience at the beginning of each experimental session, we repeat our analyses using only the data from periods 2 to 4. We discuss the resulting figures and tables (corresponding to Figures 4–6 and Tables 3–6 in the body of the article) in Section B.7 of the Supplementary Material. All results remain qualitatively unchanged. We conclude that PEAD is a robust phenomenon in our experimental markets.

The results support our earlier finding of persistent PEAD in our markets and clearly reject Hypothesis 2. Earnings autocorrelation is not a necessary condition for the occurrence of PEAD. Rather, it strengthens PEAD compared to a situation without autocorrelation.

Result Hypothesis 2. Earnings autocorrelation is not a necessary condition for PEAD, yet the drift is significantly more pronounced in its presence.

The positively autocorrelated earnings changes in treatment Corr imply 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 than a change in the opposite direction ($+-/-+$). We will label the two types of announcements "unsurprising" and "surprising," respectively. The first announcement in treatment Corr is neither unambiguously surprising nor unsurprising because earnings are as likely to increase as they are to decrease, and because the absolute size of the impact on the fundamental value is independent of the sign of the earnings surprise. We therefore exclude the first announcement from all analyses that distinguish between surprising and unsurprising announcements.

Graph C of Figure 4 plots results for surprising and unsurprising announcements. It shows greater PEAD following surprising announcements, in line with the empirical literature that documents stronger drift after larger surprises (e.g., Bernard and Thomas (1989)). Our previous finding that the drift is stronger after negative announcements continues to hold, as confirmed by the last 3 columns of Table 3. The drift is statistically significant after both surprising and unsurprising announcements, but is more pronounced after the former. The differences are significant for the long strategy and for the combined long–short strategy.

Column 2 of Table 4 presents the results of a regression that includes three separate dummy variables to identify phases in treatment Corr with unsurprising announcements, surprising announcements, and phases following the first announcement. The results confirm our previous findings, with PEAD occurring after both unsurprising and surprising announcements, but significantly more pronounced after the latter. These results support Hypothesis 3.

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

The results so far have important implications. They stem from a well-controlled laboratory experimental setting ruling out several of the factors conjectured to cause PEAD, such as informational asymmetries or changes in the riskiness of a stock that are related to earnings news. We can thus focus on the question of whether 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.[26]

C. Trading Strategies

The results so far clearly document the existence of PEAD in our markets and show that trading strategies can profitably exploit the PEAD before transaction costs. We now analyze whether the same is true when transaction costs in the form of bid–ask spreads are taken into account. Figure 5 shows the evolution of quoted bid and ask prices over time, averaged over all periods of all sessions in treatments

FIGURE 5

Development of Spreads over Time

Figure 5 presents the average difference between best ask quote and F (red), and between best bid quote and F (blue), in talers, over the trading period. Solid lines plot treatment BASE data, dashed lines treatment CORR data.

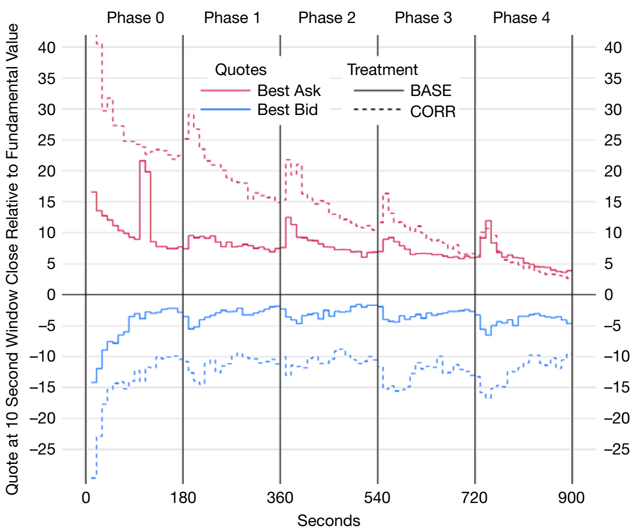

[26]Frazzini (2006) hypothesizes that PEAD may be caused by the disposition effect. We perform a battery of tests (e.g., studying the probability of net buyers in one phase to sell following a subsequent, negative announcement; studying the number of new sell and buy order submissions following positive and negative announcements; studying the number of new sell order submissions of prior phase net buyers; studying the probability of selling winners/losers following positive/negative announcements in line with Odean (1998); and studying the probability of selling over time using the hazard function methodology of Feng and Seasholes (2005)) but find no evidence that the disposition effect drives the results in our experimental markets.

BASE and CORR. To reveal pricing asymmetries, we center bid and ask prices on the fundamental value of the stock, not on the quote midpoint.

We detect five patterns in the Figure 5. First, spreads are generally wide, making trading using market orders expensive. Average spreads per 10-s window range from less than 5 to more than 30 talers. This is substantial, in particular, when compared to the change in the fundamental value caused by an earnings announcement.[27] Second, spreads tend to decline over the period. They are largest in Phase 0 and lower in later phases. The largest spreads (immediately after the start of a period) are due to the order book being empty at the beginning of Phase 0. Orders submitted early, far from the fundamental value, can thus establish a wide spread that soon narrows as more orders arrive. Third, spreads tend to widen after an announcement and then decline during the remainder of the phase. The greatest contribution to the widening of the spread following an announcement stems from an increase in limit order executions, followed by a reduction in new limit order submissions, and, finally, an increase in limit order cancelations.[28] Furthermore (and unsurprisingly), the spread widens asymmetrically. Limit orders in the book constitute free trading options which traders exploit by trading against those limit orders that are "in the money" after an announcement (i.e., against bids higher than F and asks lower than F). Thus, we observe more transactions against, and cancelations of, limit asks [bids] following positive [negative] earnings changes. The reduction in new limit order submissions, in contrast, is roughly symmetrical. A possible explanation for the strong and symmetric decline in limit order submissions is that, after an announcement, traders focus on trading against stale limit orders in the book or on canceling their own stale limit orders and submit fewer new limit orders. Fourth, spreads are consistently wider in treatment CORR than in BASE, in line with the greater variability of both fundamental values (standard deviation 25.2 vs. 12.5 talers) and prices (standard deviation 23.7 vs. 11.9 talers) in treatment CORR. Fifth, ask prices tend to be farther from the fundamental value than bid prices, suggesting slight overpricing. We study this phenomenon in more detail in Section IV.D. A regression analysis confirms the visual impressions from Figure 5. For details on model specification and results, see Section B.4 of the Supplementary Material.

To analyze whether the PEAD in our markets can be profitably exploited, we analyze two trading strategies. Our calculations use the actual tick-by-tick order book data from our experiment. The only assumption we make is that the opening trades necessary to implement the proposed strategies do not affect the prices at which the positions can be unwound at the end of the phase.

[27]This change is 10 talers in treatment BASE and can range up to 28.125 talers after surprising announcements in treatment CORR. The spreads in our 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 an 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.

[28]Transactions outnumber cancelations 6:1 in the announcement window. Furthermore, transactions outnumber spread-widening cancelations (i.e., cancelations of the best bid or ask) approximately 30:1. These factors reduce the number of limit orders in the book and are accompanied by a reduction in new limit order submissions by more than half compared to the average over all windows.

TABLE 5

Trading Strategy Returns per Phase

Table 5 shows the average single-phase aggregated log returns (in percent) to the long and short legs of two trading strategies. "Delay" is the time (in seconds) between the earnings announcement and the opening of the position. Standard errors in parentheses. *, **, and *** indicate significance at the 5%, 1%, and 0.5% levels, respectively.

	BASE		CORR	
Delay	Mean	SE	Mean	SE
Strategy *Market Orders*	Open at best bid/ask, close at best bid/ask.			
0	−5.33***	(0.52)	−17.53***	(2.59)
2	−5.78***	(0.50)	−17.65***	(2.53)
4	−7.44***	(0.61)	−19.41***	(2.55)
6	−8.29***	(0.60)	−21.10***	(2.52)
8	−8.82***	(0.64)	−22.98***	(2.53)
10	−9.24***	(0.65)	−24.68***	(2.55)
Strategy *Limit Orders*	Open using limit order at quote midpoint, close using limit order at quote midpoint after 90 s or market order at 170 s.			
0	2.20***	(0.29)	3.51***	(0.60)
2	1.83***	(0.29)	2.93***	(0.66)
4	1.75***	(0.28)	2.94***	(0.54)
6	1.39***	(0.24)	2.26***	(0.55)
8	1.22***	(0.26)	1.85***	(0.54)
10	1.05***	(0.25)	1.41*	(0.55)

The first trading strategy (Strategy *Market Orders*) consists of buying a share at the best ask price following 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).

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

The Strategy *Market Orders* panel of Table 5 reports separate profitability results for treatment BASE and for unsurprising and surprising announcements (pooled) in treatment CORR.[29] The table lists log changes (in percent) in the value of a given position from its opening to its close, relating the profit or loss to the share price paid or obtained when opening the position. The 6 lines in 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 that it is made 2 s after the announcement, and so on.

We find that a strategy based on market orders is unprofitable. Due to the bid–ask spread, the strategy yields significantly negative returns throughout. Furthermore, the longer the delay between announcement and initial trade, the more unprofitable the strategy becomes. This is due to i) prices drifting in the direction of the announcement, thus diminishing the potential profit available from trading in the direction of the announcement, and ii) spreads widening in the 10 s following the announcement, thus increasing the transaction costs.

[29]We again exclude the results for the phase following the first announcement in treatment CORR from the analysis. A separate analysis of this phase (results not shown) reveals that profits are negative.

However, traders are not restricted to using market orders, which force them to pay the full bid–ask spread (half each in the opening and closing trades). Instead, they can trade using limit orders, which potentially offer better prices. Their downside is that they introduce non-execution risk (a downside that we account for in our analysis). We propose the following strategy (Strategy *Limit Orders*): 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 until halfway through the trading phase (90 s), the trader cancels the order and we record the profit as 0. If the limit order executes, the trader submits a second limit order to close the position 90 s after the start of the phase, again with a price limit equal to the then current midpoint. If this order is executed at any time up until 170 s after the announcement, we calculate and record the profit or loss. If it is not, we assume that the trader cancels the limit order 10 s before the subsequent announcement and closes the position at the then current best bid or ask price using a market order.

Strategy *Limit Orders.* 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 90 s of the phase, cancel it and do nothing else. If it gets executed, wait until 90 s have elapsed since the announcement and submit a limit sell order, priced at the quote midpoint. If this limit order gets executed within the next 80 s, do nothing else. If it does not get executed until 170 s 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 Strategy *Limit Orders* panel of Table 5 presents the profits from following Strategy *Limit Orders*. These profits are not conditional on execution and thus reflect what traders would actually earn on average when implementing this strategy. The proportion of initial limit orders that execute is 60.3% (when the limit order is placed 10 s after the announcement). The profits are positive and significant in all cases. They are around 50% higher in treatment Corr than in treatment Base at all time delays (the difference, however, is only significant for some delays, and only at the 10% level). The profits decrease in the delay between the announcement and the opening of the position, yet they all retain an economically relevant magnitude, ranging from 1.05% to 3.51%.[30]

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

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

[30]Our analysis ignores brokerage commissions, which are absent in our experimental markets. Goldstein, Irvine, Kandel, and Wiener (2009) provide estimates of institutional brokerage commissions ranging from 9 to 12 basis points, which is much less than the percentage profits available to traders in our experimental markets.

We check to ensure that our results are robust to variations in the trading strategies and report the results in Section B.6 of the Supplementary Material. First, we split the Corr results into "surprising" and "unsurprising" announcements. Both generate significantly positive returns for Strategy *Limit Orders* in almost all cases. Second, we vary the timing of the cutoff point for position opening as well as the timing of the closing order submission from 90 s to 60 s or 120 s. Our results do not change materially (longer delays, e.g., 120 s, actually increase profits because opening and closing price levels tend to be farther apart, even though we more frequently have to close the position with market orders). Third, we vary the price levels of the opening and closing orders. Instead of placing both at the bid–ask midpoint, we can set them between the midpoint and the best same-side order (the best bid in the case of a buy order, or the best ask in the case of a sell order), or between the midpoint and the best opposite-side order (the best ask in the case of a buy order, or the best bid in the case of a sell order). We find that profitability increases the farther apart we set our opening and closing price levels. The reason is that, through these less aggressive orders, we earn part of the spread, which more than compensates for the decreased execution rate.

D. Pricing Efficiency

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

Before we discuss price adjustment, we analyze the pricing efficiency of our markets more generally. Following Powell (2016), we define two measures of mispricing as follows:

$$GD \equiv 100 \cdot \left(exp\left(\frac{1}{N} \sum_{t=0}^{N-1} ln\left(\frac{P_{k,t}}{F_k} \right) \right) - 1 \right), \tag{5}$$

and

$$GAD \equiv 100 \cdot \left(exp\left(\frac{1}{N} \sum_{t=0}^{N-1} \left| ln\left(\frac{P_{k,t}}{F_k} \right) \right| \right) - 1 \right), \tag{6}$$

where $P_{k,t}$ is the quote midpoint ("price") at the end of window t in phase k, F_k is the fundamental value, and $N = 18$ is the number of 10-s windows in each phase. GD thus measures the average geometric deviation of market prices from the fundamental value in the 170 s beginning after the announcement window. It is a signed measure of the percentage by which average prices exceed or fall short of the fundamental value. GAD measures the 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 6
Mispricing

Table 6 compares the measures of relative (GD) and absolute (GAD) mispricing relative to F. “Starting phase” is the phase prior to the first earnings announcement. “First announcement” is the phase following the first announcement in treatment Corr.

Treatment	Phases	GD	GAD	No. of Obs.
All	All	2.06%	8.77%	800
	Starting phase	3.59%	9.95%	160
	Positive earnings change	−3.51%	6.60%	320
	Negative earnings change	6.86%	10.36%	320
Base	All	1.67%	6.66%	400
	Starting phase	0.99%	7.11%	80
	Positive earnings change	−2.54%	5.34%	160
	Negative earnings change	6.23%	7.76%	160
Corr	All	2.44%	10.88%	400
	Starting phase	6.18%	12.78%	80
	First announcement	4.75%	11.05%	80
	Positive earnings change	−4.49%	7.86%	160
	First announcement	−2.09%	8.15%	40
	Surprise	−4.29%	8.33%	29
	No surprise	−5.61%	7.58%	91
	Negative earnings change	7.50%	12.96%	160
	First announcement	11.58%	13.94%	40
	Surprise	8.02%	11.56%	29
	No surprise	5.54%	12.98%	91

Table 6 presents the pricing efficiency results pooled over both treatments and separately for treatments Base and Corr. We report measures of overall mispricing, of mispricing in 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 also reports results for phases following surprising and unsurprising announcements.

Prices on average exceed the fundamental value by 2.06% in the pooled data. The assets in our experimental markets are thus slightly overpriced. This tendency is already visible in the phase leading up to the first announcement, and is more pronounced in treatment Corr than in Base (average GD of 6.18% vs. 0.99%). Most importantly, the mispricing is negative after positive earnings surprises and vice versa. Thus, while prices drift in the direction of the earnings news, they tend to stay below [above] the fundamental value after positive [negative] announcements. Prices thus underreact to earnings news.

Absolute mispricing as measured by GAD is, by definition, greater than or equal to signed mispricing GD because in the latter, cases of positive and negative mispricing cancel out. We find greater GAD after negative than after positive earnings surprises, a result that is driven by the general tendency toward slight overpricing. Mechanically, when market prices lie above F, underreaction to a positive shock to F is partially mitigated by the pre-existing overpricing, while underreaction to a negative shock to F is exacerbated. Interestingly, the absolute mispricing in treatment Corr does not differ markedly after surprising versus after unsurprising announcements. In other words, even though the absolute change in F is substantially greater in surprising announcements, price adjustment (in percentage terms) is similar. We use regressions to gain a better understanding of the dynamics of GAD and report the results in Appendix C. The regressions confirm that mispricing is greater in treatment Corr and following negative earnings surprises. They also show that mispricing in a given market diminishes over time.[31]

Having discussed price levels, we now turn to price adjustment. Price *changes* can correctly reflect changes in fundamental value even when price *levels* deviate from fundamental values. Whether this is the case in a given market is, again, difficult to analyze with field data because neither levels nor changes in fundamental values can be observed. In contrast, we readily observe both in our markets. We analyze the price changes in response to the changes in fundamental values caused by the earnings announcements as follows: We normalize the price prior to an announcement to 0 and the sum of the pre-announcement price and the change in fundamental value to 100% [−100%] in the case of a positive [negative] announcement. A price change of (±)100% then implies that the price change equals the change in fundamental value, whereas a price change of less than [more than] (±)100% indicates underreaction [overreaction]. Because we study price changes, the analysis is unaffected by the slight overpricing in the experimental markets documented earlier.

Graph A of Figure 6 presents the results for the pooled observations from both treatments. It shows that the initial price reaction to an announcement is smaller than the implied change in 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 news but fails to adjust fully. Following positive surprises, prices reflect 33% of ΔF within the 10-s announcement window and reach around 46% within 20 s of the announcement. However, by the end of the phase, they still only reflect 65% of the change in fundamental value. Following negative earnings news, prices reflect only 18% of ΔF within the post-announcement window and take another 30 s to surpass 50%. Yet, by the end of the phase following negative surprises, prices reflect about 87% of ΔF.

The result that prices adjust more fully to changes in fundamental value following negative earnings surprises seems to contradict our finding of greater mispricing following negative surprises. Two factors partly resolve this apparent contradiction. First, as shown before, average prices slightly exceed fundamental values. Thus, even after a price change that equals the change in fundamental value following a negative surprise, prices can still exceed fundamental values, contributing to mispricing. The dashed gray lines in Figure 6 plot the fundamental value prior to the announcement. Their position implies that prices adjust roughly symmetrically relative to this pre-announcement fundamental value. Second, the mispricing reported in

[31]The regression results are reminiscent of the saw-tooth pattern of price adjustment discussed in Plott (2008). He discusses price adjustment over several periods of an experiment. In each period, prices move toward the equilibrium price (but not monotonically so, e.g., because of bid–ask bounce). The opening price of the next period is then closer to the equilibrium price than the opening price of the previous period, but farther from the equilibrium price than the closing price of the previous period, resulting in a saw-tooth-like pattern of deviations from equilibrium. Our regression results in Table C.1 mirror this pattern. Mispricing decreases by approximately 0.85 talers per phase in the pooled data set (results are numerically different but qualitatively similar when we consider positive and negative announcements separately). Over the 4 announcements of a period, the mispricing therefore decreases by approximately 3.4 talers. Mispricing also decreases across periods, by approximately 1.4 talers each. Thus, opening prices of a new period are closer to equilibrium than opening prices of the previous period but farther away than the closing prices of the previous period, following the pattern described in Plott (2008).

FIGURE 6

Adjustment of Stock Price as a Percentage of the Change in Fundamental Value Induced by an Announcement

Graph A of Figure 6 plots results for the pooled data from all treatments; Graph B reports results separately for treatments Base (left) and Corr (right); Graph C plots only Corr data and reports results separately for unsurprising (left) and surprising (right) earnings news. The blue, upward trending [orange, downward trending] lines plot price adjustment following positive [negative] earnings news. The bold, black, horizontal lines indicate full adjustment of prices to the change in F induced by the earnings announcement. The thin, dotted horizontal line at 0 indicates the price level at the moment of the earnings announcement. The dashed horizontal line indicates F prior to the earnings announcement.

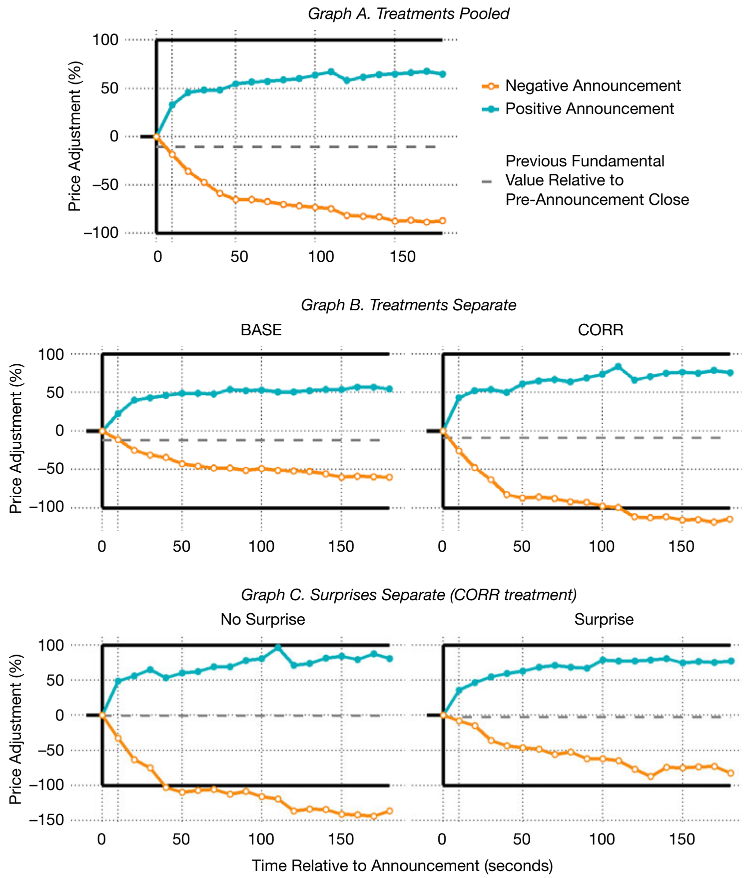

Table 6 is an average over the entire phase, while the 87% price adjustment mentioned in the last paragraph reflects the price level at the end of the phase.

Graph B of Figure 6 presents 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, particularly after negative earnings surprises. Here, the price change fully reflects the change in fundamental value after around 110 s. We conjecture that the more complete price adjustment in Corr may be driven by participants devoting greater attention to announcements in this treatment. Participants know that there can be two types of news (surprising and unsurprising) with very different implications for the asset value, and therefore evaluate the announcements more thoroughly.[32] We summarize the findings of this section 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 amount of price adjustment that is eventually achieved is greater in treatment CORR.

The final question we analyze is whether experimental participants correctly assess the effect of earnings autocorrelation on asset values. Section IV.B documented that the PEAD is more pronounced after surprising than after unsurprising announcements. At the same time, Graph C of Figure 6 shows that prices adjust faster and (particularly in the case of negative announcements) more completely to changes in fundamental value after unsurprising news. In fact, we even observe overshooting after negative unsurprising announcements (i.e., a price change exceeding the change in fundamental value).

This pattern is consistent with underestimation of the effect of earnings autocorrelation. Recall that, in treatment CORR, unsurprising announcements lead to smaller absolute changes in fundamental value than do announcements in treatment BASE, because the earnings change in CORR can be partly anticipated. Participants who underestimate the effect of earnings autocorrelation *overestimate* the effect of *unsurprising* news on the fundamental value. This would imply stronger adjustment of prices to the change in fundamental value and less pronounced PEAD after unsurprising news. By the same argument, participants who underestimate the effect of earnings autocorrelation *underestimate* the effect of *surprising* news on the fundamental value. This would imply weaker adjustment of prices to the change in fundamental value and more pronounced PEAD after surprising news. This is precisely the pattern we documented earlier. Our results thus support Hypothesis 6:

Result Hypothesis 6. In the presence of earnings autocorrelation, prices adjust more completely to the information content of unsurprising than of surprising news. This is consistent with underestimation of the implications for asset values of earnings autocorrelation.

Note that the statement above relates to adjustment as a percentage of the change in fundamental value. As we reported in Section IV.A, the *price reaction* (in talers) to surprising earnings news is significantly stronger than the reaction to unsurprising news, as is also evident from Graph C of Figure 4. Traders thus partially reflect the impact of earnings autocorrelation in their pricing decisions, they just fail to do so sufficiently.

[32] The incomplete adjustment in BASE could be driven by participants underreacting to streaks of same-sign news. When we follow Kieren, Müller-Dethard, and Weber (2020) and analyze announcements as confirming, disconfirming, or corrections, however, we find no evidence for the patterns observed by Kieren et al. in our data.

V. Conclusion

The PEAD is one of 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 observed directly and therefore need to be estimated. These difficulties are absent under the controlled conditions of the experimental laboratory. We report results of a series of experimental markets designed to analyze the importance of earnings autocorrelation for the emergence and strength of the PEAD. We carefully design the experiments to rule out several other potential causes of the drift, such as risk-based explanations and informational asymmetries. The PEAD in our markets persists even in the absence of aggregate risk and of asymmetric information, as well as in the presence of full information about the fundamental value and the earnings-generating process.

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

These results matter for investors and researchers alike. Our finding that PEAD arises in the absence of aggregate risk strengthens the evidence that PEAD is a mispricing phenomenon that can be targeted by appropriately designed trading strategies. The result that profitable trading strategies exist despite the high bid–ask spreads in the experimental markets complements the empirical evidence that PEAD-based trading strategies may yield abnormal returns in the field. Finally, the observation that drift also occurs in a setting without earnings autocorrelation suggests that additional research is warranted on the underlying, likely behavioral, causes of the drift.

That the PEAD can be replicated in the simplified and controlled environment of the experimental lab is a valuable finding in its own right. It shows that the PEAD is not driven by the idiosyncrasies of the institutional environment of securities trading. Rather, exchange regulations, brokers, analysts, news services, etc., are incidental to the phenomenon.

PEAD is considered to derive from underreaction to news, and our experimental results provide supporting evidence. However, underreaction is a much more general phenomenon. Prices tend to underreact to stock splits (Ikenberry and Ramnath (2002)), to bad headline news (Chan (2003)), to analyst recommendations (Ben-Rephael, Da, and Israelsen (2017)), and to 10 K filings (You and Zhang (2009)), to name but a few examples. Empirical evidence also suggests that underreaction is linked to investor (in)attention (e.g., Hirshleifer et al. (2009), Ben-Rephael et al. (2017), and Chen, He, Tao, and Yu (2022)). Our experimental results suggest that empirically observed underreaction phenomena can be replicated in the laboratory. We believe that the opportunities offered by the experimental method should thus be further exploited in the future. Pursuing explanations based on investor (in)attention may be a particularly promising avenue for future research.

Appendix A. Participant Comprehension of the Instructions

The post-experiment questionnaire yielded data about participants' understanding of the experiment. Participants answered an open question asking them for feedback on the clarity of the instructions. Out of 238 participants, 162 (68%) provided an answer. We coded these answers on a 5-point scale using the following categorization: −2 (did not really understand), −1 (partly understood but with difficulty), 0 (more or less understood), +1 (understood, but suggested feedback for improving clarity), and +2 (fully understood). The categorization was performed by two researchers working independently of one another. They arrived at different categorizations in 22 cases (13.6% of all valid 162 answers), with an average absolute difference of 1.36. A third independent coder broke the tie in these cases. The average understanding in BASE (CORR) is 1.65 (1.67), with 81% (86%) of the participants professing to have "fully understood" (category +2), and only 1% (4%) of the participants professing to not have understood (category −2).

We further checked whether a participant's prior stock trading experience was associated with higher payouts in our experiment. We did not find any significant differences and take this as an indication that the instructions provided a level playing field for all participants in terms of understanding.

Appendix B. Trading Screen

FIGURE B.1

The Trading Interface

Translations (in red) were not present in the experiment.

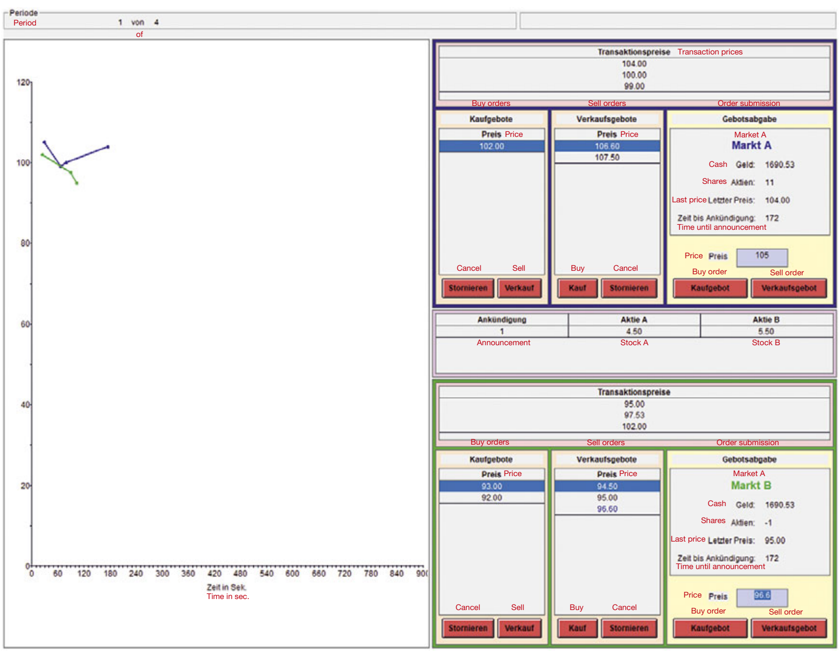

Appendix C. Mispricing

We use regressions to gain a better understanding of the dynamics of absolute mispricing (GAD).[33] Our independent variables are dummies for observations from phases following unsurprising announcements in treatment CORR, for observations from phases following surprising announcements in treatment CORR, and for 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. Model 1 includes data from all earnings announcements, model 2 data from phases following positive announcements, and model 3 data from phases following negative announcements.

Table C.1 presents the results. The coefficient of the constant in model 1, 14.44%, measures the average GAD at the beginning of the first period in treatment BASE

TABLE C.1

Regression Analysis of Absolute Mispricing

OLS regressions of mispricing at the close of consecutive 10-s 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 percentage, calculated using closing midpoints for each window as in Powell (2016) and equation (6). 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 in treatment CORR. 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 10-s window starting at the time of the announcement). *, **, and *** indicate significance at the 5%, 1%, and 0.5% levels, respectively. Standard errors, clustered at the session level, in parentheses.

	Pooled	Positive Announcements	Negative Announcements
	1	2	3
CONSTANT	14.442***	10.436***	16.307***
	(1.846)	(1.338)	(1.946)
CORRELATED_NO_SURPRISE	4.485*	2.624	6.256*
	(2.257)	(1.995)	(2.557)
CORRELATED_SURPRISE	4.134***	3.618*	4.327**
	(1.267)	(1.740)	(1.678)
CORRELATED_FIRST_ANNOUNCEMENT	3.500*	2.077	3.812
	(1.672)	(2.100)	(2.080)
PERIOD0	−1.394***	−1.277***	−0.729
	(0.474)	(0.386)	(0.511)
PHASE0	−0.848**	−0.613***	−1.575***
	(0.324)	(0.186)	(0.398)
WINDOW	−0.749***	−0.473***	−0.970***
	(0.111)	(0.070)	(0.134)
WINDOW2	0.026***	0.018***	0.034***
	(0.005)	(0.003)	(0.006)
R^2	0.063	0.073	0.124
Adj. R^2	0.062	0.071	0.123
No. of obs.	13,769	5680	5464

[33] A detailed analysis of the dynamics of GD does not add insights beyond those of the analysis of GAD. The only exception to this statement is the finding that GD is mainly driven by the sign of the earnings news, with positive [negative] news associated with negative [positive] GD values.

https://doi.org/10.1017/S0022109023000881 Published online by Cambridge University Press

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. 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. This tendency is, of course, a direct reflection of the existence of a post-earnings-announcement drift. Finally, a comparison of the results for models 2 and 3 confirms our earlier finding of greater mispricing following negative earnings surprises.

We conduct a robustness check at the level of the phase instead of at the level of the individual window and report the results in Section B.5 of the Supplementary Material. They remain qualitatively unchanged.

Appendix D. Arbitrage

As a robustness check, we report results for 10 sessions of a third treatment, which we will term treatment Arbi. Sessions of Arbi used the Corr design in all respects except that we ruled out arbitrage opportunities. There are two types of clear (i.e., risk-free) arbitrage opportunities in our experimental design. One arises whenever the sum of the lowest ask prices of stocks A and B falls short of 200 talers. Arbitrageurs in this situation can use market orders to buy 1 A and 1 B share and hold them until the end of the period. Since such a mixed two-share bundle always pays 200 talers at the end of a period, the arbitrageur earns a risk-free profit of 200 talers minus the cost of buying the two shares. A similar mechanism applies for the second type of arbitrage opportunity: whenever the sum of the highest bid prices of stocks A and B exceeds 200 talers, an arbitrageur can profit by (short) selling one share each.

In Arbi, an automated arbitrageur eliminated all instances of arbitrage opportunities precisely 1 s after their coming into existence. The arbitrage algorithm did not face any liquidity constraints and thus could effect arbitrage trades whenever they arose. As a result, arbitrage opportunities were available only 0.6% of the time in treatment Arbi, while traders could have effected arbitrage trades 15.8% and 11.7% of the time in treatments Base and Corr, respectively.[34]

We start our analysis by preparing a modified version of Figure 4 in the body of the article, using Corr and Arbi data and presenting it in Figure D.1. Comparing the Corr and Arbi results shows that arbitrage opportunities did not materially affect our results. In particular, we still find our main patterns of i) observing PEAD in general, of ii) less PEAD following unsurprising announcements, and of iii) greater PEAD following surprising announcements. Furthermore, the price paths following unsurprising announcements look very similar in the two treatments. Only following surprising announcements do we see slightly different trajectories, with PEAD appearing to be somewhat attenuated in Arbi.

[34]Due to a problem of how our experimental software updated data records, our data contains records of persisting arbitrage opportunities in 2 (out of 40) periods of Arbi, even though the arbitrage algorithm properly traded to eliminate the arbitrage opportunity. Since traders in the experiment were unaffected by this software problem (i.e., they correctly saw arbitrage-free prices), but our analysis would be affected, we eliminate these two periods from our analysis.

FIGURE D.1

Price Changes in Talers Relative to the Quote Midpoint

Figure D.1 graphs the average price changes relative to the quote midpoint at the time of the announcement, using the closing quote midpoint for each 10 s window following the announcement. Graph A reports results separately for treatments CORR (left) and ARBI (right); Graph B plots only CORR data and reports results separately for unsurprising (left) and surprising (right) earnings news, using only data from Phases 2 to 4; Graph C plots the same as Graph B, but for ARBI instead of for CORR data. The blue [orange] lines plot 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 10-s announcement window.

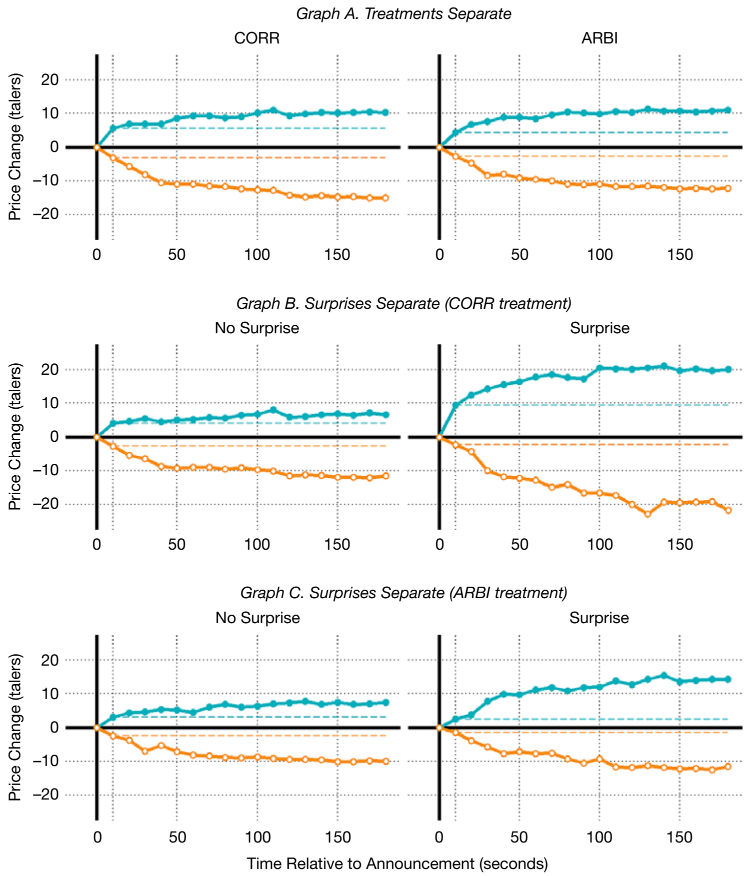

We support these visual impressions with two regressions. Table D.1 is a replication of Table 4 in the body of the article, but using data from all three treatments and including dummy variables to capture the effect of the ARBI treatment. Comparing the coefficient of CORRELATED to that of ARBITRAGE in model 1 reveals essentially no differences between the two treatments. Furthermore, comparing model 1 in Table D.1 to model 1 in Table 4 also shows essentially the same picture. Moving on to model 2, we find that the coefficients of the ARBITRAGE variables are again comparable to the coefficients of the CORRELATED variables, with somewhat lower statistical significance in ARBI in the case of surprising announcements.

We also prepared an exact replication of Table 4, only using data from BASE and ARBI instead of data from BASE and CORR and again find that the ARBI results

TABLE D.1

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). 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 treatment CORR. CORRELATED_SURPRISE is a dummy for an earnings change carrying the opposite sign as in the previous announcement in treatment CORR. CORRELATED_FIRST_ANNOUNCEMENT is a dummy for returns stemming from the phase following the first announcement in treatment CORR (which is neither unambiguously surprising nor unsurprising). The ARBITRAGE variables serve the same function for treatment ARBI as the CORRELATED variables do for treatment CORR. POSITIVE_EARNINGS_CHANGE is a dummy variable for a positive earnings change. 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). 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). *, **, and *** indicate significance at the 5%, 1%, and 0.5% levels, respectively. Standard errors, clustered at the session level, in parentheses.

	Model 1	Model 2
CONSTANT	1.273***	1.209***
	(0.111)	(0.112)
CORRELATED	0.227***	
	(0.061)	
ARBITRAGE	0.222***	
	(0.071)	
CORRELATED_NO_SURPRISE		0.053
		(0.062)
CORRELATED_SURPRISE		0.560***
		(0.135)
CORRELATED_FIRST_ANNOUNCEMENT		0.381***
		(0.045)
ARBITRAGE_NO_SURPRISE		0.059
		(0.087)
ARBITRAGE_SURPRISE		0.443*
		(0.197)
ARBITRAGE_FIRST_ANNOUNCEMENT		0.417***
		(0.109)
POSITIVE_EARNINGS_CHANGE	−0.187***	−0.185***
	(0.048)	(0.048)
PERIOD0	−0.016	−0.017
	(0.032)	(0.034)
PHASE0	−0.017	0.028
	(0.020)	(0.020)
WINDOW0	−0.227***	−0.228***
	(0.024)	(0.024)
WINDOW02	0.010***	0.010***
	(0.001)	(0.001)
R^2	0.005	0.005
Adj. R^2	0.004	0.004
No. of obs.	14,955	14,955
RMSE	6.013	6.012

are very similar to the CORR results. We are happy to provide this regression table upon request. We conclude that the arbitrage opportunities we observe in our main experimental treatments (BASE and CORR) are not materially linked to the main conclusions of our article.

Supplementary Material

To view supplementary material for this article, please visit http://doi.org/10.1017/S0022109023000881.

References

Ackert, L. F.; B. K. Church; and L. Qi. "An Experimental Examination of Portfolio Choice." *Review of Finance*, 20 (2016), 1427–1447.

Andries, M.; M. Bianchi; K. K. Huynh; and S. Pouget. "Return Predictability, Expectations, and Investment: Experimental Evidence." Working Paper, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3710466 (2022).

Asparouhova, E.; M. Hertzel; and M. Lemmon. "Inference from Streaks in Random Outcomes: Experimental Evidence on Beliefs in Regime Shifting and the Law of Small Numbers." *Management Science*, 55 (2009), 1766–1782.

Assenza, T.; T. Bao; C. Hommes; and D. Massaro. "Experiments on Expectations in Macroeconomics and Finance." *Research in Experimental Economics*, 17 (2014), 11–70.

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

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

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

Ball, R.; S. Kothari; and R. Watts. "The Economic Determinants of the Relation Between Earnings Changes and Stock Returns." *Accounting Review*, 68 (1993), 622–638.

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

Bathke Jr, A. W.; K. S. Lorek; and G. L. Willinger. "The Security Market's Reaction to Firms' Quarterly Earnings Evidencing Varying Degrees of Autocorrelation." *Advances in Accounting*, 22 (2006), 29–43.

Battalio, R. H., and R. R. Mendenhall. "Earnings Expectations, Investor Trade Size, and Anomalous Returns Around Earnings Announcements." *Journal of Financial Economics*, 77 (2005), 289–319.

Battalio, R. H., and R. R. Mendenhall. "Post-Earnings Announcement Drift: Bounds on Profitability for the Marginal Investor." *Financial Review*, 46 (2011), 513–539.

Ben-Rephael, A.; Z. Da; and R. D. Israelsen. "It Depends on Where You Search: Institutional Investor Attention and Underreaction to News." *Review of Financial Studies*, 30 (2017), 3009–3047.

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

Bernard, V. *Stock Price Reactions to Earnings Announcements: A Summary of Recent Anomalous Evidence and Possible Explanations*, Vol. 1. New York, NY: Russell Sage Foundation (1993).

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

Bernard, V., and J. Thomas. "Evidence that Stock Prices Do Not Fully Reflect the Implications of Current Earnings for Future Earnings." *Journal of Accounting and Economics*, 13 (1990), 305–340.

Bhattacharya, N.; P. Olsson; and H. Park. "Predictability of Analyst Earnings Forecast Errors and Institutional and Individual Investors' Reactions to Earnings News." *Journal of Accounting, Auditing and Finance*, 36 (2020), 1–28.

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

Booth, G. G.; J.-P. Kallunki; and T. Martikainen. "Delayed Price Response to the Announcements of Earnings and its Components in Finland." *European Accounting Review*, 6 (1997), 377–392.

Bossaerts, P., and C. Plott. "The CAPM in Thin Experimental Financial Markets." *Journal of Economic Dynamics and Control*, 26 (2002), 1093–1112.

Bossaerts, P., and C. Plott. "Basic Principles of Asset Pricing Theory: Evidence from Large-Scale Experimental Financial Markets." *Review of Finance*, 8 (2004), 135–169.

Caginalp, G.; D. Porter; and L. Hao. "Asset Market Reactions to News: An Experimental Study." ESI Working Paper, Chapman University (2011).

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

Chakrabarty, B.; P. Moulton; and X. Wang. "Attention: How High-Frequency Trading Improves Price Efficiency Following Earnings Announcements." *Journal of Financial Markets*, 59 (2022), 100690.

Chan, W. C. "Stock Price Reaction to News and No-News: Drift and Reversal After Headlines." *Journal of Financial Economics*, 70 (2003), 223–260.

Chen, X.; W. He; L. Tao; and J. Yu. "Attention and Underreaction-Related Anomalies." *Management Science*, 69 (2022), 636–659.

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

Chordia, T.; A. Subrahmanyam; and Q. Tong. "Have Capital Market Anomalies Attenuated in the Recent Era of High Liquidity and Trading Activity?" *Journal of Accounting and Economics*, 58 (2014), 41–58.

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

DellaVigna, S., and J. M. Pollet. "Investor Inattention and Friday Earnings Announcements." *Journal of Finance*, 64 (2009), 709–749.

Doyle, J. T.; R. J. Lundholm; and M. T. Soliman. "The Extreme Future Stock Returns Following I/B/E/S Earnings Surprises." *Journal of Accounting Research*, 44 (2006), 849–887.

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

Fama, E. F., and K. R. French. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics*, 33 (1993), 3–56.

Feng, L., and M. S. Seasholes. "Do Investor Sophistication and Trading Experience Eliminate Behavioral Biases in Financial Markets?" *Review of Finance*, 9 (2005), 305–351.

Fink, J. "A Review of the Post-Earnings-Announcement Drift." *Journal of Behavioral and Experimental Finance*, 29 (2021) 100446.

Fischbacher, U. "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics*, 10 (2007), 171–178.

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

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

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

Freeman, R. N., and S. Tse. "The Multiperiod Information Content of Accounting Earnings: Confirmations and Contradictions of Previous Earnings Reports." *Journal of Accounting Research*, 27 (1989), 49–79.

Friedman, D.; G. W. Harrison; and J. W. Salmon. "The Informational Efficiency of Experimental Asset Markets." *Journal of Political Economy*, 92 (1984), 349–408.

Frydman, C., and G. Nave. "Extrapolative Beliefs in Perceptual and Economic Decisions: Evidence of a Common Mechanism." *Management Science*, 63 (2017), 2340–2352.

Goldstein, M.; P. Irvine; E. Kandel; and Z. Wiener. "Brokerage Commissions and Institutional Trading Patterns." *Review of Financial Studies*, 22 (2009), 5175–5212.

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

Hlavac, M. *stargazer: Well-Formatted Regression and Summary Statistics Tables*. Bratislava: Central European Labour Studies Institute (CELSI) (2018). R package version 5.2.2.

Hou, K.; W. Xiong; and L. Peng. "A Tale of Two Anomalies: The Implications of Investor Attention for Price and Earnings Momentum." Working Paper, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=976394 (2009).

Hung, M.; X. Li; and S. Wang. "Post-Earnings-Announcement Drift in Global Markets: Evidence from an Information Shock." *Review of Financial Studies*, 28 (2015), 1242–1283.

Ikenberry, D. L., and S. Ramnath. "Underreaction to Self-Selected News Events: The Case of Stock Splits." *Review of Financial Studies*, 15 (2002), 489–526.

Janssen, D.-J.; J. Li; J. Qiu; and U. Weitzel. "The Disposition Effect and Underreaction to Private Information." *Journal of Economic Dynamics and Control*, 113 (2020), 103856.

Jiang, H.; S. Li; and H. Wang. "Pervasive Underreaction: Evidence from High-Frequency Data." *Journal of Financial Economics*, 141 (2021), 573–599.

Karpoff, J. M. "The Relation Between Price Changes and Trading Volume: A Survey." *Journal of Financial and Quantitative Analysis*, 22 (1987), 109–126.

Ke, B., and S. Ramalingegowda. "Do Institutional Investors Exploit the Post-Earnings Announcement Drift?" *Journal of Accounting and Economics*, 39 (2005), 25–53.
Kieren, P.; J. Müller-Dethard; and M. Weber. "Can Agents Add and Subtract When Forming Beliefs?" Working Paper, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3644226 (2020).
Kim, D., and M. Kim. "A Multifactor Explanation of Post-Earnings Announcement Drift." *Journal of Financial and Quantitative Analysis*, 38 (2003), 383–398.
Kirchkamp, O. "Importing z-Tree Data into R." *Journal of Behavioral and Experimental Finance*, 22 (2019), 1–2.
Kleinlercher, D., and T. Stöckl. "Thou Shalt Not Trade—An Analysis of the Violations of No-Trade Predictions in Experimental Asset Markets." *Journal of Behavioral and Experimental Finance*, 32 (2021), 100590.
Kothari, S. P.; J. Lewellen; and J. B. Warner. "Stock Returns, Aggregate Earnings Surprises and Behavioral Finance." *Journal of Financial Economics*, 79 (2006), 537–568.
Leifeld, P. "texreg: Conversion of Statistical Model Output in R to LaTeX and HTML Tables." *Journal of Statistical Software*, 55 (2013), 1–24.
Leitner, J., and U. Leopold-Wildburger. "Experiments on Forecasting Behavior with Several Sources of Information - A Review of the Literature." *European Journal of Operational Research*, 213 (2011), 459–469.
Li, J. "Slow Price Adjustment to Public News in After-Hours Trading." *Journal of Trading*, 11 (2016), 16–31.
Louhichi, W. "Adjustment of Stock Prices to Earnings Announcements: Evidence from Euronext Paris." *Review of Accounting and Finance*, 7 (2008), 102–115.
Maines, L., and J. Hand. "Individuals' Perceptions and Misperceptions of Time Series Properties of Quarterly Earnings." *Accounting Review*, 71 (1996), 317–336.
Martineau, C. "Rest in Peace Post-Earnings Announcement Drift." *Critical Finance Review*, 11 (2022), 613–646.
McLean, R. D., and J. Pontiff. "Does Academic Research Destroy Stock Return Predictability?" *Journal of Finance*, 71 (2016), 5–32.
Meursault, V.; P. Liang; B. Routledge; and M. Scanlon. "PEAD.txt: Post-Earnings-Announcement Drift Using Text." *Journal of Financial and Quantitative Analysis*, 58 (2023), 2299–2326.
Milgrom, P., and N. Stokey. "Information, Trade and Common Knowledge." *Journal of Economic Theory*, 26 (1982), 17–27.
Nam, J.; J. Wang; and G. Zhang. "Strategic Trading Against Retail Investors with Loss-Aversion." *International Review of Economics & Finance*, 17 (2008), 45–55.
NASDAQ. "NASDAQ Nordic Market Model." Available at https://www.nasdaq.com/docs/2023/07/03/Nasdaq-Nordic-Market-Model-2023_06_.pdf (2023). (accessed 3 August 2023).
Ng, J.; T. O. Rusticus; and R. S. Verdi. "Implications of Transaction Costs for the Post-Earnings Announcement Drift." *Journal of Accounting Research*, 46 (2008), 661–696.
Noussair, C. N., and S. Tucker. "Cash Inflows and Bubbles in Asset Markets with Constant Fundamental Values." *Economic Inquiry*, 54 (2016), 1596–1606.
Odean, T. "Are Investors Reluctant to Realize Their Losses?" *Journal of Finance*, 53 (1998), 1775–1798.
Palan, S. "A Review of Bubbles and Crashes in Experimental Asset Markets." *Journal of Economic Surveys*, 27 (2013), 570–588.
Palan, S. "GIMS - Software for Asset Market Experiments." *Journal of Behavioral and Experimental Finance*, 5 (2015), 1–14.
Pavlova, I., and A. Parhizgari. "In Search of Momentum Profits: Are They Illusory?" *Applied Financial Economics*, 21 (2011), 1617–1639.
Plott, C., and S. Sunder. "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models." *Journal of Political Economy*, 90 (1982), 663–698.
Plott, C. R. "Properties of Disequilibrium Adjustment in Double Auction Markets." In *Handbook of Experimental Economics Results*, Vol. 1, C. R. Plott; and V. L. Smith, eds. Amsterdam, Netherlands: North-Holland (2008), 16–21.
Powell, O. "Numeraire Independence and the Measurement of Mispricing in Experimental Asset Markets." *Journal of Behavioral and Experimental Finance*, 9 (2016), 56–62.
R Core Team. "R: A Language and Environment for Statistical Computing." Vienna, Austria: R Foundation for Statistical Computing (2017).
Rangan, S., and R. Sloan. "Implications of the Integral Approach to Quarterly Reporting for the Post-Earnings-Announcement Drift." *Accounting Review*, 73 (1998), 353–371.
Rendleman, R.; C. Jones; and H. Latane. "Further Insight into the Standardized Unexpected Earnings Anomaly: Size and Serial Correlations Effect." *Financial Review*, 22 (1987), 131–144.

Richardson, S.; I. Tuna; and P. Wysocki. "Accounting Anomalies and Fundamental Analysis: A Review of Recent Research Advances." *Journal of Accounting and Economics*, 50 (2010), 410–454.

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

Smith. "Markets as Economizers of Information: Experimental Examination of the 'Hayek Hypothesis'." *Economic Inquiry*, 20 (1982), 165–179.

Sun, Q. "Stock Price Discovery in Earnings Season." *International Journal of Business and Finance Research*, 9 (2015), 1–15.

Tirole, J. "On the Possibility of Speculation Under Rational Expectations." *Econometrica*, 50 (1982), 1163–1181.

Truong, C. "Post Earnings Announcement Drift and the Roles of Drift-Enhanced Factors in New Zealand." *Pacific-Basin Finance Journal*, 18 (2010), 139–157.

Truong, C. "Post-Earnings Announcement Abnormal Return in the Chinese Equity Market." *Journal of International Financial Markets, Institutions and Money*, 21 (2011), 637–661.

Weber, M., and F. Welfens. "How Do Markets React to Fundamental Shocks?: An Experimental Analysis on Underreaction and Momentum." Working Paper, University of Mannheim (2007).

Wickham, H. *ggplot2: Elegant Graphics for Data Analysis*. New York: Springer-Verlag (2016).

You, H., and X. Zhang. "Financial Reporting Complexity and Investor Underreaction to 10-k Information." *Review of Accounting Studies*, 14 (2009), 559–586.

Zhang, G., and S. Zhang. "Information Efficiency of the US Credit Default Swap Market: Evidence from Earnings Surprises." *Journal of Financial Stability*, 9 (2013), 720–730.

Zhu, H. *kableExtra: Construct Complex Table with 'kable' and Pipe Syntax* (2019). R package version 1.1.0.