Revisiting Market Efficiency: Evidence from Post-Earnings Announcement Drift in U.S. Equities

Osman Idris Zeinu

Advisor: Professor Alessandro Citanna

May 6, 2026

Thesis submitted to the Division of Social Sciences at New York University Abu Dhabi for the Economics Capstone Project in 2026

Academic Integrity Declaration

I submit this Thesis for the Capstone project in Economics to the Division of Social Sciences at New York University Abu Dhabi. The Thesis provides sufficient information on sources and methods, including mathematical derivations, data sources, sample periods, manipulations, and transformations, to allow for a full replication of all findings reported in the Thesis.

I hereby declare that this Thesis is solely my work and contains no material I have previously submitted for assessment at New York University or elsewhere, without the express approval of all instructors including my Capstone advisor. To the best of my knowledge and belief, the Thesis contains no material previously published or written by another person except where due reference is made in the Thesis.

I declare that I have read, and in undertaking this research, have complied with the Policy for Academic Integrity of Students at NYU, the Economics Capstone guidelines on Responsible Use of AI, and all IRB policies on research with human subjects. I also declare that I understand what is meant by plagiarism and that it is unacceptable. Any data sources, mathematical derivations, and computer code used to derive the results presented in this Thesis, including digital and electronic media sources, have been properly acknowledged.

Signature: Osman Idris Zeinu Date: May 6, 2026

AI-Software Usage Self-Report for Capstone in Economics

Student Name: Osman Idris Zeinu

Student Group: Individual project

Thesis Title: Revisiting Market Efficiency: Evidence from Post-Earnings Announcement Drift in U.S. Equities

Date: May 6, 2026

Please check the appropriate boxes to indicate your use of any AI-software in the preparation of your capstone thesis.

General Usage

- ☒ I used AI for brainstorming and generating ideas related to my thesis topic.
- ☒ I employed AI for clarifying and explaining complex economic concepts.
- ☒ I used AI to draft and edit portions of my thesis.
- ☒ I used AI to identify relevant academic sources and references for my thesis.

Data Analysis and Interpretation

- ☒ I used AI for help in analyzing data and statistical results.
- ☒ I consulted AI to interpret the economic aspects of the data.
- ☒ I used AI to generate hypotheses or potential explanations for observed economic behaviors or outcomes.

Literature Review and Citations

- ☒ I used AI to find relevant research papers and articles related to my thesis topic.

☒ I employed AI to help write summaries and critiques of academic literature.

☒ I used AI to generate citations and references in the appropriate format.

Ethical Considerations

☒ I ensured that I used AI responsibly and that my usage of AI adhered to the AI guidelines and did not involve plagiarism.

☒ I critically assessed and verified any information obtained from AI for accuracy.

Additional explanation: AI tools were used as a research and writing assistant to support brainstorming, conceptual clarification, source discovery, data interpretation, citation formatting, and LaTeX debugging. All substantive claims, citations, data construction choices, regression outputs, tables, figures, and final interpretations were reviewed and verified before inclusion in the thesis. I remain solely responsible for the content of the thesis.

Signature: Osman Idris Zeinu Date: May 6, 2026

Contents

List of Figures

List of Tables

Revisiting Market Efficiency: Evidence from Post-Earnings Announcement Drift in U.S. Equities[1]

Osman Idris Zeinu

May 6, 2026

Abstract

This paper revisits post-earnings announcement drift (PEAD) in modern U.S. equity markets, asking whether earnings surprises still predict returns after announcements and why the adjustment may be temporary. I develop a narrative-diffusion framework in which investors remain Bayesian, but the share of narrative-sensitive traders evolves over time, generating front-loaded but incomplete price adjustment. Empirically, I test the prediction using analyst-based earnings surprises, event-time market-adjusted abnormal returns, and price-formation regressions for U.S. equities from 2018 to 2024. Earnings news is incorporated strongly around the announcement window, but surprise ranks continue to predict short- and medium-horizon post-announcement abnormal returns. PEAD remains statistically visible, though calendar-time and factor-adjusted alpha tests provide more conservative evidence on tradability.

Keywords: Post-earnings announcement drift; market efficiency; earnings surprises; narrative diffusion; abnormal returns; asset pricing.

[1] I am grateful to Professor Alessandro Citanna for his guidance and advice throughout this capstone project. Any remaining errors are my own.

1 Introduction

Post-earnings announcement drift (PEAD) is one of the most robust empirical anomalies in financial economics. First documented by Ball and Brown (1968) and later formalized by Bernard and Thomas (1989, 1990), PEAD refers to the tendency of stock prices to continue moving in the direction of an earnings surprise after the announcement date. Firms that report earnings above expectations tend to earn positive abnormal returns after the announcement, while firms that report earnings below expectations tend to underperform.

This pattern is important because it challenges a strict interpretation of the Efficient Market Hypothesis (EMH). If earnings announcements are public, scheduled, and widely followed, then prices should adjust rapidly once the information becomes available. Yet the PEAD literature shows that earnings surprises can continue to predict returns after the announcement window. The puzzle is not simply that markets fail to react. In many cases, they react strongly and immediately. The deeper puzzle is why the initial reaction appears incomplete.

This paper asks whether PEAD remains visible in modern U.S. equity markets and how it should be interpreted theoretically. The central argument is that PEAD is best understood as a pattern of *front-loaded but incomplete price adjustment*. Earnings news is incorporated quickly, but the post-announcement return path still reflects delayed adjustment. This creates a middle ground between two extreme interpretations: markets are not completely slow or inattentive, but they are also not perfectly efficient in the immediate aftermath of earnings announcements.

1.1 Why existing explanations remain incomplete

Existing explanations of PEAD generally fall into two broad categories. Risk-based explanations argue that the post-announcement return spread may compensate investors for bearing systematic risk. However, it is difficult for standard risk-premium explanations to fully account for the timing and direction of the drift, especially because PEAD is closely linked to the sign and magnitude of earnings surprises.

Behavioral and limited-attention explanations provide a more natural account of delayed reaction. These models emphasize investor underreaction, limited attention, conservatism, or gradual information diffusion. They help explain why public information may not be incorporated immediately. However, many of these explanations treat investor composition or sentiment as fixed. They do not fully explain why the impact of narratives or sentiment should strengthen for a period and then fade.

Thus, the central puzzle remains:

> *Prices react quickly, but not completely. Drift appears after the announcement, yet it does not last indefinitely.*

1.2 Research question

This paper asks:

> Can PEAD be explained by a dynamic process in which narratives temporarily spread through the market, amplify the price impact of sentiment, and then dissipate as more sophisticated investors dominate?

The theoretical section develops a model in which all agents are Bayesian, but investors differ in attention, sophistication, and sensitivity to narratives. The empirical section then tests the observable implication of this mechanism: if adjustment is incomplete, earnings-surprise ranks should continue to predict post-announcement abnormal returns even after the announcement-window reaction.

1.3 Contribution

The paper makes two contributions. First, it develops a theoretical explanation of PEAD based on narrative diffusion and evolving market composition. The model contains three investor types:

- **Retail (R):** low attention, easily influenced;
- **Narrative (N):** preferences shaped by prevailing narratives;
- **Sophisticated (H):** institutional, fully attentive and not impacted by sentiment.

All agents are Bayesian; none hold systematically biased beliefs about fundamentals. The key feature is that the population share of narrative traders evolves over time. Retail investors may adopt narratives after interacting with narrative traders, while narrative traders may transition into high-precision types after interacting with sophisticated investors. This random-matching structure generates a hump-shaped path for the narrative-trader share.

Because only narrative traders respond to sentiment, the price impact of sentiment depends on the size of the narrative-trader group. When the narrative share rises, sentiment has a stronger effect on demand and prices, creating post-announcement drift. When the narrative share falls, the sentiment multiplier weakens and the drift fades. PEAD therefore emerges as a temporary amplification mechanism rather than as permanent irrational mispricing.

Second, the paper provides new empirical evidence on PEAD in a modern sample from 2018 to 2024. The empirical analysis uses analyst based earnings surprises, event time abnormal returns, SUE decile portfolios, price formation regressions, transaction cost checks, and factor adjusted alpha tests. This allows the paper to distinguish statistical event time drift from economically tradable alpha. The results show that PEAD remains visible in event time, but the evidence for clean factor adjusted trading profits is more conservative.

1.4 Structure of the paper

The rest of the paper proceeds as follows. Section 2 reviews the literature on PEAD, limited attention, and narrative based explanations of market anomalies. Section 3 presents the theoretical model of narrative diffusion and derives its main implications. Section 4 describes the data, sample construction, earnings-surprise measure, and event-time return methodology. Section 5 presents the empirical evidence on price formation, post-announcement drift, tradability, and factor-adjusted alpha. Section 6 concludes.

2 Literature Review

PEAD is one of the most persistent deviations from the Efficient Market Hypothesis. Ball and Brown (1968) first documented that stock prices incorporate earnings news only gradually rather than instantaneously. Subsequent work by Bernard and Thomas (1989, 1990) demonstrated that firms with positive (negative) earnings surprises continue to experience positive (negative) abnormal returns for several months after the announcement, even after controlling for known risk factors. These foundational studies established PEAD as a robust empirical regularity and highlighted the central puzzle of slow information diffusion in otherwise competitive markets.

A recent and especially relevant contribution is Martineau (2021), who argues that PEAD has substantially weakened in modern financial markets. Martineau shows that price discovery around earnings announcements has become increasingly concentrated on the announcement date, with PEAD disappearing earlier for large stocks and only more recently for microcap stocks. This finding is important for this paper because it shifts the empirical question from whether PEAD existed historically to whether any form of post-announcement drift remains detectable in recent data. I therefore treat Martineau's evidence as the key modern benchmark: rather than assuming that PEAD persists in its classic form, I test whether earnings-surprise ranks continue to predict post-announcement abnormal returns in the 2018–2024 period.

A substantial literature connects PEAD to behavioral or limited-attention mechanisms. Hong and Stein (1999) showed that when investors process information at different speeds, prices adjust only gradually. Barberis et al. (1998) linked conservatism biases to underreaction, while Hirshleifer and Teoh (2003) and Peng and Xiong (2006) formalized how limited attention leads investors to overweight salient signals and underreact to more complex components of earnings news. Together, these models show that heterogeneous information processing rather than irrationality can generate persistent return predictability and momentum-type anomalies such as PEAD.

A parallel and increasingly influential literature emphasizes the role of narratives and social transmission. Shiller (2017) argues that economic stories spread epidemiologically, shaping expectations, beliefs, and market dynamics. Rational-inattention models such as Gabaix (2014) provide microfoundations for selective information processing, where agents optimally ignore certain signals. More recently, Bordalo et al. (2022) develop a diagnostic-expectations framework in which individuals overweight salient information, generating systematic forecast errors and boom–bust patterns in macroeconomic and financial variables. These advances move beyond static behavioral assumptions and emphasize how beliefs evolve in response to changes in the information environment.

Despite these contributions, two key gaps remain central to understanding PEAD:

1. **Sentiment is typically exogenous.** Existing models do not fully explain why its price impact varies systematically over time.

2. **Investor composition is often fixed.** Heterogeneous-agent models rarely allow population shares to evolve, limiting their ability to generate changing amplification of sentiment.

This paper addresses both limitations. Drawing on narrative economics and attention-based theories, it develops a framework in which investors differ by attention rather than by biased beliefs, and all agents remain Bayesian. What distinguishes them is the extent to which they respond to narrative-driven motives.

Crucially, the model incorporates endogenous population dynamics. Retail traders may adopt narrative-based behavior after interacting with narrative investors, while narrative traders may transition into sophisticated types when exposed to more disciplined informational environments. These transitions, modeled through a contagion-like random matching process, generate a hump-shaped evolution of narrative participation in the market.

This endogenous narrative diffusion interacts with sentiment shocks to produce PEAD mechanically. Sentiment influences only narrative traders' demand, and its aggregate price impact depends on the size of the narrative group. As narratives spread, the sentiment multiplier increases, generating delayed price response; as narratives fade, the multiplier contracts, causing drift to disappear.

Thus, the model yields PEAD as a temporary amplification of sentiment driven by evolving market composition, without invoking irrational expectations or non-

Bayesian behavior. It embeds classical signal-extraction insights (Ball and Brown, 1968; Bernard and Thomas, 1989, 1990) within a modern attention-oriented framework inspired by Hong and Stein (1999), Shiller (2017), Gabaix (2014), and Bordalo et al. (2022). This theoretical framework motivates the empirical analysis that follows, which tests whether earnings-surprise ranks generate a front-loaded but incomplete pattern of price adjustment in modern U.S. equity markets.

3 The Model

This section presents the theoretical framework used to study information processing, sentiment transmission, and post-announcement return predictability following earnings announcements. The model remains fully Bayesian: no agent holds irrational beliefs. Instead, post-earnings announcement drift arises from endogenous population dynamics and sentiment-driven demand among narrative traders.

The environment is discrete time, $t = 0, 1, 2, \ldots$. Markets are competitive, and there is a single risky asset in fixed supply normalized to one. The equilibrium price at time t depends on three forces: beliefs formed through public and private information, the composition of investor types, and sentiment shocks that affect only narrative traders.

3.1 Economic Environment

There is a single risky asset that pays the next-period fundamental payoff θ_{t+1}. Investors trade a claim to this payoff at price P_t. At time t, investors form beliefs about θ_{t+1} using public and private information.

Fundamentals

At each date t, investors share a common prior about the next-period payoff:

$$\theta_{t+1} \sim \mathcal{N}(\mu_t, \sigma_t^2). \tag{1}$$

The prior precision is:

$$\pi_t \equiv \frac{1}{\sigma_t^2}. \tag{2}$$

Information Structure

Investors observe noisy signals about θ_{t+1}.

Public earnings signal.

$$a_t = \theta_{t+1} + \varepsilon_{a,t}, \qquad \varepsilon_{a,t} \sim \mathcal{N}(0, \sigma_{a,t}^2). \tag{3}$$

The public-signal precision is:

$$\pi_{a,t} \equiv \frac{1}{\sigma_{a,t}^2}. \tag{4}$$

Private signal of investor i.

$$s_{i,t} = \theta_{t+1} + u_{i,t}, \qquad u_{i,t} \sim \mathcal{N}(0, \sigma_{u,t}^2). \tag{5}$$

The private-signal precision is:

$$\pi_{u,t} \equiv \frac{1}{\sigma_{u,t}^2}. \tag{6}$$

All noise processes are mutually independent and independent of fundamentals.

3.2 Investor Types

The market contains a continuum of investors partitioned into three types. Population shares at time t are:

$$(r_t, n_t, h_t), \qquad r_t + n_t + h_t = 1. \tag{7}$$

- **Retail traders (R):** low-attention investors with limited information-processing capacity;

- **Narrative traders (N):** investors who update beliefs rationally but whose demand is also affected by sentiment;
- **High-precision traders (H):** sophisticated investors who process information more precisely and are not directly affected by sentiment.

For each type $k \in \{R, N, H\}$, the effective public and private signal variances are:

$$\sigma_{a,t}^{k\,2}, \qquad \sigma_{u,t}^{k\,2}. \tag{8}$$

The corresponding precisions are:

$$\pi_{a,t}^k = \frac{1}{\sigma_{a,t}^{k\,2}}, \qquad \pi_{u,t}^k = \frac{1}{\sigma_{u,t}^{k\,2}}. \tag{9}$$

3.3 Bayesian Beliefs

Posterior precision for type k is:

$$\Pi_t^k = \pi_t + \pi_{a,t}^k + \pi_{u,t}^k, \qquad v_t^k = (\Pi_t^k)^{-1}. \tag{10}$$

Posterior mean for investor i of type k is:

$$m_{i,t}^k = \frac{\pi_t \mu_t + \pi_{a,t}^k a_t + \pi_{u,t}^k s_{i,t}}{\Pi_t^k}. \tag{11}$$

With a continuum of investors, private noise averages out within each type:

$$\bar{m}_t^k = \frac{\pi_t \mu_t + \pi_{a,t}^k a_t + \pi_{u,t}^k \theta_{t+1}}{\Pi_t^k}. \tag{12}$$

Sentiment. Sentiment ξ_t is payoff-irrelevant:

$$\mathbb{E}[\xi_t] = 0, \qquad \mathrm{Cov}(\xi_t, \theta_{t+1}) = 0. \tag{13}$$

It affects the demand of narrative traders but not their beliefs about fundamentals.

3.4 Preferences and Demand

All investors have CARA utility with risk aversion γ. Wealth evolves as:

$$W_{t+1} = W_t + X_{i,t}(\theta_{t+1} - P_t). \tag{14}$$

For non-narrative traders, optimal demand is:

$$X_{i,t}^k = \frac{m_{i,t}^k - P_t}{\gamma v_t^k}. \tag{15}$$

Narrative traders have utility:

$$U_N(W_{t+1}) = -\exp\{-\gamma(W_{t+1} + \beta\xi_t X_{i,t})\}. \tag{16}$$

This generates an additional sentiment-driven demand term:

$$X_t^{N,\text{sent}} = \frac{\beta\xi_t}{\gamma v_t^N}. \tag{17}$$

For clarity, I separate the rational component of narrative-trader demand from the aggregate sentiment component. The rational demand of narrative traders is:

$$X_t^{N,rat} = \frac{\bar{m}_t^N - P_t}{\gamma v_t^N}. \tag{18}$$

Aggregate sentiment demand is:

$$X_t^{sent} = \beta n_t \xi_t, \tag{19}$$

where β absorbs the constant scaling term from risk aversion and posterior variance. The key point is that aggregate sentiment demand is proportional to the narrative-trader share n_t.

3.5 Market Clearing and Price

Market clearing requires:

$$r_t X_t^R + n_t X_t^{N,rat} + h_t X_t^H + X_t^{sent} = 1. \tag{20}$$

Solving the market-clearing condition yields the price decomposition:

$$P_t = P_t^{fund} + \phi_t \xi_t. \tag{21}$$

Aggregate precision weight

$$\Omega_t \equiv \frac{r_t}{v_t^R} + \frac{n_t}{v_t^N} + \frac{h_t}{v_t^H} = r_t \Pi_t^R + n_t \Pi_t^N + h_t \Pi_t^H. \tag{22}$$

Fundamental price

$$P_t^{fund} = \frac{\frac{r_t \bar{m}_t^R}{v_t^R} + \frac{n_t \bar{m}_t^N}{v_t^N} + \frac{h_t \bar{m}_t^H}{v_t^H} - \gamma}{\Omega_t}. \tag{23}$$

Sentiment multiplier

$$\phi_t = \frac{\gamma \beta n_t}{\Omega_t}. \tag{24}$$

The term $\phi_t \xi_t$ is the sentiment wedge. Its strength depends on the narrative-trader share n_t. When the narrative-trader share rises, sentiment has a stronger impact on prices. When the narrative-trader share falls, sentiment loses influence.

3.6 Random-Matching Population Dynamics

Each period, agents meet one random partner.

Contagion $(R \to N)$**.**

$$\Pr(R \to N \mid \text{meet } N) = \lambda \quad \Rightarrow \quad \text{Flow}_{R \to N} = \lambda r_t n_t. \tag{25}$$

Learning from sophistication $(N \to H)$**.**

$$\Pr(N \to H \mid \text{meet } H) = \alpha \quad \Rightarrow \quad \text{Flow}_{N \to H} = \alpha n_t h_t. \tag{26}$$

Absorbing state.

$$\Pr(H \to H) = 1. \tag{27}$$

Laws of motion.

$$r_{t+1} = r_t - \lambda r_t n_t, \tag{28}$$

$$n_{t+1} = n_t + \lambda r_t n_t - \alpha n_t h_t, \tag{29}$$

$$h_{t+1} = h_t + \alpha n_t h_t. \tag{30}$$

The system generates a hump-shaped evolution of the narrative share n_t. In the long run:

$$h_t \to 1, \tag{31}$$

$$r_t \to 0, \tag{32}$$

$$n_t \to 0. \tag{33}$$

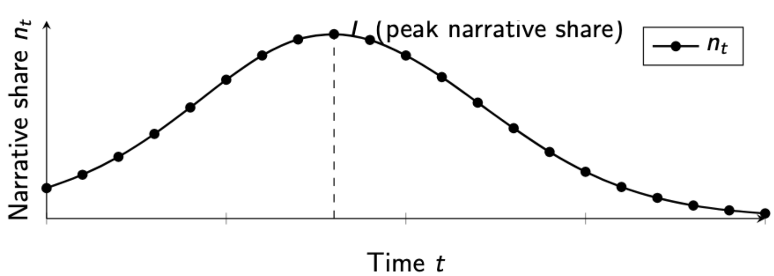

Figure 1: Share of Narrative Traders Through Time

3.7 Main Theoretical Predictions

Proposition 1: Positive Drift Before the Peak. When the narrative-trader share is increasing, the sentiment multiplier rises. Under the maintained condition that changes in Ω_t do not dominate changes in n_t, if $n_{t+1} > n_t$, then $\phi_{t+1} > \phi_t$. Following a positive earnings announcement that generates positive narrative sentiment ($\xi_t > 0$), expected price changes are positive:

$$\mathbb{E}[\Delta P_{t+1}] > 0. \tag{34}$$

Proposition 2: Drift Disappears After the Peak. When the narrative-trader share declines, the sentiment multiplier weakens. If $n_{t+1} < n_t$, then:

$$\phi_{t+1} - \phi_t \to 0 \quad \Rightarrow \quad \mathbb{E}[\Delta P_{t+1}] \to 0. \tag{35}$$

Narratives dissipate, high-precision investors dominate, and drift fades.

3.8 Transition to the Empirical Analysis

The model provides a theoretical mechanism for why earnings related price adjustment may be front-loaded but incomplete. Earnings announcements generate an immediate rational response through the fundamental component of price, P_t^{fund}. However, if the announcement also activates narrative-sensitive demand, the sentiment wedge $\phi_t \xi_t$ can generate temporary post-announcement drift.

The empirical section tests the observable implication of this mechanism. If prices adjust immediately and completely, earnings surprise ranks should explain announcement window returns but should not predict returns after the announcement. If adjustment is incomplete, earnings surprise ranks should continue to predict post-announcement abnormal returns. The following section therefore constructs analyst based SUE ranks, event time abnormal return windows, price formation regressions, and tradability checks to evaluate whether modern U.S. equity markets exhibit this front loaded but incomplete adjustment pattern.

4 Data and Empirical Design

This section describes the construction of the dataset used to test whether post-earnings announcement drift remains present in U.S. equities during the modern 2018–2024 period. The empirical design follows an event-study structure. I first identify earnings announcement events, measure the analyst-based earnings surprise associated with each event, and then track daily market-adjusted stock returns around the announcement date. The central question is whether firms with stronger positive earnings surprises continue to earn higher abnormal returns after the announcement window, and whether firms with negative earnings surprises continue to underperform.

The analysis is built around three layers of evidence. First, I construct an event-level panel linking earnings announcements, analyst-based earnings surprises, daily stock returns, market returns, and firm characteristics. Second, I sort firms into earnings-surprise portfolios and examine whether compounded post-announcement abnormal returns differ systematically across surprise deciles. Third, I use the resulting event-time patterns to motivate formal regression tests that distinguish immediate price formation from delayed post-announcement drift. This structure allows the paper to evaluate not only whether PEAD is statistically visible, but also whether the pattern survives more conservative tests related to size, transaction costs, calendar-time portfolio returns, and factor-adjusted alpha.

A key feature of the design is that the analysis separates the announcement-window reaction from the post-announcement drift window. The announcement window captures the immediate market response to earnings news, while the post-announcement window tests whether prices continue adjusting after the initial reaction. This distinction is important because evidence of a large announcement reaction alone would be consistent with fast information incorporation, while continued abnormal returns after the announcement would suggest incomplete price adjustment. The empirical design is closely related to Martineau (2021), who emphasizes the importance of separating announcement-date price formation from post-announcement drift when studying PEAD in modern markets. Following this logic, I distinguish the immediate announcement-window reaction, measured by

$BHAR[0, 1]$, from the post-announcement drift windows, especially $BHAR[2, 15]$ and $BHAR[2, 60]$. I also use analyst-based earnings surprises, ranked surprise portfolios, market-cap heterogeneity, and tradability checks to evaluate whether any remaining PEAD is economically meaningful or mainly concentrated in less liquid stocks.

4.1 Data Sources and Sample Construction

The empirical analysis is built from a firm-event panel that combines earnings announcement information, analyst-based earnings surprises, daily stock returns, market returns, and firm size classifications. The main unit of observation is a firm-earnings-announcement event. For each earnings announcement, I observe the announcement date, the firm identifier, the analyst-based earnings surprise, daily stock returns around the announcement, and the corresponding market return used to construct abnormal returns.

The final event panel covers the period from October 2, 2018 to December 31, 2024. The event-time window runs from 10 trading days before the earnings announcement to 62 trading days after the announcement, or $t = -10$ to $t = +62$. This structure allows the analysis to separate short pre-announcement price movement, the immediate announcement-window reaction, and post-announcement drift. The full daily event panel contains 4,392,191 firm-day observations, 60,167 earnings announcement events, and 4,369 unique firms.

The dataset is organized in event time. For each firm i and earnings announcement j, event day $t = 0$ corresponds to the earnings announcement date. Negative event days capture the short pre-announcement period, while positive event days capture the market reaction after the announcement. The main post-announcement drift window is defined as BHAR[2,60], which tracks cumulative abnormal performance from trading day +2 to trading day +60. Starting the main drift window at day +2 avoids mixing the immediate announcement reaction with the delayed price adjustment that PEAD is meant to capture.

The raw daily event panel is then converted into an event-level dataset. For each earnings event, I compute daily market-adjusted abnormal returns and compound them over selected event windows. The main event-level sample requires nonmissing earnings surprise information, return data, market return data, and market-cap classification. Because some specifications require additional filters, such as previous-quarter surprise-rank cutoffs or factor-return availability, the exact number of observations differs slightly across descriptive, event-time, calendar-time, and factor-adjusted tests. These differences are expected and reflect the specific data requirements of each test.

A key advantage of this sample is that it focuses on a recent period of U.S. equity markets. Unlike classic PEAD studies that rely heavily on earlier decades, this sample covers the COVID/stimulus period and the 2022–2024 rate-hike regime. This makes the dataset useful for testing whether PEAD remains visible in a modern market environment with faster information dissemination, greater algorithmic trading, and wider use of analyst forecast data.

Table 1: Sample Construction Summary

Sample Item	**Value**
Sample period	2018-10-02 to 2024-12-31
Event-time window	-10 to +62 trading days
Daily firm-day observations	4,392,191
Earnings announcement events	60,167
Unique firms	4,369
Main drift window	BHAR[2,60]

Table 1 summarizes the main event-study dataset used in the empirical analysis. The main drift window begins on event day +2 to isolate post-announcement drift from the immediate earnings-announcement reaction.

Figure 2: Earnings Events per Year

Figure 2 reports the annual distribution of earnings announcement events in the sample. The sample spans October 2018 through December 2024 and includes both the COVID/stimulus period and the 2022–2024 rate-hike regime. This provides a modern testing environment for evaluating whether post-earnings announcement drift remains present in recent U.S. equity data.

4.2 Earnings Surprise Measure: Analyst-Based SUE

The main earnings-surprise measure is analyst-based standardized unexpected earnings, or SUE. In this paper, SUE is implemented as a price-scaled analyst forecast error. For each firm i and earnings announcement j, I measure the earnings surprise as the difference between actual reported earnings per share and the consensus analyst forecast of earnings per share, scaled by the firm's stock price:

$$SUE_{i,j} = \frac{Actual\ EPS_{i,j} - Consensus\ Analyst\ Forecast\ EPS_{i,j}}{Stock\ Price_{i,j}}$$

I refer to this variable as SUE throughout the paper, but more precisely it is a price-scaled analyst-based earnings surprise. A positive SUE indicates that the firm reported earnings above analyst expectations, while a negative SUE indicates that the firm reported earnings below analyst expectations. Scaling the forecast error by stock price makes the surprise measure more comparable across firms with different share prices and earnings levels.

The intuition is straightforward. If markets fully and immediately incorporate earnings news, then stock prices should adjust quickly when actual earnings are released. In that case, after the announcement window, high-SUE firms should not continue to earn systematically higher abnormal returns than low-SUE firms. However, if investors underreact to earnings news, then firms with positive surprises may continue to outperform after the announcement, while firms with negative surprises may continue to underperform. This continued return predictability is the core prediction of post-earnings announcement drift.

In the empirical analysis, I do not rely only on raw SUE. Raw earnings surprises can be noisy, skewed, and sensitive to extreme observations. Instead, the main analysis sorts firms into SUE deciles. Decile 10 contains the strongest positive earnings surprises, while Decile 1 contains the most negative earnings surprises. This ranked-portfolio approach reduces the influence of outliers and makes the economic interpretation clearer: the key comparison is whether the best-surprise portfolio continues to outperform the worst-surprise portfolio after earnings announcements.

The decile approach is also closer to how PEAD is commonly tested in empirical finance. Rather than asking whether one unit of raw SUE mechanically predicts returns, the analysis asks whether firms with relatively better earnings news earn higher subsequent abnormal returns than firms with relatively worse earnings news. In the formal price-formation regressions, surprise ranks are formed using prior-quarter cutoffs where appropriate, which reduces look-ahead bias by avoiding the use of future information when assigning surprise ranks. This is important because the early diagnostic regressions show that raw SUE is weak in simple linear specifications, while ranked SUE deciles produce stronger and more stable evidence of post-announcement drift.

Portfolio	Earnings Surprise Type	Interpretation
Decile 1	Most negative SUE	Worst earnings surprises
Deciles 2–9	Intermediate SUE	Moderate negative to moderate positive surprises
Decile 10	Most positive SUE	Best earnings surprises

Table 2: Interpretation of SUE Deciles

Table 2 summarizes the earnings-surprise portfolio construction. Firms are sorted into SUE deciles, where Decile 1 contains the lowest earnings surprises and Decile 10 contains the highest earnings surprises. The main PEAD comparison is the post-announcement abnormal return spread between Decile 10 and Decile 1.

4.3 Abnormal Returns and Compounded Market-Adjusted Returns

After measuring the earnings surprise, I measure the stock market response using abnormal returns. For each firm i, earnings announcement j, and event day t, the daily abnormal return is defined as the firm's stock return minus the corresponding market return:

$$AR_{i,j,t} = R_{i,t} - R_{m,t}$$

where $R_{i,t}$ is the daily stock return of firm i, and $R_{m,t}$ is the market return on the same trading day. This market-adjusted return removes broad market movement and helps isolate firm-specific return behavior around the earnings announcement. This adjustment is especially important because the sample covers volatile macroeconomic periods, including the COVID/stimulus period and the 2022–2024 rate-hike regime, when market-wide returns could otherwise distort the measured response to earnings news.

I then compound daily abnormal returns over different event windows. Throughout the paper, I refer to this measure as BHAR, but more precisely it is a compounded market-adjusted abnormal return:

$$BHAR[a,b]_{i,j} = \prod_{t=a}^{b}(1 + AR_{i,j,t}) - 1$$

where a and b define the beginning and ending event days around earnings announcement j. For example, BHAR[2,60] measures the compounded market-adjusted abnormal return from trading day +2 through trading day +60 after the earnings announcement.

This construction captures the cumulative abnormal performance of a stock over the event window after removing market-wide movements. A positive BHAR[2,60] means that the stock outperformed the market over the post-announcement window, while a negative value means that it underperformed. In the PEAD setting, the key question is whether high-SUE firms earn higher post-announcement compounded abnormal returns than low-SUE firms.

The main drift measure is BHAR[2,60]. The window begins on day +2, rather than day 0, because the goal is to study drift after the immediate earnings-announcement reaction. If prices fully adjust to earnings news immediately, then earnings surprise should strongly explain the announcement-window return but should not continue to predict BHAR[2,60]. Therefore, significant post-announcement return predictability is evidence of incomplete price adjustment.

Strictly speaking, this paper implements BHAR as compounded daily market-adjusted abnormal returns, rather than as the difference between the firm's raw buy-and-hold return and a benchmark buy-and-hold return. I use the term BHAR because the measure captures the same event-study intuition of cumulative abnormal performance over a buy-and-hold event window.

4.4 Event Windows and Empirical Design

The empirical design separates price formation into several event windows around the earnings announcement. This is necessary because PEAD is not only about whether stock prices react to earnings news, but also about when that reaction occurs. The goal is to distinguish immediate announcement-window price formation from delayed post-announcement drift.

The event date, $t = 0$, is the earnings announcement date. I use the short pre-announcement window BHAR[-10,-1] to test whether prices move in the direction of the future earnings surprise before the formal announcement. This window may capture anticipation, information leakage, analyst revisions, or gradual incorporation of information before the event date. Because the available event panel begins at event day -10, this is a short pre-announcement test rather than a long-run pre-announcement run-up measure.

The announcement-window reaction is measured using BHAR[0,1]. This window captures the immediate stock-price response to earnings news. A strong relation between earnings surprise and BHAR[0,1] indicates that the market reacts sharply when earnings information is released. This is important because the analysis does not assume that markets ignore earnings news. Instead, the key question is whether the initial reaction fully incorporates the information.

The main post-announcement drift windows are BHAR[2,15] and BHAR[2,60]. The shorter window, BHAR[2,15], tests whether drift remains during the first few weeks after the announcement. The longer window, BHAR[2,60], is the main medium-term PEAD outcome. Both windows begin on day +2 to avoid mixing the immediate announcement reaction with subsequent drift.

Finally, I use BHAR[0,60] in the price-formation and unbiasedness analysis. This window combines the immediate announcement response with the subsequent post-announcement adjustment. It helps evaluate whether the announcement-window return contains most of the information about the full 60-day price formation process. If BHAR[0,1] strongly explains BHAR[0,60], price formation is front-loaded. If earnings surprise also predicts BHAR[2,60], then the reaction is front-loaded but incomplete.

Table 3 summarizes the event windows used in the empirical analysis. The main PEAD test focuses on BHAR[2,60], which begins after the immediate announcement window and therefore isolates post-announcement drift.

4.5 Data Quality and Sample Diagnostics

I conduct several data quality checks before testing for post-earnings announcement drift. These checks focus on event matching, missing values, the distribution of earnings surprises, return availability, market-cap classification, and potential illiquidity. This step is important because PEAD tests are sensitive to noisy earnings surprise measures, missing return observations, extreme SUE values, and microcap trading frictions.

Table 3: Event Windows Used in the Empirical Analysis

Event Window	**Interpretation**	**Purpose**
BHAR[−10,−1]	Short pre-announcement window	Tests anticipation or pre-announcement price formation
BHAR[0,1]	Announcement-window reaction	Measures immediate market response to earnings news
BHAR[2,15]	Short post-announcement window	Tests short-run PEAD after the initial reaction
BHAR[2,60]	Main post-announcement drift window	Main outcome for medium-term PEAD
BHAR[0,60]	Full price-formation window	Compares immediate reaction with full-window adjustment

The first diagnostic verifies whether earnings announcement events are successfully matched to the return panel. The main BHAR[2,60] construction retains approximately 60,135 events, implying a clean event match rate of about 99.95% relative to the original earnings event sample. This indicates that almost all earnings events are linked to the return observations needed to compute post-announcement abnormal returns. The high match rate reduces the risk that the results are driven by event-matching failures.

I also examine missing values in the core variables used throughout the analysis: earnings surprise, stock return, market return, and market-cap bucket. These variables are required to construct SUE deciles, abnormal returns, and market-cap heterogeneity tests. The missing-value audit also helps explain why observation counts differ across specifications. Some regressions require additional filters, such as previous-quarter surprise-rank cutoffs, factor-return availability, or complete event-window observations.

A second diagnostic concerns the distribution of SUE. Earnings surprises are often skewed and can contain extreme values. This matters because raw SUE regressions may be sensitive to outliers. In this dataset, the SUE distribution is centered close to zero, suggesting that analyst forecast errors are not mechanically biased in one direction. However, the tails of the distribution remain important. This motivates the use of ranked SUE deciles rather than relying only on raw continuous

SUE.

I also track zero-return days as an illiquidity diagnostic. Zero-return observations are especially relevant in smaller stocks, where trading may be infrequent and bid-ask spreads may be large. This matters because a PEAD pattern can be statistically visible but economically difficult to trade if it is concentrated in illiquid securities. For this reason, the later analysis separates event-time return predictability from tradable alpha.

Finally, I examine average portfolio characteristics across SUE deciles. This diagnostic checks whether the earnings-surprise portfolios mechanically line up with other firm characteristics, especially market-cap bucket and, where available, book-to-market or momentum. The goal is not to prove factor independence at this stage. Instead, the purpose is to identify whether the SUE sorting appears to be merely a simple size, value, or momentum sort. Formal factor-adjusted tests are reported later.

Table 4: Data Quality and Sample Diagnostics

Diagnostic Item	**Value**
Main event-level observations for BHAR[2,60]	60,135
Clean event match rate	99.95%
Missing SUE values	2,336
Missing return/market values	0
Missing market-cap bucket values	0
Zero-return days	64,299
Extreme SUE outliers ($\lvert SUE \rvert > 1$)	274
SUE median	0.000849
SUE 1st percentile	-0.264595
SUE 99th percentile	0.185341

Table 4 reports key diagnostic checks for the event-study dataset. The high event match rate indicates strong alignment between earnings announcements and daily return data. The missing-value, SUE distribution, and illiquidity checks identify potential limitations relevant for interpreting PEAD and tradability.

5 Empirical Evidence and Results

5.1 Descriptive Evidence: Drift by Earnings Surprise Decile

After constructing the event-level sample, I first examine whether post-announcement abnormal returns vary systematically across earnings-surprise deciles. This provides a direct descriptive test of post-earnings announcement drift. If PEAD is present, firms in the highest earnings-surprise deciles should earn higher post-announcement abnormal returns than firms in the lowest earnings-surprise deciles.

Figure 3 plots the mean BHAR[2,60] for each SUE decile. The pattern provides the first visual evidence of PEAD in the sample. The highest SUE decile earns the strongest positive post-announcement abnormal return, while most lower and middle deciles earn negative abnormal returns. This suggests that earnings surprise continues to contain information about returns after the immediate announcement window.

The most important comparison is between Decile 10 and Decile 1. Decile 10 represents the strongest positive earnings surprises, while Decile 1 represents the most negative earnings surprises. In the descriptive table, Decile 10 earns a mean 60-day abnormal return of 2.867%, compared with -1.282% for Decile 1, implying a simple Decile 10–Decile 1 spread of approximately 4.149 percentage points. This spread is the core descriptive signature of PEAD: good-news firms continue to outperform bad-news firms after earnings announcements.

At the same time, the pattern should be described carefully. The decile profile is not perfectly monotonic across all ten portfolios. Instead, the strongest evidence appears in the tails: the highest surprise portfolio performs clearly better than the low- and middle-surprise portfolios. This supports the use of ranked SUE portfolios while also motivating the formal regression tests that follow later.

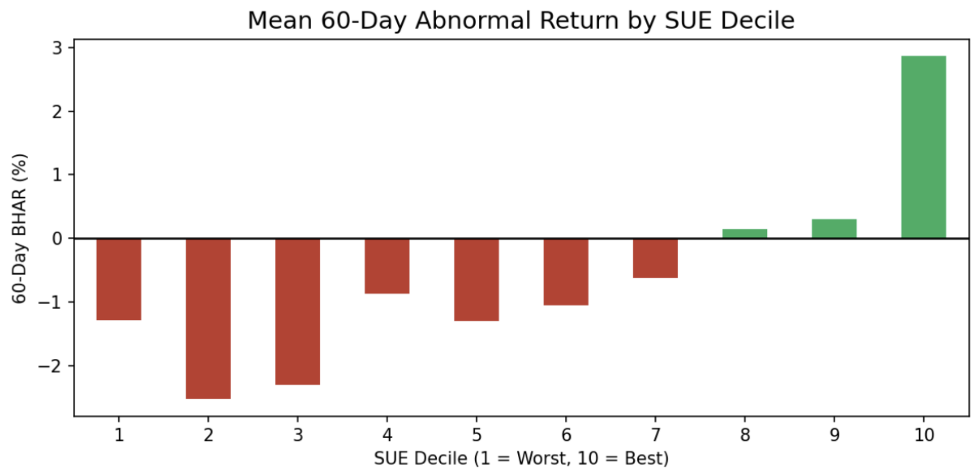

Figure 3: Mean 60-Day Abnormal Return by SUE Decile. The figure reports mean BHAR[2,60] by analyst-based SUE decile. Decile 1 contains the worst earnings surprises, while Decile 10 contains the best earnings surprises. The highest SUE decile earns the strongest positive post-announcement abnormal return, providing descriptive evidence that good-news firms continue to outperform after earnings announcements.

5.2 Drift Dynamics: Daily CAR Path by SUE Decile

The decile-level bar chart shows the terminal 60-day abnormal return, but it does not show how the drift develops over time. To examine the mechanism of PEAD more directly, I plot the cumulative abnormal return path from event day +2 to event day +60 for Decile 1, Decile 5, and Decile 10.

Figure 4 shows that the spread between the best-surprise and worst-surprise portfolios widens gradually after the announcement. This matters because PEAD is not simply a one-day reaction to earnings news. If prices fully adjusted immediately, the abnormal return spread should appear at the announcement and then stop widening. Instead, the figure shows that high-surprise firms continue to earn higher abnormal returns over the post-announcement window.

This gradual separation is the central mechanism behind the anomaly. The market reacts to earnings news, but the adjustment appears incomplete. The Decile 10 portfolio continues to outperform, while the Decile 1 portfolio remains weaker. The result is a growing long-short spread over the post-announcement period.

The CAR path therefore complements the terminal decile bar chart. It shows not only that high-SUE firms outperform over the full 60-day window, but also that the outperformance accumulates gradually. This is consistent with delayed price adjustment rather than a purely instantaneous announcement effect.

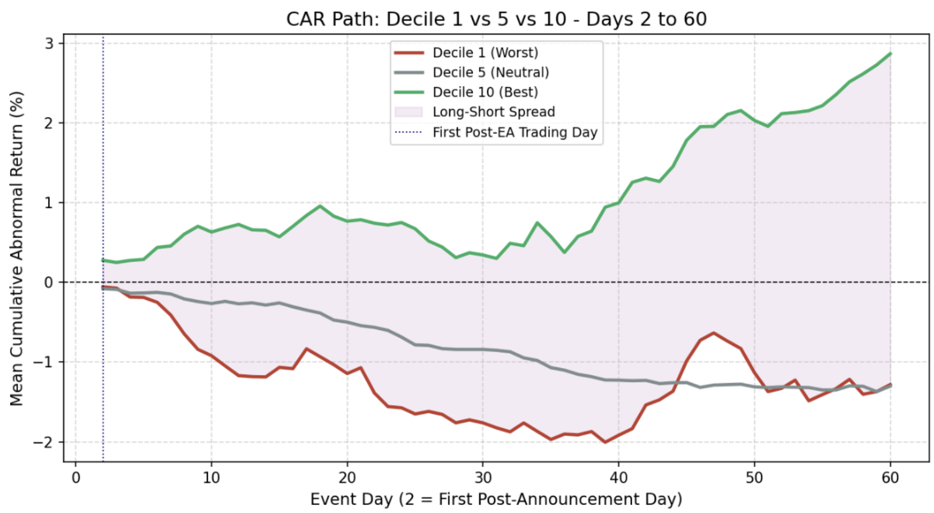

Figure 4: Daily Cumulative Abnormal Return Path by SUE Decile. The figure plots the mean cumulative market-adjusted abnormal return from event day +2 to event day +60 for Decile 1, Decile 5, and Decile 10. The gradual widening between the best-surprise and worst-surprise portfolios suggests that price adjustment continues after the announcement window, consistent with post-earnings announcement drift.

5.3 Formal Price-Formation Tests

The descriptive evidence shows that post-announcement abnormal returns differ across earnings-surprise deciles. To test this pattern more formally, I estimate price-formation regressions that relate abnormal returns over different event windows to the firm's earnings-surprise rank. The baseline specification is:

$$BHAR[a,b]_{i,j} = \alpha_i + \alpha_q + \beta SurpriseRank_{i,j} + \varepsilon_{i,j}, \tag{36}$$

where $BHAR[a,b]_{i,j}$ is the compounded market-adjusted abnormal return for firm i around earnings announcement j, $SurpriseRank_{i,j}$ is the earnings-surprise decile rank, α_i represents firm fixed effects, and α_q represents earnings-quarter fixed effects. Standard errors are clustered by firm and earnings quarter.

The purpose of this regression is to identify where price formation occurs. If earnings information is incorporated immediately and completely, then $SurpriseRank$ should strongly predict announcement-window returns but should not predict abnormal returns after the announcement. If $SurpriseRank$ continues to predict $BHAR[2,15]$ or $BHAR[2,60]$, then the evidence supports incomplete price adjustment after the earnings announcement.

The results show a front-loaded but incomplete price reaction. The coefficient on $SurpriseRank$ is largest for $BHAR[0,1]$, indicating a strong immediate response to earnings news. However, the coefficients remain positive and statistically significant for both $BHAR[2,15]$ and $BHAR[2,60]$. This means that although the market reacts quickly to earnings information, the initial response does not fully eliminate return predictability after the announcement.

Table 5: Price Formation Around Earnings Announcements

Outcome	**Coef. on SurpriseRank**	**T-stat**	**P-value**	**D10–D1 Spread**	**Interpretation**
$BHAR[0,1]$	0.933	22.65	< 0.001	8.40 pp	Strong immediate reaction
$BHAR[2,15]$	0.135	3.22	0.001	1.22 pp	Short-run drift remains
$BHAR[2,60]$	0.378	3.52	< 0.001	3.40 pp	Medium-run PEAD remains

Note: D10–D1 Spread refers to the implied Decile 10–Decile 1 spread. The spread equals nine times the coefficient because moving from Decile 1 to Decile 10 is a nine-rank increase.

Table 5 reports price-formation regressions of compounded market-adjusted abnormal returns on earnings-surprise rank. The dependent variable is measured in percentage points. The implied Decile 10–Decile 1 spread equals nine times the coefficient because moving from Decile 1 to Decile 10 is a nine-rank increase. Firm fixed effects and earnings-quarter fixed effects are absorbed, and standard errors are clustered by firm and earnings quarter.

The economic interpretation is straightforward. The $BHAR[0, 1]$ coefficient of 0.933 means that moving up by one SUE decile is associated with approximately 0.933 percentage points higher announcement-window abnormal return. Moving from Decile 1 to Decile 10 therefore implies an announcement-window spread of about $9 \times 0.933 = 8.40$ percentage points. This confirms that earnings news is incorporated strongly and quickly. However, the post-announcement coefficients are also positive and statistically significant. The $BHAR[2, 15]$ coefficient implies a 1.22 percentage-point Decile 10–Decile 1 spread, while the $BHAR[2, 60]$ coefficient implies a 3.40 percentage-point spread over the main drift window. Therefore, the evidence supports a nuanced conclusion: price formation is front-loaded, but not complete.

I also test whether the post-announcement pattern is simply a continuation of pre-announcement price movement. Because the event panel begins at event day -10, I use $BHAR[-10, -1]$ as a short pre-announcement window. The pre-announcement return is positively related to future earnings-surprise rank, suggesting that some price discovery occurs before the formal announcement. However, the post-announcement coefficient on $SurpriseRank$ remains positive and statistically significant after controlling for $BHAR[-10, -1]$. This indicates that the main PEAD result is not explained solely by short pre-announcement price formation.

5.4 Tradability and Factor-Adjusted Alpha

The event-time evidence shows that PEAD remains statistically visible, but statistical drift is not the same as tradable alpha. To address this distinction, I conduct two additional checks. First, I examine the PEAD long-short spread by market-cap tier after estimated transaction costs. Second, I form a calendar-time long-short PEAD portfolio and test whether it earns alpha after standard factor controls.

The market-cap analysis sorts firms into SUE deciles within each size bucket, then compares Decile 10 with Decile 1. This is a more direct tradability test than simply comparing average returns by market-cap group, because a PEAD strategy requires going long good-surprise firms and short bad-surprise firms. I then compare the long-short spread with estimated round-trip transaction costs. This test asks whether the return spread is large enough to survive realistic trading frictions.

Mkt Cap	D10 (%)	D1 (%)	L/S Spread (%)	RT TC (%)	Net Spread (%)	D10 N	D1 N
Nano	2.168%	-2.742%	4.910%	4.00%	**0.910%**	2,280	2,280
Micro	5.386%	-1.415%	6.801%	1.90%	**4.901%**	1,227	1,227
Small	2.387%	-1.652%	4.039%	1.10%	**2.939%**	925	925
Mid	1.441%	-1.283%	2.724%	0.40%	**2.324%**	814	814
Large	1.450%	-1.258%	2.709%	0.16%	**2.549%**	770	770

** RT TC = estimated round-trip transaction cost, approximated as 2x one-way bid-ask spread. Net Spread = D10-D1 Long-Short Spread minus estimated round-trip transaction cost. This is a friction-adjusted event-time spread test, not definitive proof of implementable trading alpha.*

Figure 5: PEAD Long-Short Spread and Transaction Costs by Market-Cap Tier. The figure reports the Decile 10–Decile 1 PEAD spread within each market-cap tier and compares the spread with estimated round-trip transaction costs. The figure distinguishes statistical return predictability from economic tradability by asking whether the long-short PEAD spread survives transaction-cost estimates.

The market-cap tradability test shows that the PEAD long-short spread remains positive after rough round-trip transaction-cost estimates across all reported size buckets. The strongest estimated net spread appears in microcap stocks, while nano stocks are more heavily reduced by transaction costs. However, this table should still be interpreted cautiously. The transaction-cost adjustment captures only an estimated bid-ask spread effect and does not fully account for market impact, short-sale constraints, borrow costs, capacity limits, slippage, or implementation timing. Therefore, the evidence suggests that PEAD may remain economically meaningful in event time, but it does not by itself prove easily implementable trading alpha.

To test tradable alpha more rigorously, I also estimate calendar-time and factor-adjusted alpha tests. These tests are stricter because they convert the event-time PEAD pattern into a time-series long-short portfolio and then test whether the portfolio earns alpha after accounting for standard risk factors. The results are positive but not statistically significant across the main daily and monthly specifications. This means that PEAD is easier to detect in event time than as a clean tradable factor-adjusted strategy.

Table 6: Calendar-Time and Factor-Adjusted PEAD Alpha

Test	Alpha	T-statistic	P-value
Daily calendar-time alpha	0.032%	1.10	0.269
Daily CAPM alpha	0.032%	1.09	0.275
Daily Fama-French 3 alpha	0.030%	1.10	0.269
Daily Fama-French 5 alpha	0.026%	0.96	0.336
Monthly compounded alpha	0.662%	1.10	0.271
Monthly FF5 + Momentum alpha	0.349%	0.58	0.559

Table 6 reports calendar-time and factor-adjusted PEAD alpha tests. Daily tests use Newey-West standard errors with 60 lags, while monthly tests use Newey-West standard errors with 3 lags. The alphas are positive but not statistically significant, suggesting that event-time PEAD is more visible than tradable factor-adjusted alpha.

5.5 Summary of Empirical Evidence

Overall, the empirical evidence supports a front-loaded but incomplete price-adjustment interpretation of PEAD. Earnings-surprise rank has its largest effect around the announcement window, showing that markets react strongly to earnings news. However, surprise rank continues to predict abnormal returns in the post-announcement windows, especially $BHAR[2, 15]$ and $BHAR[2, 60]$, indicating that the initial reaction is not complete. At the same time, the tradability evidence is more conservative: calendar time and factor adjusted alpha tests are positive but not statistically significant. The evidence therefore supports statistical event time PEAD more strongly than clean, implementable trading alpha.

6 Conclusion

This paper revisits post-earnings announcement drift in modern U.S. equity markets from both a theoretical and empirical perspective. The theoretical framework develops a narrative-diffusion mechanism in which all investors remain Bayesian, but the market share of narrative-sensitive traders evolves over time. Retail traders may adopt narrative-driven behavior through interaction, while narrative traders may eventually transition into more sophisticated types. This random-matching structure generates a hump-shaped evolution of narrative participation. Because only narrative traders respond directly to sentiment, the price impact of sentiment becomes time varying: it strengthens when narratives spread and weakens when narratives dissipate. The model therefore provides a mechanism for a temporary jump–drift–fade pattern without requiring permanently irrational beliefs or permanent mispricing.

The empirical analysis then tests whether the core return pattern associated with PEAD remains visible in a modern U.S. equity sample from 2018 to 2024. Using analyst based earnings surprises, event time abnormal returns, SUE decile portfolios, and price formation regressions, the paper examines whether firms with stronger positive earnings surprises continue to outperform firms with weaker or negative surprises after the announcement date.

The main empirical finding is that PEAD has not disappeared, but its modern form is more nuanced than the classic version of the anomaly. Earnings news is incorporated quickly around the announcement window. The coefficient on earnings-surprise rank is largest for $BHAR[0,1]$, showing that most price formation occurs immediately when earnings information becomes public. This result is important because it shows that modern markets do react strongly to earnings news; the evidence is not consistent with a simple story in which investors ignore earnings announcements.

At the same time, the announcement-window reaction is not complete. Earnings surprise rank continues to predict abnormal returns in both the short post-announcement window, $BHAR[2,15]$, and the main medium-term drift window, $BHAR[2,60]$. The descriptive decile evidence, cumulative abnormal return paths, and formal price formation regressions all point to the same conclusion: high surprise firms continue to outperform after the announcement, while lower surprise firms perform worse. The most accurate interpretation is therefore front-loaded but incomplete price adjustment.

This empirical pattern is consistent with the theoretical mechanism developed in the model. The model predicts that prices should respond immediately to fundamental earnings information, but that additional temporary drift can arise when sentiment sensitive demand is amplified by evolving market composition. The empirical evidence does not directly prove the narrative diffusion channel, because the final tests focus on earnings surprises and event time abnormal returns rather than direct sentiment measures. However, the observed pattern of strong immediate price formation followed by statistically significant post-announcement drift is exactly the kind of return dynamic that the model is designed to explain.

The tradability evidence is more conservative. Event time return predictability does not automatically imply easy trading profits. The market cap tradability analysis suggests that PEAD spreads can remain positive after rough transaction cost estimates, but these estimates do not fully capture market impact, short sale constraints, borrow costs, capacity limits, slippage, or implementation timing. When the analysis is moved into calendar time and adjusted for standard risk factors, the estimated alphas remain positive but are not statistically significant. This means that PEAD is easier to detect as an event time price formation pattern than as a clean, implementable factor adjusted trading strategy.

Overall, the safest conclusion is that PEAD is not fully dead, but it is weaker, faster, and harder to monetize than the classic anomaly. In the 2018–2024 sample, earnings information is absorbed rapidly, but not perfectly. The anomaly survives most clearly in ranked earnings surprise portfolios and event time abnormal returns, especially among firms with strong positive earnings surprises. However, once transaction costs, liquidity frictions, calendar time portfolio construction, and factor controls are considered, the evidence for easily harvestable alpha becomes less conclusive.

This paper therefore contributes to the market efficiency debate by showing that modern markets are neither perfectly efficient nor plainly inefficient. They process earnings news quickly, but some predictable adjustment remains after the announcement. The theoretical model offers one explanation for why this adjustment may be temporary: narrative sensitive demand can amplify earnings news for a period before fading as market composition changes. The empirical evidence supports the broader interpretation that PEAD remains statistically visible and economically meaningful, but is harder to convert into reliable trading alpha in modern markets.

References

Ball, R. and Brown, P. (1968). An empirical evaluation of accounting income numbers. *Journal of Accounting Research*, 6(2):159–178.

Barberis, N., Shleifer, A., and Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, 49(3):307–343.

Bernard, V. L. and Thomas, J. K. (1989). Post-earnings-announcement drift: Delayed price response or risk premium? *Journal of Accounting Research*, 27(Suppl.):1–36.

Bernard, V. L. and Thomas, J. K. (1990). Evidence that stock prices do not fully reflect the implications of current earnings for future earnings. *Journal of Accounting and Economics*, 13(4):305–340.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2022). Overreaction and diagnostic expectations in macroeconomics. *Journal of Economic Perspectives*, 36(3):223–244.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *Quarterly Journal of Economics*, 129(4):1661–1710.

Hirshleifer, D. and Teoh, S. H. (2003). Limited attention, information disclosure, and financial reporting. *Journal of Accounting and Economics*, 36(1–3):337–386.

Hong, H. and Stein, J. C. (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance*, 54(6):2143–2184.

Martineau, C. (2021). Rest in peace post-earnings announcement drift. *SSRN Working Paper*. Available at SSRN: https://ssrn.com/abstract=3111607.

Peng, L. and Xiong, W. (2006). Investor attention, overconfidence and category learning. *Journal of Financial Economics*, 80(3):563–602.

Shiller, R. J. (2017). Narrative economics. *American Economic Review*, 107(4):967–1004.

Appendix A CARA–Normal Demand

This appendix derives the optimal asset demand for each investor type under CARA utility and normally distributed beliefs. All algebraic steps omitted from the main text are shown here.

A.1 Setup

Investor of type $k \in \{R, N, H\}$ has beliefs

$$\theta_{t+1} \mid I_{i,t}^k \sim N(m_{i,t}^k,\ v_t^k). \tag{A1}$$

Wealth evolves according to

$$W_{t+1} = W_t + X_{i,t}^k(\theta_{t+1} - P_t). \tag{A2}$$

Preferences are CARA:

$$U(W_{t+1}) = -e^{-\gamma W_{t+1}}. \tag{A3}$$

A.2 Objective

The investor solves

$$\max_{X_{i,t}^k}\ \mathbb{E}_t\left[-e^{-\gamma W_{t+1}}\right]. \tag{A4}$$

Using (A2),

$$-e^{-\gamma W_t}\ \mathbb{E}_t\left[e^{-\gamma X_{i,t}^k(\theta_{t+1}-P_t)}\right]. \tag{A5}$$

This is equivalent to

$$\min_{X_{i,t}^k} \mathbb{E}_t\left[e^{-\gamma X_{i,t}^k(\theta_{t+1}-P_t)}\right]. \tag{A6}$$

A.3 CARA–Normal Formula

Since

$$\theta_{t+1} - P_t \sim N(m_{i,t}^k - P_t,\ v_t^k),$$

define

$$Y = X_{i,t}^{k}(\theta_{t+1} - P_t). \tag{A7}$$

Then

$$Y \sim N\left(X_{i,t}^{k}(m_{i,t}^{k} - P_t),\ (X_{i,t}^{k})^2 v_t^k\right). \tag{A8}$$

The CARA–Normal expectation satisfies

$$\mathbb{E}\left[e^{-\gamma Y}\right] = \exp\left(-\gamma X_{i,t}^{k}(m_{i,t}^{k} - P_t) + \frac{\gamma^2}{2}(X_{i,t}^{k})^2 v_t^k\right). \tag{A9}$$

A.4 First-Order Condition

Minimize the exponent in (A9):

$$\frac{\partial}{\partial X}\left[-\gamma X(m_{i,t}^{k} - P_t) + \frac{\gamma^2}{2}X^2 v_t^k\right] = -\gamma(m_{i,t}^{k} - P_t) + \gamma^2 X v_t^k. \tag{A10}$$

FOC:

$$-\gamma(m_{i,t}^{k} - P_t) + \gamma^2 X_{i,t}^{k} v_t^k = 0. \tag{A11}$$

Solution:

$$X_{i,t}^{k} = \frac{m_{i,t}^{k} - P_t}{\gamma v_t^k}. \tag{A12}$$

A.5 Type-Average Demand

With a continuum of investors, replace $m_{i,t}^{k}$ with type-k mean belief $\bar{m}_t^k$:

$$X_t^k = \frac{\bar{m}_t^k - P_t}{\gamma v_t^k}. \tag{A13}$$

This is the expression used in the main model section.

Appendix B Sentiment Demand for Narrative Traders

This appendix derives the optimal demand of narrative traders, who experience a sentiment shift $\beta\xi_t$ in perceived payoffs. All algebraic steps omitted from the main text are provided below.

B.1 Utility with Sentiment Shift

Narrative traders have utility

$$U_N(W_{t+1}) = -\exp\left\{-\gamma\left(W_{t+1} + \beta\xi_t X_{i,t}^N\right)\right\}. \tag{B1}$$

Substituting the wealth equation (A2),

$$U_N(W_{t+1}) = -\exp\Big(-\gamma W_t - \gamma X_{i,t}^N(\theta_{t+1} - P_t + \beta\xi_t)\Big). \tag{B2}$$

Equivalent objective:

$$\min_{X_{i,t}^N} \mathbb{E}_t\left[e^{-\gamma X_{i,t}^N(\theta_{t+1} - P_t + \beta\xi_t)}\right]. \tag{B3}$$

B.2 Shifted Payoff Distribution

Define the shifted payoff:

$$\theta_{t+1} - P_t + \beta\xi_t \sim N\Big(m_{i,t}^N - P_t + \beta\xi_t,\ v_t^N\Big). \tag{B4}$$

This is the only modification relative to the CARA–Normal case in Appendix A.

B.3 First-Order Condition

Applying the CARA–Normal expectation formula, the derivative of the exponent gives

$$-\gamma(m_{i,t}^N - P_t + \beta\xi_t) + \gamma^2 X_{i,t}^N v_t^N = 0. \tag{B5}$$

The solution is

$$X_{i,t}^N = \frac{m_{i,t}^N - P_t}{\gamma v_t^N} + \frac{\beta\xi_t}{\gamma v_t^N}. \tag{B6}$$

The first term is standard CARA demand; the second term is sentiment-driven excess demand.

B.4 Aggregate Sentiment Demand

With a continuum of narrative traders of mass n_t, aggregate sentiment-driven demand is:

$$X_t^{sent} = \beta n_t \xi_t. \tag{B7}$$

The term β absorbs the constant $\frac{1}{\gamma v_t^N}$ for notational simplicity.

This expression is used in the equilibrium price characterization in the main model.

Appendix C Market Clearing and Price

This appendix derives the market-clearing price and the decomposition of price into a fundamental component and a sentiment component.

C.1 Total Demand

Type-k demand from Appendix A is:

$$X_t^k = \frac{\bar{m}_t^k - P_t}{\gamma v_t^k}, \qquad k \in \{R, N, H\}. \tag{C1}$$

Aggregate rational demand is:

$$X_t^{\text{rat}} = r_t \frac{\bar{m}_t^R - P_t}{\gamma v_t^R} + n_t \frac{\bar{m}_t^N - P_t}{\gamma v_t^N} + h_t \frac{\bar{m}_t^H - P_t}{\gamma v_t^H}. \tag{C2}$$

Sentiment demand from narrative traders is:

$$X_t^{\text{sent}} = \beta n_t \xi_t. \tag{C3}$$

Total demand is:

$$X_t^{\text{tot}} = X_t^{\text{rat}} + X_t^{\text{sent}}. \tag{C4}$$

Market clearing, with net supply normalized to one, requires:

$$X_t^{\text{tot}} = 1. \tag{C5}$$

Substituting (C2)–(C3) into (C5):

$$r_t \frac{\bar{m}_t^R - P_t}{\gamma v_t^R} + n_t \frac{\bar{m}_t^N - P_t}{\gamma v_t^N} + h_t \frac{\bar{m}_t^H - P_t}{\gamma v_t^H} + \beta n_t \xi_t = 1. \tag{C6}$$

Multiplying (C6) by γ:

$$r_t \frac{\bar{m}_t^R}{v_t^R} - r_t \frac{P_t}{v_t^R} + n_t \frac{\bar{m}_t^N}{v_t^N} - n_t \frac{P_t}{v_t^N} + h_t \frac{\bar{m}_t^H}{v_t^H} - h_t \frac{P_t}{v_t^H} + \gamma \beta n_t \xi_t = \gamma. \tag{C7}$$

C.2 Aggregate Precision Weight

Collect the coefficients on P_t and define:

$$\Omega_t \equiv \frac{r_t}{v_t^R} + \frac{n_t}{v_t^N} + \frac{h_t}{v_t^H}. \tag{C8}$$

Then (C7) can be written as:

$$-P_t \Omega_t + \left(r_t \frac{\bar{m}_t^R}{v_t^R} + n_t \frac{\bar{m}_t^N}{v_t^N} + h_t \frac{\bar{m}_t^H}{v_t^H} \right) + \gamma \beta n_t \xi_t = \gamma. \tag{C9}$$

C.3 Price Decomposition

Rearranging (C9) and solving for P_t:

$$P_t \Omega_t = r_t \frac{\bar{m}_t^R}{v_t^R} + n_t \frac{\bar{m}_t^N}{v_t^N} + h_t \frac{\bar{m}_t^H}{v_t^H} + \gamma \beta n_t \xi_t - \gamma. \tag{C10}$$

Hence:

$$P_t = \frac{r_t \frac{\bar{m}_t^R}{v_t^R} + n_t \frac{\bar{m}_t^N}{v_t^N} + h_t \frac{\bar{m}_t^H}{v_t^H}}{\Omega_t} - \frac{\gamma}{\Omega_t} + \frac{\gamma \beta n_t}{\Omega_t} \xi_t. \tag{C11}$$

Define the fundamental component:

$$P_t^{\text{fund}} \equiv \frac{r_t \frac{\bar{m}_t^R}{v_t^R} + n_t \frac{\bar{m}_t^N}{v_t^N} + h_t \frac{\bar{m}_t^H}{v_t^H}}{\Omega_t} - \frac{\gamma}{\Omega_t}, \tag{C12}$$

and the sentiment multiplier:

$$\phi_t \equiv \frac{\gamma\beta n_t}{\Omega_t}. \tag{C13}$$

Then the market-clearing price admits the decomposition:

$$\boxed{P_t = P_t^{\text{fund}} + \phi_t \xi_t}. \tag{C14}$$

Appendix D Dynamics of Type Shares

This appendix derives the transition properties of the population shares introduced in the main text. All algebraic steps and qualitative implications are shown explicitly.

D.1 Change in the Narrative Share

From the law of motion:

$$n_{t+1} = n_t + \lambda r_t n_t - \alpha n_t h_t,$$

the change in the narrative share is:

$$\Delta n_t = n_{t+1} - n_t = \lambda r_t n_t - \alpha n_t h_t.$$

Factoring out n_t:

$$\Delta n_t = n_t(\lambda r_t - \alpha h_t). \tag{D1}$$

D.2 Hump-Shaped Dynamics

Because $n_t \geq 0$, the sign of Δn_t depends on:

$$\lambda r_t - \alpha h_t.$$

Growth occurs when:

$$\lambda r_t > \alpha h_t \quad \Rightarrow \quad n_{t+1} > n_t.$$

The peak occurs when:

$$\lambda r_t = \alpha h_t \quad \Rightarrow \quad n_{t+1} = n_t.$$

Decline occurs when:

$$\lambda r_t < \alpha h_t \quad \Rightarrow \quad n_{t+1} < n_t.$$

Since r_t decreases monotonically and h_t increases monotonically, the inequality eventually reverses. Therefore, the narrative-trader share follows a hump-shaped path:

$$n_t \text{ is hump-shaped.} \tag{D2}$$

D.3 Monotonic Decline of Retail Share

From:

$$r_{t+1} = r_t - \lambda r_t n_t,$$

we obtain:

$$\Delta r_t = -\lambda r_t n_t < 0 \quad \text{whenever } n_t > 0. \tag{D3}$$

Thus, the retail share declines over time:

$$r_t \rightarrow 0. \tag{D4}$$

D.4 Monotonic Rise of High-Precision Share

High-precision types evolve as:

$$h_{t+1} = h_t + \alpha n_t h_t,$$

implying:

$$\Delta h_t = \alpha n_t h_t \geq 0. \tag{D5}$$

With $r_t \to 0$ and $n_t \to 0$, the high-precision share converges to one:

$$h_t \to 1. \tag{D6}$$

D.5 Mass Conservation

Using the transition system:

$$r_{t+1} = r_t - \lambda r_t n_t,$$

$$n_{t+1} = n_t + \lambda r_t n_t - \alpha n_t h_t,$$

$$h_{t+1} = h_t + \alpha n_t h_t,$$

sum the three equations:

$$r_{t+1} + n_{t+1} + h_{t+1} = (r_t - \lambda r_t n_t) + (n_t + \lambda r_t n_t - \alpha n_t h_t) + (h_t + \alpha n_t h_t).$$

The right-hand side simplifies to:

$$r_t + n_t + h_t. \tag{D7}$$

Thus:

$$r_{t+1} + n_{t+1} + h_{t+1} = 1 \quad \text{whenever} \quad r_t + n_t + h_t = 1. \tag{D8}$$

Therefore, the total mass of agents is conserved.

D.6 Long-Run Convergence

Given:

$$\Delta r_t < 0, \qquad \Delta h_t > 0,$$

and the hump-shaped path of n_t, the system converges to:

$$r_t \to 0, \qquad n_t \to 0, \qquad h_t \to 1. \tag{D9}$$

These dynamics imply that narrative amplification is temporary, causing medium-run drift but ensuring that PEAD eventually fades as high-precision types dominate.

Appendix E Price Dynamics, Drift, and Empirical Link

This appendix connects the equilibrium pricing equation:

$$P_t = P_t^{\text{fund}} + \phi_t \xi_t$$

to one-period drift, cumulative abnormal return dynamics, and the empirical design used in the paper.

E.1 One-Period Price Change

From equilibrium:

$$P_t = P_t^{\text{fund}} + \phi_t \xi_t, \qquad P_{t+1} = P_{t+1}^{\text{fund}} + \phi_{t+1} \xi_{t+1}.$$

Define the price change:

$$\Delta P_{t+1} \equiv P_{t+1} - P_t.$$

Substitute the decomposition:

$$\Delta P_{t+1} = (P_{t+1}^{\text{fund}} - P_t^{\text{fund}}) + (\phi_{t+1}\xi_{t+1} - \phi_t \xi_t). \tag{E1}$$

Add and subtract $\phi_{t+1}\xi_t$:

$$\Delta P_{t+1} = \Delta P_{t+1}^{\text{fund}} + (\phi_{t+1} - \phi_t)\xi_t + \phi_{t+1}(\xi_{t+1} - \xi_t). \tag{E2}$$

This expresses price changes as a sum of rational updating, a drift term caused by changes in the sentiment multiplier, and sentiment innovations.

E.2 Conditional Expected Drift

Take expectation conditional on information at time t. Since ξ_t is known:

$$\mathbb{E}[(\phi_{t+1} - \phi_t)\xi_t \mid I_t] = \xi_t \, \mathbb{E}[\phi_{t+1} - \phi_t \mid I_t]. \tag{E3}$$

Sentiment innovations are mean-zero:

$$\mathbb{E}[\xi_{t+1} - \xi_t \mid I_t] = 0,$$

implying:

$$\mathbb{E}[\phi_{t+1}(\xi_{t+1} - \xi_t) \mid I_t] = 0. \tag{E4}$$

Thus expected drift is:

$$\boxed{\mathbb{E}[\Delta P_{t+1} \mid I_t] = \Delta P_{t+1}^{\text{fund}} + (\phi_{t+1} - \phi_t)\xi_t}. \tag{E5}$$

This is the drift equation used in the main text.

E.3 Sign of the Drift

The sentiment multiplier is:

$$\phi_t = \frac{\gamma\beta n_t}{\Omega_t}, \qquad \Omega_t = \frac{r_t}{v_t^R} + \frac{n_t}{v_t^N} + \frac{h_t}{v_t^H}. \tag{E6}$$

Because $\gamma\beta > 0$, the sentiment multiplier is increasing in the narrative-trader share n_t, holding the aggregate precision weight Ω_t fixed. More generally, the sign of $\phi_{t+1} - \phi_t$ depends on both the change in n_t and the change in Ω_t. Under the maintained condition that changes in Ω_t do not dominate changes in n_t, an increase in the narrative-trader share raises the sentiment multiplier:

$$n_{t+1} > n_t \quad \Rightarrow \quad \phi_{t+1} > \phi_t. \tag{E7}$$

Combined with the population dynamics in Appendix D, this implies that sentiment-driven drift is strongest while the narrative share is rising and fades as the narrative share declines. Therefore, the model generates temporary drift rather than permanent mispricing.

E.4 Cumulative Abnormal Return Dynamics

Define cumulative price change over the interval $[0, \tau]$:

$$CAR[0,\tau] = \sum_{t=0}^{\tau-1} \Delta P_{t+1}. \tag{E8}$$

Substitute (E2):

$$CAR[0,\tau] = \sum_{t=0}^{\tau-1} \Delta P_{t+1}^{\text{fund}} + \sum_{t=0}^{\tau-1} (\phi_{t+1} - \phi_t)\xi_t + \sum_{t=0}^{\tau-1} \phi_{t+1}(\xi_{t+1} - \xi_t). \tag{E9}$$

Using the fact that sentiment innovations satisfy $\mathbb{E}[\xi_{t+1} - \xi_t] = 0$:

$$\mathbb{E}\left[\sum_{t=0}^{\tau-1} \phi_{t+1}(\xi_{t+1} - \xi_t)\right] = 0. \tag{E10}$$

Thus expected cumulative abnormal return is:

$$\mathbb{E}[CAR[0,\tau]] = \sum_{t=0}^{\tau-1} \mathbb{E}[\Delta P_{t+1}^{\text{fund}}] + \sum_{t=0}^{\tau-1} (\phi_{t+1} - \phi_t)\xi_t. \tag{E11}$$

Because fundamental updates telescope:

$$\sum_{t=0}^{\tau-1} \Delta P_{t+1}^{\text{fund}} = P_\tau^{\text{fund}} - P_0^{\text{fund}}, \tag{E12}$$

the CAR decomposition becomes:

$$\boxed{\mathbb{E}[CAR[0,\tau]] = (P_\tau^{\text{fund}} - P_0^{\text{fund}}) + \sum_{t=0}^{\tau-1} (\phi_{t+1} - \phi_t)\xi_t}. \tag{E13}$$

E.5 Link to the Empirical Analysis

The theoretical model predicts a front-loaded but incomplete pattern of price adjustment. Earnings announcements generate an immediate rational response through the fundamental component of price, P_t^{fund}. However, when sentiment-sensitive demand is amplified by a rising narrative-trader share, the sentiment wedge $\phi_t \xi_t$ can generate temporary post-announcement drift.

In the empirical section, I do not directly observe ξ_t or n_t. Therefore, the main empirical tests focus on the observable implication of the model: if price adjustment is incomplete, earnings-surprise ranks should continue to predict abnormal returns after the announcement window. The baseline empirical specification is:

$$BHAR[a,b]_{i,j} = \alpha_i + \alpha_q + \beta SurpriseRank_{i,j} + \varepsilon_{i,j}, \tag{E14}$$

where $BHAR[a,b]_{i,j}$ is the compounded market-adjusted abnormal return for firm i around earnings announcement j, $SurpriseRank_{i,j}$ is the earnings-surprise decile rank, α_i denotes firm fixed effects, and α_q denotes earnings-quarter fixed effects.

The model implies the following empirical interpretation:

- A large positive coefficient for $BHAR[0,1]$ reflects immediate price formation around the earnings announcement.
- A positive and statistically significant coefficient for $BHAR[2,15]$ or $BHAR[2,60]$ indicates incomplete post-announcement price adjustment.
- If the post-announcement coefficients are positive but smaller than the announcement-window coefficient, the evidence supports front-loaded but incomplete adjustment.

Appendix F Supplemental Data and Robustness Tables

This appendix reports supplemental diagnostics and robustness tests supporting the main data and empirical evidence. The main paper focuses on the core event-time PEAD results, the price-formation regressions, and the tradability/factor-alpha interpretation. This appendix provides additional detail on portfolio composition, pre-announcement controls, price-formation robustness, market-cap and regime heterogeneity, and factor loadings.

F.1 SUE Decile Portfolio Composition and Post-Announcement Drift

Table 7: SUE Decile Portfolio Composition and Mean Post-Announcement Drift

SUE Decile	Mean SUE	Avg. Market-Cap Bucket	Mean BHAR[2,60]	Observations
1	-0.3450	0.305	-1.282%	6,014
2	-0.0065	0.840	-2.522%	6,013
3	-0.0014	1.578	-2.301%	6,014
4	0.0000	2.175	-0.866%	6,013
5	0.0005	2.586	-1.302%	6,014
6	0.0012	2.192	-1.046%	6,013
7	0.0023	1.822	-0.622%	6,013
8	0.0043	1.437	0.152%	6,014
9	0.0094	0.989	0.299%	6,013
10	0.2011	0.371	2.867%	6,014

Table 7 reports the composition and post-announcement abnormal return of each SUE decile. The Mean SUE column confirms that the ranking procedure separates bad-news and good-news earnings announcements. The return column shows that the strongest positive abnormal return occurs in Decile 10, while most lower and middle deciles have negative post-announcement abnormal returns. The average market-cap bucket also shows that the extreme SUE portfolios are tilted toward smaller firms, which is why the main analysis treats size and tradability carefully.

F.2 Pre-Announcement Return Controls

Table 8: Controlling for Short Pre-Announcement Returns

Outcome	Sample	Variable	Coef.	T-stat	P-value
BHAR[2,60]	Full sample	SurpriseRank	0.392	3.76	< 0.001
BHAR[2,60]	Full sample	BHAR[-10,-1]	-0.093	-1.81	0.071
BHAR[2,60]	All-but-microcap	SurpriseRank	0.401	2.83	0.005
BHAR[2,60]	All-but-microcap	BHAR[-10,-1]	-0.031	-0.56	0.576
BHAR[2,60]	Microcap proxy	SurpriseRank	0.302	2.28	0.023
BHAR[2,60]	Microcap proxy	BHAR[-10,-1]	-0.092	-1.90	0.057

Table 8 reports regressions of BHAR[2,60] on earnings-surprise rank and short pre-announcement abnormal returns. The coefficient on SurpriseRank remains positive and statistically significant after controlling for BHAR[-10,-1], suggesting that the post-announcement drift is not explained solely by short pre-announcement price movement.

F.3 Unbiasedness Regressions

Table 9: Unbiasedness Regressions: Announcement Reaction and Full-Window Price Formation

Sample	Beta on BHAR[0,1]	T-stat	P-value	N	R-squared
Full sample	0.949	25.38	< 0.001	54,382	0.0949
All-but-microcap	0.983	46.99	< 0.001	33,965	0.1482
Microcap proxy	0.927	15.02	< 0.001	20,417	0.0738

Table 9 reports regressions of BHAR[0,60] on BHAR[0,1]. The coefficient close to one indicates that the announcement-window return is highly informative about the full 60-day price-formation window. This supports the paper's main interpretation that price formation is front-loaded, while the remaining post-announcement predictability indicates incomplete adjustment.

F.4 D10–D1 PEAD Effect by Market-Cap Group and Macro Regime

Table 10: D10–D1 PEAD Effect by Market-Cap Group

Market-Cap Group	Coef.	T-stat	P-value	N
Nano	1.762	0.92	0.360	22,792
Micro	5.117	2.92	0.003	12,261
Small	1.919	1.61	0.108	9,241
Mid	1.597	1.69	0.091	8,134
Large	2.248	3.00	0.003	7,697

Table 10 reports the Decile 10 minus Decile 1 PEAD effect within market-cap groups. The results show that the D10–D1 spread is not confined to one size group, although statistical strength varies across buckets.

Table 11: D10–D1 PEAD Effect by Macro Regime

Regime	Coef.	T-stat	P-value	N
PreCOVID_2019	0.145	0.08	0.938	8,278
COVID_2020_2021	2.760	1.36	0.172	21,215
RateHike_2022_2024	3.301	2.77	0.006	30,632

Table 11 reports the Decile 10 minus Decile 1 PEAD effect across macro regimes. The D10–D1 spread is strongest and statistically significant in the 2022–2024 rate-hike period, suggesting that the strength of PEAD may vary across macroeconomic environments.

F.5 Daily Factor Loadings

Table 12: Daily Factor Loadings for the PEAD Long-Short Portfolio

Model	Variable	Coef.	T-stat	P-value
CAPM	Intercept	0.032	1.09	0.275
CAPM	mkt_rf	-0.004	-0.16	0.870
FF3	Intercept	0.030	1.10	0.269
FF3	mkt_rf	0.014	0.96	0.337
FF3	smb	-0.109	-3.40	0.001
FF3	hml	0.095	1.85	0.065
FF5	Intercept	0.026	0.96	0.336
FF5	mkt_rf	0.028	1.88	0.060
FF5	smb	-0.048	-1.34	0.179
FF5	hml	0.013	0.24	0.808
FF5	rmw	0.132	2.13	0.033
FF5	cma	0.126	1.76	0.078

Table 12 reports daily factor loadings for the calendar-time PEAD long-short portfolio. The intercepts are positive but not statistically significant. The factor exposures show that the daily PEAD portfolio has some exposure to standard risk factors, which motivates the factor-adjusted alpha interpretation in the main text.

F.6 Monthly Factor Loadings

Table 13: Monthly Factor Loadings for the PEAD Long-Short Portfolio

Model	Variable	Coef.	T-stat	P-value
Alpha only	Intercept	0.662	1.10	0.271
CAPM	Intercept	0.534	0.82	0.413
CAPM	mkt_rf	0.116	1.12	0.263
Fama-French 3	Intercept	0.455	0.76	0.450
Fama-French 3	mkt_rf	0.154	1.56	0.119
Fama-French 3	smb	-0.237	-1.24	0.215
Fama-French 3	hml	0.301	3.86	< 0.001
Fama-French 5	Intercept	0.405	0.68	0.499
Fama-French 5	mkt_rf	0.125	1.10	0.270
Fama-French 5	smb	-0.181	-0.76	0.445
Fama-French 5	hml	0.318	2.43	0.015
Fama-French 5	rmw	0.172	0.64	0.520
Fama-French 5	cma	-0.110	-0.58	0.563
FF5 + Momentum	Intercept	0.349	0.58	0.559
FF5 + Momentum	mkt_rf	0.146	1.27	0.202
FF5 + Momentum	smb	-0.109	-0.45	0.651
FF5 + Momentum	hml	0.343	2.62	0.009
FF5 + Momentum	rmw	0.219	0.88	0.379
FF5 + Momentum	cma	-0.166	-0.85	0.393
FF5 + Momentum	mom	0.126	1.09	0.277

Table 13 reports monthly factor loadings for the PEAD long-short portfolio. The monthly portfolio has a positive and statistically significant HML loading, indicating that the PEAD strategy partly behaves like a value-oriented strategy in monthly regressions. However, the intercepts remain statistically insignificant after factor controls, supporting the main paper's conclusion that event-time PEAD is more visible than tradable factor-adjusted alpha.