

Short-Horizon Market Efficiency Following Positive Earnings Surprises: A Post-Earnings-Announcement-Drift Event Study

Dmitri De Freitas

Washington University in St. Louis — B.S. Data Science & Financial Engineering
Working paper · July 2026 · Interactive companion: dmitridefreitas.com/lab/pead

ABSTRACT. Post-earnings-announcement drift (PEAD) — the tendency of prices to continue moving in the direction of an earnings surprise for weeks after the announcement — is among the oldest and most cited anomalies in empirical asset pricing. This study tests whether the short-horizon component of the drift survives in modern U.S. equity data. For a cross-section of 110 stocks spanning 11 sectors, I compute market-model-adjusted abnormal returns around positive earnings surprises constructed from Compustat fundamentals, CRSP prices, and I/B/E/S consensus estimates, and test three-day post-announcement alpha per name with both parametric (t) and non-parametric (Wilcoxon signed-rank) tests. Only 11 of 110 stocks — 10% — exhibit statistically significant three-day abnormal returns at the 5% level, a rate close to the size of the test itself. The evidence is consistent with near-instantaneous incorporation of earnings news into prices at short horizons, and suggests the classical tradable PEAD effect has substantially attenuated in liquid large-cap names.

1 Introduction

Ball and Brown (1968) first documented that stock prices continue to drift in the direction of earnings news well after announcement day; Bernard and Thomas (1989, 1990) established the effect's magnitude, persistence, and its concentration around subsequent announcements, interpreting it as underreaction to the autocorrelation structure of earnings. PEAD earned a rare distinction: Fama (1998) called it the anomaly "above suspicion." Yet anomalies decay. Rising institutional ownership, algorithmic news parsing, and cheap execution all push toward faster price discovery, and several studies document attenuation of PEAD in large, liquid stocks (e.g., Chordia et al. 2009 on the role of liquidity and arbitrage costs).

This study asks a deliberately narrow question: *after a positive earnings surprise, is there exploitable short-horizon drift left?* The three-day post-announcement window isolates the immediate underreaction channel — the component a systematic trader could realistically harvest with modest holding periods — and tests it name-by-name rather than in pooled portfolios, where a few extreme observations can manufacture apparent drift.

2 Data

The cross-section comprises 110 U.S. equities, 10 from each of 11 sectors, filtered for continuous price history and analyst coverage over the sample window. Three institutional sources are merged: **Compustat** (reported fundamentals and announcement dates), **CRSP** (daily returns and market

capitalization), and **I/B/E/S** (consensus analyst estimates). A positive earnings surprise is defined as reported EPS exceeding the latest pre-announcement I/B/E/S consensus. Standardizing by the dispersion of estimates (SUE-style scaling) leaves the qualitative conclusions unchanged.

3 Methodology

For each stock i and announcement event, abnormal returns are computed under the market model, with parameters estimated over a pre-event window ending before the announcement:

$$AR_{i,t} = r_{i,t} - (\hat{\alpha}_i + \hat{\beta}_i r_{m,t}) \quad (1)$$

$$CAR_i(t_1, t_2) = \sum_{t=t_1}^{t_2} AR_{i,t} \quad (2)$$

The object of interest is the three-day post-announcement window $CAR(+1, +3)$ — drift, not the announcement-day jump itself. For each name, the null of zero mean abnormal return is tested two ways: a standard t-test, and the Wilcoxon signed-rank test, which is robust to the heavy tails and skew of daily abnormal returns. A stock is classified as exhibiting significant drift only at $p < 0.05$. Event windows out to +60 days are examined in the interactive companion tool for the longer-horizon component.

4 Results

Statistic	Value
Stocks tested (11 sectors $\times$ 10)	110
Significant 3-day alpha ($p < 0.05$, both tests)	11 (10.0%)
Expected false-positive rate at $\alpha = 0.05$	~5% (≈ 5 –6 names)
Sectors with any significant name	scattered; no sector concentration

Table 1: Name-level tests of three-day post-announcement abnormal returns after positive surprises.

A 10% rejection rate against a 5% test size leaves little room for a broad, exploitable effect: even before multiple-testing corrections, the excess over the false-positive baseline is a handful of names, and applying a Benjamini–Hochberg control at the 10% FDR level thins the significant set further. The absence of sector clustering argues against an industry-specific information channel. The point estimate of the average three-day CAR across all 110 names is economically small — well inside plausible transaction costs for the marginal trader.

5 Interpretation

The results do not contradict the historical PEAD literature; they date it. Bernard and Thomas's effect was measured in 1974–1986 data, before Reg FD, decimalization, and machine-readable news. In the modern microstructure, the mechanism they identified — slow diffusion of earnings information — has

been largely arbitrated down to the noise floor at the three-day horizon for liquid names. This is the standard lifecycle of an anomaly (McLean and Pontiff 2016): publication attracts capital, and capital compresses the premium. What remains of PEAD in recent data appears concentrated in small, illiquid, hard-to-arbitrage names — precisely where implementation costs are highest.

6 Limitations

Three caveats bound the conclusion. First, the universe of 110 covered, liquid names is selected for data quality, which biases against finding drift; the result should be read as "no short-horizon PEAD *where one could trade it*," not "no PEAD anywhere." Second, only positive surprises are tested; the negative-surprise leg, historically the stronger side, is left to future work. Third, name-level testing sacrifices the power of pooled portfolio sorts in exchange for robustness to aggregation artifacts — a deliberate trade.

7 Conclusion

For 110 liquid U.S. equities, positive earnings surprises are priced essentially on arrival: only 10% of names show statistically significant three-day drift, barely above the false-positive floor, and with no sector pattern. At short horizons and in tradeable names, the market looks efficient with respect to earnings news — and the burden of proof now rests on any strategy claiming otherwise.

References

- Ball, R., and P. Brown (1968). "An Empirical Evaluation of Accounting Income Numbers." *Journal of Accounting Research* 6(2), 159–178.
- Bernard, V. L., and J. K. Thomas (1989). "Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium?" *Journal of Accounting Research* 27, 1–36.
- Bernard, V. L., and J. K. Thomas (1990). "Evidence that Stock Prices Do Not Fully Reflect the Implications of Current Earnings for Future Earnings." *Journal of Accounting and Economics* 13(4), 305–340.
- Chordia, T., A. Goyal, G. Sadka, R. Sadka, and L. Shivakumar (2009). "Liquidity and the Post-Earnings-Announcement Drift." *Financial Analysts Journal* 65(4), 18–32.
- Fama, E. F. (1998). "Market Efficiency, Long-Term Returns, and Behavioral Finance." *Journal of Financial Economics* 49(3), 283–306.
- McLean, R. D., and J. Pontiff (2016). "Does Academic Research Destroy Stock Return Predictability?" *Journal of Finance* 71(1), 5–32.