ANALYSIS / ANALYSIS

Analysis · DEEP ANALYSIS · AI ASSISTED · 12 min read · Oct 31, 2026

Generative AI announcements delivered a +1.84% average abnormal return, event study finds

A 2026 paper says markets price AI news quickly but not completely. We checked its numbers against the paper and the price data.

Updated 4 Nov 2026

Read online: https://traderange.net/analysis/generative-ai-announcements-event-study-tmnycde2/



When a major generative-AI announcement lands, AI-linked stocks move quickly and, on average, in the right direction. But they don't finish moving on the first day. That is the central claim of *Market Reactions to Generative AI Announcements: An Event Study Analysis of Abnormal Returns, 2022–2025* by Kabir Bhushan, published on June 22, 2026 in the *Journal of Economic Research & Reviews* (vol. 6, issue 1, pp. 01–11). The paper lists the author's affiliation as Marquette High School in Chesterfield, Missouri.

The study follows six stocks (Nvidia, Microsoft, Alphabet, Meta, AMD and Palantir) through 18 generative-AI events between November 2022 and January 2025. Its headline result: on announcement day, the average abnormal return was **+1.84%** ($t = 4.32$, $p < 0.01$; Bhushan 2026, Table 3). An abnormal return is the part of a stock's move that the market's own move doesn't explain. In effect, it is the market's verdict on the news itself.

KEY TAKEAWAYS

- Day 0: the average abnormal return across 18 events and six stocks was +1.84%, significant at p < 0.01 (Table 3).
- Day +1 added +0.71% (p < 0.05), so the reaction was not complete at the first close (Table 3).
- Over the five-day window from day -2 to day +2, the all-event cumulative average abnormal return was +3.24% (Table 4).
- Bad news hit harder: the two negative surprises averaged -4.73% against +3.81% for positive model releases, a ratio of about 1.24 (Table 4).
- Our check found misdated events, inconsistent event counts and tables that don't reconcile. The day-0 average is the paper's most consistently reported figure.

How the study works

Bhushan uses a standard event-study design (Section 3). Expected returns come from the CAPM, estimated over 120 trading days running from day -130 to day -11 before each event, with the S&P 500 as the market and the 3-month Treasury bill as the risk-free rate. Abnormal returns are measured over two windows, [-2, +2] and [-1, +1], and post-event drift is tracked over [+1, +10]. Significance is tested with a Brown-Warner cross-sectional t-test.

The 18 events fall into three groups. Positive model or capability releases (type P) include the ChatGPT launch, GPT-4, Google's Bard and Gemini Ultra, Meta's LLaMA 2 and GPT-4o. Positive infrastructure, partnership and earnings events (type I) include Microsoft's $10B investment in OpenAI and two Nvidia earnings reports. The two negative surprises (type N) are the November 2023 removal of OpenAI chief executive Sam Altman and the January 2025 DeepSeek R1 release.

Fast, positive, and not finished on day one

AVERAGE ABNORMAL RETURN BY EVENT DAY, ALL 18 EVENTS. SOURCE: BHUSHAN (2026), TABLE 3.

EVENT DAY	AVERAGE ABNORMAL RETURN	T-STATISTIC	SIGNIFICANCE
-2	+0.21%	1.04	not significant
-1	+0.38%	1.87	p < 0.10
0	+1.84%	4.32	p < 0.01
+1	+0.71%	2.19	p < 0.05
+2	+0.16%	0.78	not significant

Day 0 dominates, but day +1 is also positive and significant at **+0.71%** (t = 2.19). Bhushan reads this follow-through as underreaction. Investors give new information too little weight relative to what they already believed and update slowly, a pattern behavioural finance calls conservatism bias. In the abstract, the author writes that the findings "sit uncomfortably with the semistrong form of the efficient market hypothesis."

The follow-through is real but short. Section 4.4 reports a cumulative 1.4% over days +1 to +10 (t = 1.83, significant only at the 10% level). Roughly 80% of it arrives by day +5, and there is no significant gain from day +6 onward. The evidence supports a day or two of drift after an announcement, not a long, steady repricing.

There is also a small move before the news: +0.38% on day -1, significant only at the 10% level. The paper raises the possibility that information leaked ahead of some announcements. An average this small, across 18 events, can't show that on its own.

Not all AI news is equal

CUMULATIVE AVERAGE ABNORMAL RETURN (CAAR) BY EVENT TYPE. SOURCE: BHUSHAN (2026), TABLE 4.

EVENT TYPE	EVENTS	CAAR DAYS -1 TO +1	CAAR DAYS -2 TO +2	SIGNIFICANCE
Positive model/capability (P)	11	+2.94%	+3.81%	p < 0.01
Positive infrastructure (I)	5	+1.63%	+2.14%	p < 0.05
Negative surprise (N)	2	-3.91%	-4.73%	p < 0.01
All events	18	+2.11%	+3.24%	p < 0.01

Model and capability releases drove the effect, at +3.81% over five days against +2.14% for infrastructure and partnership news. The two negative surprises averaged -4.73%.

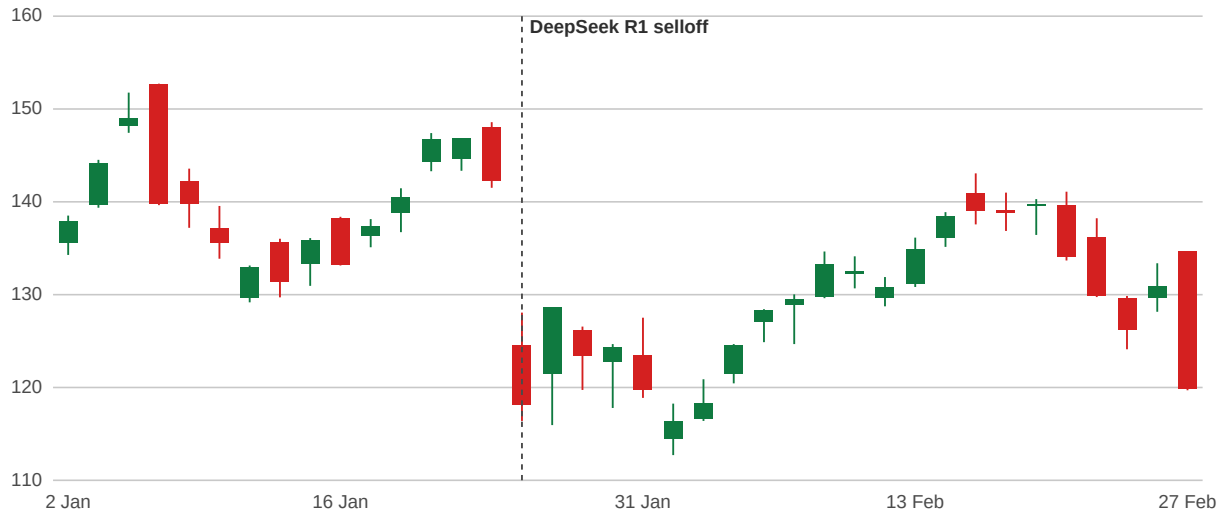
The DeepSeek asymmetry

The sharpest episode in the sample is the DeepSeek R1 release, which the study dates to January 27, 2025. The paper puts Nvidia's one-day fall at roughly 17% and reports five-day cumulative abnormal returns of -9.2% for Nvidia and -6.8% for the six-stock portfolio (Section 4.2).

TradeRange's daily price history agrees on the size of the move. Nvidia closed January 27, 2025 down 16.97% from the previous session, while the S&P 500 fell 1.46%. The simple gap between the two is about 15.5 percentage points. That is a market-adjusted figure, not the paper's CAPM abnormal return, but it gives a fair sense of how much of the drop was specific to Nvidia.

NVDA · Nvidia around the DeepSeek R1 selloff

DAILY CANDLES, JANUARY TO FEBRUARY 2025. MARKET DATA IS DELAYED AND MAY CONTAIN ERRORS. · PRICES AS OF 4 NOV 2026



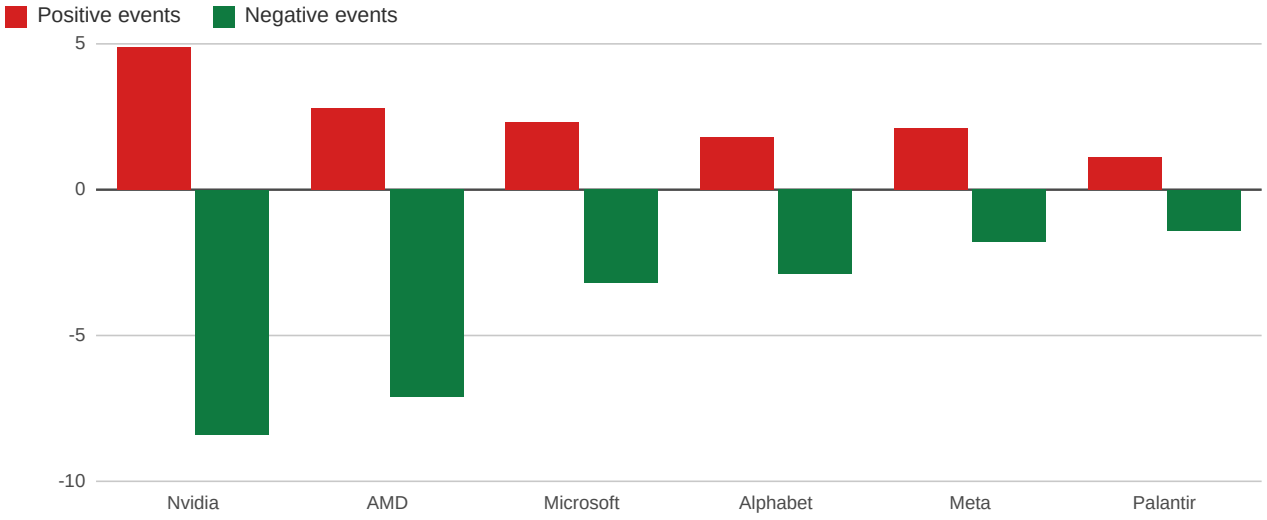
The paper's headline comparison, that negative surprises produce reactions about 24% larger in absolute terms, needs careful reading. The abstract frames it as larger than any positive-surprise event. The figures in Table 4 support a narrower claim: the average five-day CAAR of the two negative events (-4.73%, which also includes the Altman episode) is about 1.24 times the +3.81% average for positive model releases. Bhushan attributes the asymmetry to loss aversion, the tendency for losses to weigh more heavily than gains of the same size.

What it means for Microsoft

The study places Microsoft among firms that are sensitive to generative-AI news, but less so than the chipmakers. Section 4.3 reports an average positive-event CAR of +2.3% for Microsoft and -3.2% for negative events, compared with +4.9% and -8.4% for Nvidia. The paper also notes that Microsoft responds more positively than Alphabet to news about OpenAI, which fits the partnership between the two companies.

Average cumulative abnormal return by firm

POSITIVE-EVENT AND NEGATIVE-EVENT AVERAGES PER FIRM. SOURCE: BHUSHAN (2026), SECTION 4.3. THE PAPER DOES NOT STATE THE EVENT WINDOW; NVIDIA'S FIGURE IS DESCRIBED AS A THREE-DAY AVERAGE.



Two caveats apply to these per-firm numbers. First, the paper doesn't state their event window. Second, they are averages across events, and individual events vary a lot. On January 23, 2023, the day of the $10B OpenAI investment announcement, Microsoft's shares rose 0.98%. On the DeepSeek day they fell 2.14%, against Nvidia's 16.97%. Those are raw price moves from TradeRange data, not abnormal returns.

Reading the paper closely

We checked the paper's event list and arithmetic against public records and against its own tables. Several things don't line up.

- Three event dates differ from the public record. Microsoft introduced Copilot+ PCs on May 20, 2024, not January 19, 2024. Nvidia reported its Q2 fiscal 2024 results on August 23, 2023, not August 22. OpenAI released o1 as a preview on September 12, 2024 and in full on December 5, 2024; the paper uses November 20, 2024. A wrong date means returns are measured on the wrong days.
- The Bard event, which the paper dates to February 8, 2023, is coded as positive. Alphabet's own shares fell 7.68% that day while the S&P 500 fell 1.11%, according to TradeRange data. The paper lists event classification as a judgment call among its limitations.
- The event counts disagree. The text describes 11 positive, 5 infrastructure and 2 negative events; the event table lists 10, 6 and 2.
- The tables don't reconcile. Table 3's daily averages sum to +2.93% over days -1 to +1, but Table 4 reports +2.11% for all events in that window. Weighting Table 4's type averages by event count gives about +2.4% over days -2 to +2, not the reported +3.24%.
- The paper describes a sign test, a 60-day estimation-window check and a market-adjusted robustness test, but reports no results for any of them.
- The introduction lists Claude among the model releases studied, but no Claude release appears in the

event table.

Misdated events affect every window, including day 0, so the exact figures are best read as approximate. Still, the day-0 average is large relative to its standard error and is reported consistently throughout the paper. The multi-day and per-type figures are where the tables disagree.

The paper states that the author has no financial or non-financial interests to disclose and received no funding from the organisations it mentions. It also discloses that large language models were used in drafting and editing. It doesn't include a statement about peer review.

Limits the author acknowledges

The sample is small: 18 events and six firms. The CAPM is a simple model, and three- or five-factor models might fit better. Classifying events as positive or negative involves judgment. The 2022 to 2025 period was unusually enthusiastic about AI, so the results may not carry over to other periods. The study covers US stocks only. And event studies measure short windows around news, not lasting changes in value.

The bottom line

As a description of its sample, the study's core finding is plausible and clearly stated. Across 18 announcements, AI-linked stocks moved meaningfully on the day of the news, kept moving for a day or so, and fell harder on bad surprises than they rose on good ones. The day-0 average is the paper's strongest number. The event-dating problems and inconsistent tables are worth fixing before anyone builds on the rest. Nothing here is investment advice.

Cite as: Bhushan, K. (2026). Market Reactions to Generative AI Announcements: An Event Study Analysis of Abnormal Returns, 2022–2025. *Journal of Economic Research & Reviews*, 6(1), 01–11. Full text (PDF).

SOURCES

- Bhushan (2026), full text (PDF) >
- Bhushan (2026) on ResearchGate >
- Nvidia Form 8-K: Q2 fiscal 2024 results >
- Microsoft: Introducing Copilot+ PCs >
- OpenAI o1 release history >

TradeRange https://traderange.net/analysis/generative-ai-announcements-event-study-tmnycde2/