

Market Reactions to Generative AI Announcements: An Event Study Analysis of Abnormal Returns, 2022–2025

Kabir Bhushan* 

Marquette High School, Chesterfield, Missouri,
United States of America.

*Corresponding Author

Kabir Bhushan, Marquette High School, Chesterfield, Missouri, United States of America.

Submitted: 2026, May 25; Accepted: 2026, Jun 15; Published: 2026, Jun 22

Citation: Bhushan, K. (2026). Market Reactions to Generative AI Announcements: An Event Study Analysis of Abnormal Returns, 2022–2025. *J Eco Res & Rev*, 6(1), 01-11.

Abstract

This paper asks a straightforward question: when a major generative AI announcement drops, do markets actually price it right away, or does something more complicated happen? Using standard event study methodology applied to six AI-linked equities (Nvidia, Microsoft, Alphabet, Meta, AMD, and Palantir) across 18 significant generative AI announcements from November 2022 through January 2025, I find that the answer is mostly the latter. Markets react fast and in the right direction, with announcement-day abnormal returns averaging +1.84% and are highly significant, but they do not get it fully right on day one. A statistically significant positive drift on day +1 (+0.71%, $p < 0.05$) and persistent post-announcement drift following the largest positive-surprise events suggest that generative AI information is incorporated gradually, consistent with conservatism bias among investors. On the negative side, the Deep Seek R1 release in January 2025 triggered a sector-wide selloff with Nvidia dropping roughly 17% in a single day and the six-firm portfolio averaging a -6.8% five-day CAR, a reaction roughly 24% larger in absolute terms than any positive-surprise event, which is exactly what loss aversion theory would predict. These findings sit uncomfortably with the semi-strong form of the efficient market hypothesis and point to real behavioral mechanisms at work in how investors process transformative technology news. Policy implications for AI disclosure standards and practical takeaways for investors in AI-linked equities are discussed.

Keywords: Generative Artificial Intelligence, Event Study, Abnormal Returns, Efficient Market Hypothesis, CAPM, Behavioral Finance, Loss Aversion, Stock Market Efficiency

1. Introduction

The swift commercialization of generative artificial intelligence since late 2022 stands as a landmark event in recent economic history, one of the most powerful technological inflections we have witnessed. Just weeks after Open AI's public introduction of Chat GPT in November 2022, equity markets had begun to re-price firms expected to benefit from the widespread adoption of large language models (LLMs) with a rapidity and magnitude previously unseen outside of pure speculative bubbles. The market capitalization of Nvidia, from an approximate \$300 billion in early 2023, swelled to over \$3 trillion by mid-2024, largely due to investor expectations for an AI-driven increase in demand for

graphics processing units (GPUs). Companies like Microsoft, Alphabet, Meta and a multitude of other AI-linked smaller firms also saw their valuations rise in similar dramatic fashion over the same time period. Such market movements, clustered around discernible public announcements, provide a unique natural experiment for how real-time information related to disruptive general-purpose technologies is incorporated into prices. The question of whether markets incorporate new information related to disruptive technologies efficiently, or systematically under- and over-react to it, is a classic one in financial economics. Under a semi-strong efficient market, per Fama, prices react immediately and in full to any announcement containing public information,

such that no profits can be made trading on the announcement itself; the price should change exactly one time in response to a surprise innovation, such as the release of an AI model considerably exceeding previous performance expectations, and drift neither up nor down following the announcement [1]. Systematic abnormal returns around an event signal that these price adjustments are not instant and full. Such evidence may be indicative of behavioral finance forces driving prices away from rational behavior.

The generative AI bubble serves as a particularly challenging test of market efficiency for three reasons. First, innovation has occurred at an unprecedented pace, making it difficult for investors to rationally predict fundamental value implications from successive product releases. Second, nearly every sector of the economy is affected by generative AI technology, resulting in intricate, second-order effects that even sophisticated analysts cannot easily estimate. Third, because the set of investors interested in AI-related equities includes a considerable contingent of retail investors, and the literature has shown these individuals tend to under-react to complex information about technology, the generative AI announcement cycle provides an especially high-value test of the EMH and behavioral finance concepts [2]. This paper adds to the EMH and behavioral finance literatures. Using the standard event study methodology used to investigate the market's response to corporate announcements, I calculate cumulative abnormal returns (CARs) of an AI-linked equity portfolio around 18 generative AI announcements from November 2022 to January 2025, including model releases (GPT-3.5, GPT-4, Gemini, Claude, LLaM A 2), infrastructure and partnerships (Microsoft-Open AI, AWS Triennium, Apple-Open AI), AI-driven earnings surprises (Nvidia's Q2 2023, Q2 FY2025), and significant competitive shocks (Deep Seek R1 release) [3,4].

The CAPM is employed to derive expected returns over a 120-trading-day estimation window and event windows of [-2, +2] and [-1, +1]. The results yield three major findings. First, CARs are significantly positive over the five-day event window, averaging 3.2%, confirming that generative AI announcements contain information that is important to asset prices. Second, considerable cross-event and cross-firm variation is present; hardware-focused companies are highly sensitive to infrastructure news, platform companies' market responses are somewhat smaller, and negative surprises cause cross-sector negative impacts, with absolute CARs averaging 4.7%. Third, especially among the largest positive-surprise events, post-announcement drift constitutes evidence against the semi-strong EMH, suggesting that the implications of generative AI are incorporated into prices gradually, consistent with conservatism and limited attention among investors, as documented in the behavioral finance literature. The rest of this paper proceeds as follows. Section 2 surveys the relevant literature concerning the EMH, event studies, behavioral finance, and technology announcements. Section 3 establishes the theoretical framework. The methodology is described in Section 4 and the data and events catalog are presented in Section 5. Section 6 discusses the empirical results and Section 7 covers the policy and welfare implications. Section 8 briefly details limitations of the current

work, and Section 9 discusses directions for future research and conclusions.

2. Literature Review

2.1. The Efficient Market Hypothesis

Fama proposed the efficient market hypothesis, which takes three sequentially stronger forms [1]. In its weak form, it hypothesizes that historical prices fully encompass all price information, thus technical trading cannot yield sustained abnormal returns. In the semi-strong form, all public information is completely incorporated into price instantly, including financial statements, macroeconomic data, and corporate releases, meaning fundamental analysis should fail to produce sustained abnormal returns. In its strongest form, all private information is incorporated instantly, indicating even insiders cannot generate abnormal profits. It is the semi-strong form that the event study methodology primarily tests and the focus of this paper. Empirical research has generally supported the semi-strong EMH in large-cap U.S. stocks with no market frictions or behavioral biases over the last several decades. Fama provides an overview of evidence from various asset classes, concluding overall market efficiency while also mentioning the existence of anomalies [5]. A substantial and growing anomalies literature identifies repeated and reproducible violations of semi-strong efficiency. Post-earnings announcement drift (PEAD) refers to a phenomenon where stock prices continue to drift in the direction of an earnings surprise for weeks following the earnings announcement, first observed by Ball and Brown and elaborated upon by Bernard and Thomas, suggesting that earnings information is not fully incorporated immediately [3,6].

Momentum describes a pattern where recent past winning stocks tend to outperform past losers for up to 3 to 12 months, contradicting the EMH that historical returns hold no predictive power. Value and size premia are documented by Fama and French, whose findings suggest either systematic market mispricings or unmodeled risks. There is an active debate between market efficiency and behavior-driven explanations for the anomalies [7]. Fama and French suggest that most anomalies can be explained by their three-factor model, representing a compensation for risk rather than an anomaly [8]. Behavioral economists contend that the scale and consistency of anomalies across numerous markets, asset classes, and time horizons are too great to be attributed solely to risk premia. This paper offers a contribution to this debate by examining market efficiency in a new context: the announcement of generative AI technology, and the associated extreme information and uncertainty environment. Technology-specific markets tend to have much greater and longer-lasting deviations from EMH than either commodities or earnings announcements. The late 1990s tech bubble is described as involving valuations totally inconsistent with any rational model, as pointed out by Shiller, with dot-com bubble stock prices being driven by short-sale constraints rather than rational beliefs, as observed by Ofek and Richardson [9,10]. Pastor and Veronesi build a model where uncertainty about the profitability of new technology induces rational bubbles in adoption periods without any behavioral bias and where high returns volatility and sharp downturns follow

during the process [11]. Their prediction of high volatility, high returns, and sudden price corrections reflects the pattern observed in the AI equity boom outlined in this paper.

2.2. Event Study Methodology and Technology Announcements

The event study methodology, which began with Ball and Brown and was formalized by Fama, Fisher, Jensen, and Roll, has become the primary empirical technique for assessing market reactions to corporate events [3,4]. In essence, by first using pre-event data to establish a baseline estimate of a stock's expected return, researchers can calculate the amount of abnormal return associated with the event itself. Mac Kinlay offers a comprehensive methodological survey, noting the similarity in practical results between the CAPM and the simple market model and that 3- to 11-day event windows are standard [12]. The properties of event study test statistics using daily data are discussed by Brown and Warner, demonstrating that OLS estimations of the market model are resistant to departures from normality assumptions for return distributions [13]. It is argued that the cross-sectional t-test is most appropriate when event-induced variance is expected to increase, as would be the case when a large AI announcement is made given the unpredictability of the event, and that its calculation from cross-sectional variation in CARs rather than time-series variance is better suited. This approach is implemented in the current paper.

A multitude of applications of event study methodology to technology announcements are found in the literature. Chan, Martin, and Keninger find significant positive abnormal returns of approximately 1.4% around announcements of an increase in R&D expenditure, with these effects larger in high-tech industries [14]. Kogan, Papanicolaou, Seru, and Stoffman use the timing of patent grants to infer technological innovation and find that it predicts future stock returns, consistent with the slow incorporation of the value of new intellectual property [15]. Cavusoglu, Mishra, and Raghunathan describe significant negative abnormal returns around announcement of cybersecurity breaches, illustrating that technology-specific negative shocks do cause sharp and quick market corrections [16]. An event study by Gao and Huang examines the effects of AI patent grant announcements, showing approximately 0.8% CAR in a 3-day window; however, this predates the generative AI era by over a decade [17].

The generative AI context is somewhat different from prior technology event studies. First, for many firms, announcements are coming from private entities such as Open AI, Anthropic, or DeepMind rather than from publicly traded equities. This creates spillover effects to publicly traded firms that are most likely to be the largest suppliers, platform owners, or disruptors in the generative AI market. Second, the speed and visibility of AI advancements means that these announcements are likely to be widely reported in financial mainstream media, thus influencing a large and diverse audience of investors rather than a relatively narrow group of specialists, as can be the case for R&D or patent announcements. Third, there is considerable uncertainty surrounding the economic value that improvements in models can create; this leads to non-linearity and extreme difficulty in calculating a correct, rational

value, potentially exacerbating behavioral effects.

2.3. Behavioral Finance and Investor Psychology

Developed by Shiller and De Bondt and Thaler and first formalized in Barberis, Shleifer, and Vishny, behavioral finance provides a theoretical basis for systematic departures from the EMH [18-20]. The intuition is that the cognitive biases widely documented in psychology lead investors to misinterpret information in systematic ways that manifest themselves as predictable patterns in asset prices. The three most important mechanisms for the purposes of this generative AI study are the following.

2.3.1. Conservatism Bias

As defined by Edwards and applied to the financial market by Barberis et al., conservatism bias implies that investors underweight new information relative to their prior beliefs [20,21]. If investors exhibit conservatism bias with regard to AI news, a surprise announcement of a new advanced AI model will lead conservative investors to revise their valuations of AI-related equities less than fundamentals warrant, causing the market to undershoot its fundamental value and adjust in the following days. This implies post-announcement drift (a gradual move toward new fundamentals) in the direction of the surprise. That is, we expect post-announcement drift to be positive following positive-surprise events and negative following negative-surprise events, with the drift becoming weaker as further confirmation about the surprise unfolds.

2.3.2. Representativeness Heuristic

By the representativeness heuristic, investors overreact to salient, recent news and underweight long-run fundamentals, extending short-run trends into the long run [22]. If representative investors believe that the capabilities of AI will continue to grow exponentially from current levels following an announcement of a breakthrough model, then the current pricing of AI-related stocks could reflect expectations of growth rates that are too high, leading to initial overreaction followed by an extended period of correction. Unlike conservatism bias, this mechanism predicts negative drift following positive surprises at the short-run horizon and positive drift at the long-run horizon, and can be present alongside conservatism if different segments of investors are displaying different biases at the same time.

2.3.3. Loss Aversion

The core behavioral principle in prospect theory is loss aversion, whereby investors react more strongly to losses than to equivalent gains [23]. Applied to the AI announcement setting, loss aversion suggests that negative-surprise announcements should induce larger abnormal return reversals than positive surprises. Nvidia's approximately 17% price drop following the Deep Seek R1 announcement in January 2025 may be an example of loss aversion amplified. The negative surprise announced in conjunction with the Deep Seek R1 unveiling, including the competitive threat to Nvidia, the fact that new AI models require less computing power, and the possible repricing of the entire AI hardware investment theme, could all serve to trigger disproportionately large selloffs. Finally,

limited investor attention predicts that information presented in particularly salient or visible formats has a disproportionate impact on asset prices [2]. We can therefore hypothesize that AI announcements accompanied by high-visibility media coverage and large public demonstrations (like the introduction of Chat GPT and GPT-4) should generate larger and faster market reactions than more technical releases announced only through academic papers (like LLaM A 2). We consider this possibility when examining cross-event heterogeneity in Section 6.

2.4. AI and Financial Markets: Emerging Evidence

The research literature examining AI and financial markets is burgeoning, though largely from a micro perspective. Grennan and Michaely use AI-driven trading signals to show improved return predictability for institutional investors, indicating a demand-side influence of AI on market microstructure [24]. Cao, Jiang, Wang, and Yang document evidence that greater AI adoption corresponds to higher Tobin's Q and lower idiosyncratic volatility, implying that AI adoption is a value-relevant characteristic priced by the market. Lopez de Prado shows how machine learning-based trading algorithms have transformed price discovery mechanisms in liquid equity markets [25,26]. By contrast, relatively little attention has been given to the market's response to AI capability announcements themselves, as opposed to how firms adopt AI technologies. This is unfortunate, given that AI laboratory announcements have been among the most substantial single-day price changes recorded in the U.S. equity market during the 2022-2025 period. The most comparable market anomalies in the extant literature concern platform economy announcement effects, where cross-firm spillover effects are often triggered when platform competition intensifies, forcing competitor firms' equity values downward to reflect increased competition [27]. This channel, we will argue below, is at work in many of our events: an announcement by one platform may cause a rival platform's value and hardware providers' value to decline.

2.5. Theoretical Framework

2.5.1. The Null Hypothesis of Market Efficiency

Under semi-strong form EMH, a significant announcement on generative AI would yield an immediate, full price correction for all affected securities. The magnitude of that correction should represent a rational re-assessment of future cash flows, appropriately risk-adjusted. More formally, when an announcement E occurs at time t , the equilibrium change in price is given by: $\Delta P_t = E[PV(\text{future cash flows}) | E] - E[PV(\text{future cash flows}) | \text{no } E]$ This correction should occur within the trading day the information is released. If systematic positive or negative post-day- t returns exist, this violates the semi-strong form of efficiency. Throughout the paper, we test the null hypothesis that expected CAR equals zero. We accept that positive CAR suggests underreaction and negative CAR suggests overreaction, though these alternative hypotheses differ substantively and represent different mechanisms.

2.5.2. Capital Asset Pricing Model

Expected returns are estimated by the CAPM, which is considered the basic framework relating returns to systematic risk:

$$E[R_i] = R_f + \beta_i \times (E[R_m] - R_f) \quad (1)$$

where R_i = return on security i , R_f = risk-free rate (proxied by the 3-month U.S. T-bill rate), R_m = market return (proxied by the S&P 500), and β_i = systematic risk. Beta is obtained by performing an OLS regression of excess daily returns on excess market returns during the 120-trading-day estimation window prior to the event [28,29]. The abnormal return on security i for day t is:

$$AR_{it} = R_{it} - [R_{ft} + \beta_i \times (R_{mt} - R_{ft})] \quad (2)$$

A cumulative abnormal return on security i over $[t_1, t_2]$ is obtained by summing the abnormal returns over that period:

$$CAR_i(t_1, t_2) = \sum AR_{it}, \text{ for } t = t_1 \text{ to } t_2 \quad (3)$$

The cross-sectional average CAR across the N events, known as the CAAR, is our measure of the market's reaction. The standardized CAAR follows a t -distribution under the null of zero abnormal returns, enabling conventional hypothesis testing procedures. The CAPM is chosen over more complicated models such as the Fama-French three- or five-factor model due to ease of use and conceptual transparency. As Mac Kinlay shows, empirical evidence has demonstrated little difference in event study results between the CAPM and market model tests over one- to five-day windows, and the marginal impact of extra factors in multi-factor tests is minimal [12]. A robustness test using the market-adjusted model, forcing $\beta = 1$, is considered in Section 8.

2.5.3. Behavioral Finance Mechanisms

Three distinct, testable hypotheses generated by the behavioral model differ from the EMH null. Each relies on a different cognitive bias for which there is an empirical basis from both the psychology and finance literatures.

- **Hypothesis 1 (Conservatism Bias and Underreaction)**

If investors underweight the information in AI announcements relative to prior beliefs, prices will adjust partially at the time of the announcement and subsequently "drift" in the direction of the surprise over ensuing days. The hypothesis predicts positive CARs on day 0 that are less than the price change if all public information were processed immediately. It also predicts positive drift following such events for days +1 through +5 or more. Barberis et al. present a formal model of this mechanism and derive conditions under which conservatism produces momentum over short time horizons [21].

- **Hypothesis 2 (Representativeness and Overreaction)**

If investors overreact to salient and recent events representing clear advances in AI capability and bid up AI-linked equities based on this expectation at the time of the announcement, the initial CARs will be too large and prices will thereafter "reverse," drifting downward as reality and expectation re converge. Both of the first two hypotheses can simultaneously hold if events trigger different behavioral responses depending on the novelty or salience of the announced advancement.

- **Hypothesis 3 (Loss Aversion and Asymmetric Response)**

Prospect theory implies that individuals respond differently to gains and losses [23]. The implications for event studies are that announcement-induced events associated with negative surprises should result in absolutely larger CARs than positive surprises of comparable informational content. This asymmetry is directly testable in the data: we examine whether CARs corresponding to Type P events are absolutely smaller than CARs corresponding to Type N events.

2.5.4. Cross-Firm Spillover Mechanisms

A novel characteristic of this event study is that announcements generally occur in the private sector (Open AI, DeepMind, Anthropic) and have indirect impacts on publicly traded companies whose primary business models are directly dependent on the progress of AI. We can decompose the impact channel of an Open AI capability announcement on Nvidia stock into three factors.

- **The Direct Demand Channel:** Higher-performance AI models require more processing for training and inference, raising demand for AI chips. The relative impact of such an announcement depends on whether it implies higher or lower demand for chip processing. For example, the Deep Seek R1 model's capability gains suggest greater efficiency and potentially lower future chip demand.
- **The Platform Revenue Channel:** Announcements from Open AI imply changes in the value and utility of related products, for example, the demand for Microsoft's Azure AI services or the value of its Copilot software is influenced by the capabilities of Open AI models available on Azure.
- **The Competitive Repositioning Channel:** The announcement of improved capability from one AI vendor implies immediate re-evaluation by investors of the prospects of rivals based on shifting market positions and expected revenues of competing platform companies. Each of these three channels affects certain firms disproportionately and accounts for variations in observed CARs across firms.

3. Methodology

3.1. Event Identification and Classification

Events are defined as publicly announced generative AI developments that meet three conditions. First, the announcement must have been public on a specific, verifiable date and accessible to all market participants simultaneously. This specifically excludes information that was phased out gradually or was leaked selectively to institutional investors earlier. Second, the announcement must have been plausibly unexpected in terms of its timing, content, or magnitude; only unanticipated information can earn abnormal returns on an efficient market. Third, a contemporaneously published announcement in the financial media must have treated the event as being of commercial importance to one or more sample firms, providing the justification for cross-firm spillover effects that are the aim of the study. Events are divided into three types: Type P (positive model/capability announcement), Type I (positive infrastructure/partnership announcement), and Type N (negative-surprise competitive announcement). The classification

is based on the predicted cash flow impact on the AI-linked equity at the time the announcement was made, inferred from simultaneous analysts' reports and financial media commentary. If an announcement impacts different sample firms differently, the announcement is classified according to its predicted net effect on the majority of the sample firms. The event date is taken as the first trading day on which the announcement was publicly available to market participants. In cases where announcements were made after market close (this was true of several Open AI and Google releases), the event date is taken as the next trading day. In total, 18 events were identified during November 2022-January 2025, 11 of which were designated Type P, 5 as Type I, and 2 as Type N.

3.2. Sample Firms

The sample consists of six publicly traded firms demonstrably and materially exposed to the generative AI value chain. Nvidia Corporation (NVDA) was selected as the dominant hardware provider to the generative AI value chain in terms of both training and inference hardware, with GPU revenues forming the lion's share of the data center division's income. Microsoft Corporation (MSFT) was selected as a direct and commercially relevant platform for the use of Open AI's models (via Azure AI services), including their integration into existing products, with GitHub Copilot being a direct demonstration. Alphabet Inc. (GOOGL) is the other large vertical player, being both a builder of proprietary AI models (Gemini) and proprietary hardware (Google Cloud TPUs) and deploying these at scale. Meta Platforms Inc. (META) was selected as another key AI model developer and public provider of these models (e.g., LLaM A series), where its primary application is via its advertising business, which utilizes AI for optimization. Advanced Micro Devices Inc. (AMD) was selected as the primary alternative provider of AI chips, increasingly adopted by cloud companies as a lower-cost competitor to Nvidia's products. Palantir Technologies Inc. (PLTR) was selected to represent the enterprise application layer of this market. The six sample firms fall into three layers of the AI value chain: hardware (Nvidia, AMD), platform/model (Microsoft, Alphabet, Meta), and enterprise application (Palantir), allowing for different responses across these strata. All firms maintained a market capitalization in excess of \$50 billion throughout the entire sample period, ensuring sufficient daily liquidity and efficient price discovery.

3.3. Estimation Window and Beta Estimation

For each firm-event pair, beta was estimated by OLS regression of the stock's daily excess returns on the S&P 500 index's daily excess returns over a 120-trading-day estimation window preceding the event date, from day -130 to day -11 relative to the event. A 120-day estimation window provides sufficient data to produce a relatively precise estimation of beta, yet is short enough not to be unduly influenced by significant shifts in systematic risk far in advance of the announcement. The 10-day gap between the estimation and event windows prevents information leakages in the pre-event period from directly impacting estimates. As a sensitivity analysis, results were also produced based on a 60-day window; results appear in Section 8.

3.4. Event Windows

Two event windows are chosen to measure the speed at which the market processes information. The smaller window [-1, +1] covers the three trading days including and centered around the event date; it measures the announcement-day reaction plus same-day anticipation effects on day -1. The larger window [-2, +2] covers the five trading days surrounding the event; it measures any slower transmission of information or anticipation by two days prior. The long-window measures post-event drift using the [+1, +10] window following the event, excluding the announcement day; significant positive drift following a positive-surprise event is our primary measure of underreaction through the conservatism bias mechanism.

3.5. Statistical Tests

The statistical significance of event window measures is evaluated using the cross-sectional t-test of Brown and Warner [13]. This method measures average abnormal returns by standardizing with the cross-sectional standard deviation of abnormal returns across all firms rather than the standard deviation from the estimation window:

$$t = CAAR / (s(CAR) / \sqrt{N}) \quad (4)$$

where CAAR is the cross-sectional average CAR, $s(CAR)$ is the cross-sectional standard deviation of individual event CARs, and N is the number of events. The test statistic is robust to event-induced variance increases and cross-sectional dependence of abnormal returns. A finding of statistical significance at the 1, 5, or 10 percent levels implies a rejection of the null hypothesis using a t-statistic distributed with $N-1$ degrees of freedom. Subgroup

analyses are carried out in a similar manner over the relevant subsets of observations. The sign test is also reported as a non-parametric robustness check.

3.6. Use of AI Assistance

Consistent with Digital Finance's author policies, the author discloses that large language models were used in the drafting and editing of some sections of this manuscript. All decisions regarding statistical tests, event identification and selection criteria, data sources, and interpretations were the sole responsibility of the author.

3.7. Data and Event Catalog

3.7.1. Stock Price and Market Data

The daily adjusted closing prices for each of the six sample firms, along with the S&P 500 index, were collected from Yahoo Finance for the period January 1, 2022 to March 31, 2025. The adjusted closing prices incorporate dividend payouts and stock splits so that a consistent measure of return can be computed regardless of corporate activity extraneous to the event data. The continuously compounded daily return for stock i , $r(i,t)$, is calculated as the natural logarithm of the ratio of consecutive adjusted closing prices: $r(i,t) = \ln(P(i,t) / P(i,t-1))$, where $P(i,t)$ is the adjusted closing price for stock i on day t . The daily risk-free rate is represented by the daily continuously compounded return on the 3-month U.S. Treasury bill, collected from the Federal Reserve Economic Data (FRED) database at the Federal Reserve Bank of St. Louis. The annually compounded yields are transformed into daily continuously compounded rates by dividing by 252. Table 1 contains summary statistics for the daily returns for each of the six firms as well as for the S&P 500 over the full sample.

Firm	Mean Daily Return	Std. Dev.	Min Return	Max Return	Estimated Beta
Nvidia (NVDA)	0.31%	3.42%	-16.8%	+24.4%	1.74
Microsoft (MSFT)	0.08%	1.61%	-7.2%	+8.9%	0.91
Alphabet (GOOGL)	0.06%	1.89%	-9.1%	+10.5%	1.03
Meta (META)	0.18%	2.74%	-26.4%	+23.3%	1.28
AMD	0.19%	3.21%	-17.1%	+18.6%	1.62
Palantir (PLTR)	0.22%	4.18%	-24.7%	+31.2%	1.41
S&P 500	0.04%	1.02%	-4.9%	+5.5%	1.00

Note: Beta estimates are full-sample OLS estimates of each firm's sensitivity to S&P 500 excess returns. Event-specific betas are estimated over 120-day pre-event windows as described in Section 4.3. Returns are expressed as daily percentages.

Table 1: Summary Statistics: Daily Returns, January 2022–March 2025

Several interesting observations emerge from the summary statistics. The mean daily return for Nvidia (0.31%) is by far the highest of all six firms, as is its daily standard deviation (3.42%). This suggests that while Nvidia is the target of significant demand driven by the current AI paradigm, it also exhibits very high volatility around expectations on demand. Palantir also has notably high volatility (4.18%), which may be characteristic of a smaller, emerging company driven by growth expectations for a particular technology. Microsoft shows the least volatile pattern

among its peers in terms of mean daily return and volatility, which is expected of an established technology company with a widely diversified set of revenues and products.

3.7.2. Event Catalog

The full event catalog is presented in Table 2. Events are listed chronologically with their classification, a brief description, and the primary sample firms most likely impacted by business model implications. The event catalog spans the public launch of Chat

GPT in November 2022, which marks the beginning of the current generative AI cycle, to the January 2025 launch of Deep Seek R1, which represents a significant competitive shock to the sample of AI-related stocks.

Event Date	Event Description	Type	Primary Firms Affected
Nov 30, 2022	Chat GPT public launch by Open AI	P	MSFT, NVDA
Jan 23, 2023	Microsoft \$10B Open AI investment announced	I	MSFT, NVDA
Feb 7, 2023	Microsoft Bing AI (GPT-4) integration launch	P	MSFT, GOOGL
Feb 8, 2023	Google Bard announcement (Gemini predecessor)	P	GOOGL, MSFT
Mar 14, 2023	GPT-4 release by Open AI	P	MSFT, NVDA, AMD
Mar 21, 2023	Google Workspace AI features announcement	I	GOOGL, MSFT
Jul 18, 2023	Meta LLaM A 2 open-source release	P	META, NVDA
Aug 22, 2023	Nvidia Q2 FY2024 earnings beat (AI demand)	I	NVDA, AMD
Sep 25, 2023	Microsoft Copilot general availability	P	MSFT, NVDA
Nov 7, 2023	Open AI Dev Day and GPT-4 Turbo launch	P	MSFT, NVDA
Nov 17, 2023	Sam Altman fired and reinstated at Open AI	N	MSFT, NVDA
Dec 6, 2023	Google Gemini Ultra announcement	P	GOOGL, NVDA
Jan 19, 2024	Microsoft Copilot+ PC announcement	I	MSFT, NVDA, AMD
May 13, 2024	GPT-4o multimodal release	P	MSFT, NVDA, PLTR
Jun 10, 2024	Apple-Open AI integration announced at WWDC	I	MSFT, NVDA
Aug 28, 2024	Nvidia Q2 FY2025 earnings (record revenues)	I	NVDA, AMD
Nov 20, 2024	Open AI o1 reasoning model release	P	MSFT, NVDA, PLTR
Jan 27, 2025	Deep Seek R1 open-source release (competitive shock)	N	NVDA, AMD, MSFT, GOOGL

Note: Type P = positive model or capability announcement. Type I = positive infrastructure or partnership announcement. Type N = negative-surprise competitive event. Event dates reflect the first full trading day on which the announcement was publicly available to market participants.

Table 2: Generative AI Event Catalog, November 2022–January 2025

There are two Type N events in the catalog that merit special attention. On November 17, 2023, Sam Altman was fired and reinstated, causing tremendous uncertainty regarding the stability of operations at Open AI, as well as the sustainability of Microsoft's \$10 billion investment. These fears led to a substantial downturn in MSFT stock, along with other AI stocks, though the downturn was to a large extent reversed five days later with Altman's reinstatement. The more impactful event is the January 27, 2025, launch of Deep Seek R1. The launch of an open-source model that reportedly reaches on-par performance at a fraction of the compute costs of existing U.S. frontier models poses a threat not only to the thesis that compute will be a bottleneck for AI

expansion and thereby bolster demand for AI hardware, but also to the overall valuation premium that the sample of AI-related stocks have earned over the past two years.

4. Analysis and Results

4.1. Average Abnormal Returns by Event Day

The average abnormal returns for all 18 events during the [-2, +2] event window, along with t-statistics and significance levels for each day, are presented in Table 3. As expected, the pattern of abnormal returns shows several notable features regarding the market's processing of generative AI announcements.

Event Day	AAR (%)	t-Statistic	Significance	Positive Events (%)
Day -2	+0.21%	1.04		56%
Day -1	+0.38%	1.87	*	61%
Day 0	+1.84%	4.32	***	78%
Day +1	+0.71%	2.19	**	67%
Day +2	+0.16%	0.78		50%

*Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. T-statistics computed using the Brown-Warner cross-sectional method. Positive Events (%) reports the fraction of individual event-firm observations with positive abnormal returns on that day [13].*

Table 3: Average Abnormal Returns by Event Day (All 18 Events)

The average abnormal return for the announcement date (day 0) of +1.84% ($t = 4.32, p < 0.01$) is the largest by far and highly statistically significant, indicating that generative AI announcements bring meaningful and actionable information to the market. The AAR for day -1, at +0.38% ($t = 1.87, p < 0.10$), is marginally significant and suggests a mild form of insider trading or pre-announcement information leak, possibly due to information briefings held under embargo by AI laboratories for journalists and analysts prior to formal public announcements. The average abnormal return on day +1 (+0.71%, $t = 2.19, p < 0.05$) is statistically significant, implying that price adjustment extends to the day after the announcement. This is not consistent with the semi-strong form EMH but can be explained by the mechanism of conservatism bias leading to underreaction. The high percentage of positive AAR on day +1 (67%) supports this, as the drift cannot

be attributed solely to extreme outliers. The average abnormal returns for day -2 and day +2 are both positive but statistically insignificant, indicating that there is not much pre-announcement leakage beyond day -1 and that market prices adjust substantially within a day of the announcement. The overall CAAR over the 5-day window sums to approximately 3.3%, with the majority occurring on day 0 and day +1.

4.2. Cumulative Abnormal Returns by Event Type

Table 4 breaks down cumulative abnormal returns by event type, showing that different types of announcements elicit different market reactions. This table allows a direct test of the loss aversion hypothesis (Hypothesis 3, Section 3.3) that negative-surprise events should lead to absolute CARs larger than those for positive surprises of comparable information content.

Event Type	N	CAAR [-1,+1] (%)	CAAR [-2,+2] (%)	t-Statistic	Significance
Positive Model/Capability (P)	11	+2.94%	+3.81%	3.87	***
Positive Infrastructure (I)	5	+1.63%	+2.14%	2.41	**
Negative Surprise (N)	2	-3.91%	-4.73%	-2.96	***
All Events Combined	18	+2.11%	+3.24%	3.64	***

*Note: *** $p < 0.01$, ** $p < 0.05$. T-statistics computed using the Brown-Warner method. Negative-surprise events use a two-tailed test; significance reflects absolute magnitude [13].*

Table 4: Cumulative Abnormal Returns by Event Type

The highest CARs are obtained from the announcement of new models/capabilities (Type P events), which average +3.81% for the 5-day period and are highly significant ($p < 0.01$). Type I events, which announce infrastructure or partnership deals, generate positive but smaller significant CARs (+2.14%), indicating that such information contributes to market valuation, but to a lesser extent than model/capability upgrades, possibly due to their indirect nature in contributing to overall competitive positioning in AI. The difference between the magnitudes of CARs from Type P and Type I events is about 1.7%, implying that investors value capability improvements significantly more than infrastructure investments. Negative-surprise events lead to an absolute CAR of -4.73% over the 5-day window, considerably greater in magnitude than the +3.81% for the most impactful positive events. This asymmetric reaction, where negative surprises cause 24% larger absolute reactions than positive ones of similar informational value, is directly consistent with loss aversion as predicted by prospect theory [23]. The most extreme event is the Deep Seek R1 announcement, which results in an approximately -9.2% 5-day CAR for Nvidia and a -6.8% 5-day CAR for the six-firm portfolio, demonstrating the impact of a severe competitive threat and subsequent re-evaluation of AI infrastructure theses.

4.3. Cross-Firm Heterogeneity in Announcement Reactions

Perhaps the most insightful dimension of the findings is the high cross-firm heterogeneity in announcement responses among the six firms in the sample. Nvidia earns the highest positive-event average CARs in the sample; averaged over three days across all positive

events, Nvidia earns a 4.9% return. This result makes sense given Nvidia's dominant infrastructure role in the ongoing AI capabilities revolution: AI model training and inference become increasingly compute-intensive with capability improvements, and Nvidia supplies over 80% of all AI accelerators. For the GPT-4 event specifically, Nvidia earns a 7.1% three-day return, and the Nvidia Q2 FY2025 earnings beat earns the highest Nvidia return at 11.3%. AMD shows similar directional responses; however, at +2.8%, its average positive-event CAR is systematically 2.1 percentage points lower than Nvidia's, likely reflecting AMD's competitive yet secondary position in the AI hardware space.

Microsoft and Alphabet receive more modest positive-event CARs at +2.3% and +1.8% respectively, reflecting their more diversified business models where AI revenue is not the primary component. Note that Microsoft tends to react much more positively than Alphabet to news regarding Open AI, which relates to the exclusive partnership, while the opposite is true when rival model releases are announced and appear to pose a threat to Google's core search advertising business. Meta, at an average of +2.1% for positive events, shows heterogeneity as well; the highest reactions to Meta occur when news related to open-source models is announced (LLaMA 2), while Meta appears more negatively or neutrally affected when frontier models are announced exclusively by competition. Palantir receives the most volatile returns; it posts an average positive-event CAR of +1.1% with a standard deviation of 5.2%, reflecting the binary interpretation of Palantir's AI offering. In terms of negative surprises, the order reverses itself. Nvidia

experiences the greatest average negative CAR (-8.4% for the two events) largely on the strength of the Deep Seek R1 event; the demonstration that AI frontier models were trainable at drastically lower computation costs implied lower demand for high-end GPUs. The results for AMD are similar (-7.1%), driven by similar competitive concerns. Platform firms like Microsoft (-3.2%), Alphabet (-2.9%), and Meta (-1.8%) see more muted reactions because a threat from Deep Seek R1 is not necessarily a threat to platform revenues. Palantir shows a relatively small negative CAR (-1.4%), consistent with the intuition that its competitive advantages are rooted more in software and enterprise distribution than in specific hardware configuration.

4.4. Post-Announcement Drift and the Underreaction Test

Post-announcement drift provides the strongest evidence regarding conservatism and the semi-strong EMH. Following the 7 largest positive-surprise events, defined as those with an announcement-day return of over 3%, subsequent returns in the window [+1, +10] total a cumulative 1.4%. This is statistically different from zero at the 10% significance level ($t = 1.83$), suggesting moderate underreaction in line with conservatism. The positive drift does not occur uniformly; roughly 80% of the 1.4% cumulative gain occurs within the first 5 days of post-announcement trading [+1, +5], with no significant return gain occurring from days +6 through +10. This finding indicates a rapid initial adjustment followed by a slower period of revision as new information is digested, consistent with investor limited attention [2]. In stark contrast, negative-surprise events appear to exhibit the opposite pattern: both the Deep Seek R1 and Sam Altman events show overreaction, with stock prices largely recovering roughly 40% of their day-0 announcement loss within the subsequent 5 days of trading [+1, +5]. This result suggests that loss aversion dominates the adjustment to bad news, as compared to conservatism dominating adjustments to good news. The dual behavior that appears to characterize post-announcement returns depending on the direction of the shock would be an important contribution to the literature should it hold for a larger sample of announcement events.

5. Discussion

5.1. Implications for Market Efficiency

Overall, the empirical evidence provides a mixed and rich picture of the efficiency of the market for AI announcements. The substantial and significant abnormal returns on the announcement day (AAR = +1.84%, $t = 4.32$) confirm that AI announcements represent relevant information and that markets efficiently process this information in the immediate term as per basic predictions of the semi-strong form EMH. However, two specific findings contradict predictions of full semi-strong efficiency. First, the positive day +1 abnormal return of +0.71% ($t = 2.19$, $p < 0.05$) suggests that announcement-day CARs do not fully impound all the information in AI announcements; approximately 39% of the overall information in AI announcements appears to be incorporated with a one-day delay. Second, even though the post-announcement drift effect appears minimal in this study, its directional consistency with prior results indicates that the market potentially underreacts to positive-surprise AI events and gradually incorporates more information over subsequent days [6].

5.2. Implications for Behavioral Finance

Behavioral explanations for overreaction and underreaction appear applicable to the findings on AI announcements. According to the conservatism bias argument, positive-surprise announcements are initially underreacted to by investors. The day +1 abnormal returns and post-announcement drift in this paper are in line with this explanation. In addition to conservatism, the larger absolute reaction to negative-surprise announcements compared to positive-surprise events, attributed to loss aversion, also explains part of the results. For AI hardware investors, while they may be fully compensated for positive-surprise events over a longer time horizon, the downside risks from negative-surprise announcements might be significantly higher than expected, as loss aversion could amplify negative impacts. If stock prices overreacted to the Deep Seek R1 announcement in an attempt to avoid further loss, subsequent realization of its limitations could lead to a partially reversed correction.

5.3. Implications for Disclosure Policy

The significant stock price impacts of generative AI announcements highlight concerns over disclosure standards and market fairness. While earnings releases are governed by regulations like Regulation FD to ensure simultaneous dissemination of information to all market participants, announcements of generative AI model capabilities lack similar strictures. Firms and labs are therefore able to strategically decide when and how to disclose these innovations. If these strategic disclosure patterns create a disparity in information access, particularly between insider entities and public market investors, they may warrant regulatory scrutiny. This study's evidence of possible day -1 pre-announcement market movement further indicates a need for regulatory consideration and policy adjustment, especially for information dissemination on AI technologies.

5.4. Portfolio and Investment Implications

Several inferences can be drawn for investors from this study's results. First, the spread in event-induced abnormal returns between the hardware and platform segments suggests that a well-diversified investment strategy can significantly mitigate company-specific risks compared to concentrated investments. Second, although statistically small, the presence of a potential post-announcement drift provides some room for exploiting day-after strategies involving adding exposure on positive surprises and trimming it on negative surprises, though high trading costs can easily outweigh any profits. Finally, the loss-aversion effect suggests a potential overestimation of negative surprises by option-implied volatilities, meaning investors holding AI hardware stocks should ideally hold more out-of-the-money put options to mitigate event-risk-related downsides.

5.5. Limitations

A number of limitations restrict the interpretation of these findings and provide avenues for future refinement of the methodology.

- First, and most importantly, the event sample of 18 observations is small by the standards of the event study literature. Typical event studies involve hundreds or even thousands of events,

lending sufficient statistical power to produce precise average treatment effect estimates and enabling tests of across-group heterogeneity. The current analysis is well-powered to detect large average effects but underpowered to detect less-pronounced patterns such as the post-announcement drift with any degree of certainty. The subgroup analyses by event type are subject to substantial data limitations, with the two-event Type N subsample requiring particularly careful interpretation.

- Second, the use of CAPM expected returns is a simplification, potentially leaving residual factor exposure in the abnormal return estimates. AI-linked equities have historically been characterized by high loadings on growth and momentum factors in addition to market beta; thus, a three-factor or five-factor model could more accurately represent expected returns. Robustness tests based on the market-adjusted return model (in which beta is assumed to equal 1) produce qualitatively similar findings and therefore suggest that the major results are robust to this specific modeling decision; however, a multi-factor approach would still enhance robustness.
- Third, while the identification of events is guided by transparent criteria, it requires a subjective judgment that each announcement was in fact unexpected. A more objective method would utilize option-implied volatilities over the event window to quantify the level of market anticipation associated with the announcement content, more clearly isolating genuine surprises from events with an expected component.
- Fourth, the analysis covers an event period from 2022-2025 characterized by an extraordinary surge in investor interest in AI and inflated valuations of AI-related stocks. The magnitude of CARs here likely reflects this heightened level of investor enthusiasm and may not generalize to other periods characterized by a more typical level of tech-driven investor fascination. The effect of announcements will likely decrease as AI development progresses into less novel stages.
- Fifth, this paper does not include an international focus on AI-linked markets. The Deep Seek R1 event shows that geopolitically important AI news can have large spillover effects even on non-U.S. equities, suggesting that reactions in U.S. equity markets do not fully account for the impact on non-U.S. exchanges such as China's and other EU economies. Firms such as TSMC, a critical chip provider to Nvidia and AMD, also do not fully represent the breadth of international AI market components.

6. Conclusion

This paper analyzes 18 significant AI capability and hardware announcements from November 2022 to January 2025, employing event study methodology on the basis of the efficient market hypothesis, the CAPM, and behavioral finance theory. The study reveals three principal findings, each offering implications for financial theorists and practitioners.

- First, both AI capability and infrastructure announcements generate economically and statistically significant abnormal returns. On average, investors gained 3.2% from these announcements over a 5-day event window and 1.84% on the announcement day alone. These results indicate that AI-

specific and hardware announcement events are indeed major price-moving events, analogous to a large positive or negative earnings surprise announcement for a U.S. large-cap stock. The positive drift observed one day prior to announcement (+0.38%) calls into question the even distribution of information among market participants.

- Second, cross-event and cross-firm returns vary significantly. Announcements of increased AI capability or performance have been rewarded with about 1.7 percentage points more return than those associated with AI infrastructure; AI hardware companies generally receive better returns than AI platform companies, and negative-surprise competitive events generate a CAR magnitude 24% greater than that achieved by the top positive event. The Deep Seek R1 event, by far the single most impactful announcement in this sample, tests the boundaries of the behavioral finance hypothesis of loss aversion as well as whether market assumptions about hardware supply and demand may be unrealistic over the longer term.
- Third, significant positive post-announcement drift following large positive-surprise events and a partial price reversal after the Deep Seek R1 selloff both depart from what the semi-strong EMH would predict. The data are consistent with conservatism bias driving underreaction on the upside and loss aversion driving overreaction on the downside. These are not just academic curiosities: they show up in some of the most watched stocks in the market, in real time, during one of the most significant technological shifts in recent memory.

What this paper contributes, ultimately, is a first systematic look at how markets process transformative AI news as it actually happens. Prior event studies have looked at R&D announcements, patent grants, and earnings surprises, but none have examined the generative AI announcement cycle directly, where private-sector breakthroughs by companies like Open AI and Deep Seek ripple immediately into the valuations of publicly traded firms. The findings suggest that even in a market full of sophisticated participants, behavioral biases shape how that information gets absorbed. As the AI development cycle matures and the pace of announcements continues, understanding these dynamics will only become more important for investors managing exposure to AI-linked equities, for regulators thinking about disclosure standards, and for researchers trying to understand how markets price genuinely transformative technology [30].

Declarations

Competing Interests

The author has no relevant financial or non-financial interests to disclose. The author holds no financial stake in any of the analyzed firms, nor has any funding been received from organizations mentioned herein. The author has no competing interests to declare that are relevant to the content of this article.

Author Contributions

Kabir Bhushan is the sole author. All research design, data collection, analysis, and writing were performed by the author.

JEL Classification

G14 (Information and Market Efficiency; Event Studies); G41 (Role and Effects of Psychological, Emotional, Social, and Cognitive Factors on Decision Making in Financial Markets); O33 (Technological Change: Choices and Consequences; Diffusion Processes)

Data Availability Statement

Stock price and index data are openly accessible on Yahoo Finance (<https://finance.yahoo.com>). Risk-free rate data are obtainable from the Federal Reserve Economic Data (FRED) database hosted by the Federal Reserve Bank of St. Louis (<https://fred.stlouisfed.org>). Event dates and their corresponding descriptions are derived from public announcements detailed in major financial news sources, such as Reuters, Bloomberg, and The Wall Street Journal. This study did not involve the use of any proprietary or non-public data; all data employed are readily available to the public and reproducible.

References

1. Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The journal of Finance*, 25(2), 383-417.
2. Hirshleifer, D., & Teoh, S. H. (2003). Limited attention, information disclosure, and financial reporting. *Journal of accounting and economics*, 36(1-3), 337-386.
3. Ball, R., & Brown, P. (1968). An empirical evaluation of accounting income numbers. *Journal of Accounting Research*, 6(2), 159-178.
4. Fama, E. F., Fisher, L., Jensen, M. C., & Roll, R. (1969). The adjustment of stock prices to new information. *International economic review*, 10(1), 1-21.
5. Fama, E. F. (1991). Efficient capital markets: II. *The journal of finance*, 46(5), 1575-1617.
6. Bernard, V. L., & Thomas, J. K. (1989). Post-earnings-announcement drift: delayed price response or risk premium?. *Journal of Accounting research*, 27, 1-36.
7. Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. *the Journal of Finance*, 47(2), 427-465.
8. Fama, E. F., & French, K. R. (1996). Multifactor explanations of asset pricing anomalies. *The journal of finance*, 51(1), 55-84.
9. Shiller, R. J. (2000). *Irrational Exuberance*. Princeton University Press.
10. Ofek, E., & Richardson, M. (2003). Dotcom mania: The rise and fall of internet stock prices. *the Journal of Finance*, 58(3), 1113-1137.
11. Pástor, L., & Veronesi, P. (2009). Technological revolutions and stock prices. *American Economic Review*, 99(4), 1451-1483.
12. MacKinlay, A. C. (1997). Event studies in economics and finance. *Journal of economic literature*, 35(1), 13-39.
13. Brown, S. J., & Warner, J. B. (1985). Using daily stock returns: The case of event studies. *Journal of financial economics*, 14(1), 3-31.
14. Chan, S. H., Martin, J. D., & Kensinger, J. W. (1990). Corporate research and development expenditures and share value. *Journal of Financial Economics*, 26(2), 255-276.
15. Kogan, L., Papanikolaou, D., Seru, A., & Stoffman, N. (2017). Technological innovation, resource allocation, and growth. *The quarterly journal of economics*, 132(2), 665-712.
16. Cavusoglu, H., Mishra, B., & Raghunathan, S. (2004). The effect of internet security breach announcements on market value: Capital market reactions for breached firms and internet security developers. *International Journal of Electronic Commerce*, 9(1), 70-104.
17. Gao, L., & Huang, J. (2016). The impact of artificial intelligence on firm value (Working Paper). *University of Delaware*.
18. Shiller, R. J., & MUCH, D. S. P. M. T. (1981). Justified by Subsequent Changes in Dividends?. *The American Economic Review*, 71(3), 421-436.
19. De Bondt, W. F., & Thaler, R. (1985). Does the stock market overreact?. *The Journal of finance*, 40(3), 793-805.
20. Barberis, N., Shleifer, A., & Vishny, R. (1998). A model of investor sentiment. *Journal of financial economics*, 49(3), 307-343.
21. Edwards, W. (1968). Conservatism in human information processing. *Formal representation of human judgment*.
22. Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive psychology*, 3(3), 430-454.
23. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263-292.
24. Grennan, J., & Michaely, R. (2020). Artificial intelligence and high-skilled work: Evidence from analysts. *Swiss Finance Institute Research Paper*, (20-84).
25. Cao, S., Jiang, W., Wang, J. L., & Yang, B. (2021). From man vs. machine to man + machine: The art and AI of stock analyses (NBER Working Paper No. 28800). *National Bureau of Economic Research*.
26. Lopez de Prado, M. (2018). Advances in financial machine learning (chapter 1). *Advances in Financial Machine Learning, Wiley, 1st Edition (2018)*.
27. Seamans, R., & Zhu, F. (2014). Responses to entry in multi-sided markets: The impact of Craigslist on local newspapers. *Management Science*, 60(2), 476-493.
28. Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The journal of finance*, 19(3), 425-442.
29. Lintner, J. (1975). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. In *Stochastic optimization models in finance* (pp. 131-155). Academic Press.
30. Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of finance*, 48(1), 65-91.

Copyright: ©2026 Kabir Bhushan. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.