

POLITECNICO DI TORINO

Department of Management and Production Engineering

**Master's Degree
in Engineering and Management**

Master's Thesis

The AI Capital Cycle
A Multi-method Explosive-Root analysis applied to
Equities across the AI Value Chain



Supervisor

Professor Caviggioli Federico

Candidate

Caragnano Erika Francesca

Academic Year 2025-2026

July 2026

Table of Contents

Abstract	4
1. Introduction	6
1.1 The Research Problem: Diagnosing the claims of an AI Bubble.....	6
1.1 Background: The Economic Context of the AI Expansion.....	7
1.2 Research Questions and Hypotheses Overview	9
1.3 Research Methodology Overview and Thesis Roadmap	10
2. Theoretical Foundations and Literature Review	12
2.1 Defining bubbles	12
2.1.1 What is a bubble?	12
2.1.2 Recurring Features of Speculative Bubbles: Historical Synthesis	14
2.1.3 Innovation and General-Purpose Technologies as natural drivers of exuberance	18
2.2 Empirical diagnostics of bubbles	21
2.2.1 Rational Bubble Theory	22
2.2.2 Price-Based Tests: From the Augmented Dickey-Fuller to the GSADF	
Framework	24
2.2.3 Nonlinear/dynamical models (Log-Periodic Power Law Models)	32
2.2.4 Fundamental-deviation and cross-sectional approaches	34
2.3 Leverage and financing structure during bubbles	36
2.3.1 Leverage cycles and the financial-instability hypothesis.....	36
2.3.2 Risk transfer and systemic opacity.....	37
3. The AI Industry: Structure and Financing.....	38
3.1 The AI Value Chain and Market Structure	38
3.1.1 Defining the AI industry	38
3.1.2 The AI stack: layers and economic functions	39
3.1.3 Key actors and market structure.....	43
3.2 Demand, Constraints, and Economic Sustainability	50
3.2.1 Demand, revenue, and capex projections.....	51
3.2.2 Physical and technical constraints.....	53
3.2.3 Economic sustainability of the AI model.....	57
3.3 Vendor Financing in the AI Capital Cycle.....	61
3.3.1 Conceptual definitions and the telecommunications precedent	62
3.3.2 Evidence on circular flows in the AI industry.....	64
3.3.3 Evidence on leverage and risk transfer	69
3.3.4 Synthesis: the structural pattern and the empirical hypothesis	72

4.	Research Design: Hypothesis and Methodology	75
4.1	Research question and hypotheses	75
4.1.1	Research question and sub-questions	75
4.1.2	Hypotheses	76
4.2	Data	78
4.2.1	Construction of the AI-exposed universe: indices, equities, commodities	78
4.2.2	Price data, source, and sample window	82
4.2.3	Preprocessing	83
4.3	Methodology	83
4.3.1	The multi-method strategy	83
4.3.2	Stationarity pre-tests: ADF, Phillips–Perron and KPSS	84
4.3.3	The Supremum ADF test (SADF)	85
4.3.4	The Generalised SADF and BSADF date-stamping	85
4.3.5	Monte Carlo critical values	85
4.3.6	Episode detection and the spurious-bubble problem	86
4.3.7	The volatility-robust SV-ADF test	87
4.3.8	Validation on historical bubbles	88
4.3.9	Computational implementation and reproducibility	88
5.	Empirical results	90
5.1	Preliminary Analysis	90
5.1.1	Stationarity results	90
5.2	Explosivity analysis	91
5.2.1	SADF and GSADF results on indices and benchmarks	91
5.2.2	GSADF results on individual equities	93
5.2.3	SV-ADF results and cross-method comparison	106
5.2.4	Commodity-market results	114
5.2.5	Validation results	121
5.2.6	Robustness checks	126
6.	Discussion: Insights and limits of the test	138
6.1	Interpreting and contextualizing findings	138
6.2	Where the tests fall short	140
6.3	Integrating circular financing and debt into the analysis of results	142
6.4	Implications for bubble diagnostics in systemic innovation sectors	144
7.	Conclusions and recommendations	145
7.1	Principal findings and the verdict they support	145
7.2	Recommendations and directions for future research	146

8.	Appendices	147
8.1	Tables and Figures	147
8.1.1	Index.....	147
8.1.2	Tables	148
8.2	Technical Implementation and Repository Structure.....	155
8.2.1	Overview	155
8.2.1	Design Choices.....	155
8.2.2	Environment and Dependencies.....	155
8.2.3	Repository Structure.....	156
8.2.4	Configuration: config/settings.py	158
8.2.5	Execution Pipeline.....	158
8.2.6	Data Layer: data/	160
8.2.7	Outputs: outputs/	160
8.3	Explosivity tests implemented in Python	161
8.3.1	Phillips Perron test method	161
8.3.2	SADF test method	161
8.3.3	GSADF.....	162
8.3.4	Monte Carlo simulation and caching mechanism	164
8.4	Autoregression	167
8.5	Unit Root testing	168
8.6	Cointegration.....	168
9.	Bibliography.....	171

*“Just because a technology is good for society or revolutionizes the world,
it doesn’t mean that it’s a good business proposition”*

(Phil Clifton, November 2025).

Abstract

This thesis examines whether the equities most exposed to the artificial-intelligence build-out exhibit econometrically identifiable explosive price dynamics that are consistent with speculative bubble formation, and whether those dynamics appear along the AI value chain in a pattern robust to the choice of detection method. A bubble is defined operationally, following the rational-bubble literature, as a non-stationary explosive deviation of prices from a random-walk benchmark, identified from the statistical properties of the price series in a way that is independent of the estimate of fundamental value. The empirical design rests on a constructed universe of thirty-seven publicly listed equities, mapped a priori to a thirteen-segment taxonomy of the AI value chain spanning upstream compute, midstream cloud, and downstream software, together with a panel of AI-relevant and placebo commodities. Following the diagnosis of non-stationary time series obtained by means of a set of preliminary stationarity tests (ADF, PP and KPSS), three complementary explosive-root procedures are applied over 2020-2026: the Supremum ADF (SADF), the Generalized SADF (GSADF) with BSADF date-stamping, and a volatility-robust SV-ADF that remains valid under the persistent, time-varying volatility. Critical values for GSADF are obtained by Monte Carlo simulation, and the battery is validated against the dot-com and Bitcoin episodes before being applied to the AI universe. Most of the AI-exposed equities display explosive dynamics: the GSADF rejects the random-walk null for 70.3% of the panel at the 95% level, and for 91.9% under the weaker criterion of at least one dated episode. The activity is concentrated after the public release of ChatGPT in November 2022, with the rejection rate rising from 20.0% before that event to 56.8% after it, and it propagates along the value chain in a temporal sequence: from chip design in early 2023 through foundry, memory, and cloud to application software from late 2024, which mirrors the physical capital-deployment cycle. The commodity panel yields an informative non-finding: explosive behavior does not propagate into AI-relevant real-resource markets relative to placebos. These results are interpreted with deliberate caution. Rejection of the explosive-root null is necessary but not sufficient for a bubble in the operational, explosive-root sense adopted in this thesis; it is neither necessary nor sufficient for a bubble in the broader fundamental-value sense, nor for an eventual collapse: it cannot by itself separate speculative excess from a rational revision of growth expectations. Read alongside evidence of concentration and self-referential, circular financing within the AI ecosystem, the evidence-based diagnosis is the following: the AI capital cycle meets the conditions for a bubble in the explosive-root sense adopted in this thesis. The illustrated structure of prices and financing mechanisms is not a prediction that a collapse is certain or imminent. The principal limitation of this work is represented by the confinement of price-based tests to publicly traded assets, which excludes, thus providing a natural extension to this analysis, the privately held foundation-model developers most central to the cycle.

1. Introduction

1.1 The Research Problem: Diagnosing the claims of an AI Bubble

The question of an "AI bubble" has become the trending object of an extensive public discourse, conducted largely at the level of opinion: commentary oscillates between conflicting convictions that the expansion is a rational anticipation of an eventual surplus and the apprehension that it is a financial mania destined to correct sharply, causing irreparable damage to the worldwide stability. What this discourse lacks is a shared, reproducible basis for diagnosis. There is, at present, no unanimous and comprehensive analytical framework against which competing claims can be adjudicated, and consequently no settled vocabulary for distinguishing a justified repricing of long-horizon expectations from a self-reinforcing divergence of prices from any external basis of value.

This gap is consequential because the economics of financing the construction of the "next big thing" is a matter of public interest much more than it is a matter of private portfolio positioning. The capital being committed, and the financing structures through which it is raised, increasingly intersect with the balance sheets of publicly traded firms and, through them, with the savings of agents who never made a deliberate wager on the technology. In the absence of a protective, decision-useful diagnostic, market participants and observers are left to navigate a period of genuine uncertainty by means of narrative and sentiment alone. The need, then, is for a framework capable of supporting conscious investment decisions under uncertainty, one that replaces the binary alarm of "we are in a bubble, and its collapse will be catastrophic" with a structured assessment of whether, where, and when explosive dynamics are present.

This thesis does not adopt an ethical stance on the desirability of AI expansion, nor does it pronounce on whether the capital deployed is well or badly spent. Its purpose is diagnostic: to ask whether the observable behavior of AI-exposed asset prices exhibits the statistical signature of speculative dynamics, and to relate any such signature to the sector's structural and financing characteristics. The research problem is therefore the selection and application of a diagnostic framework in a domain where the prevailing debate has so far supplied conviction in place of evidence.

1.1 Background: The Economic Context of the AI Expansion

The deployment of capital into artificial intelligence infrastructure has become one of the largest and most concentrated investment cycles in the contemporary technology sector, although its macroeconomic scale remains modest by comparison with earlier investment booms: fluctuating around 1% of US GDP, it is still “*half as large as the rise in IT investment during the dot-com boom of the 1990s*” (Aldasoro, Doerr and Rees, 2026). Aggregate capital expenditure announced by the five largest United States cloud providers (Alphabet, Amazon, Meta, Microsoft, and Oracle) has risen from approximately \$160B in 2022 to nearly \$450B in 2025¹; they already have committed for 2026 lying in the USD 660–725B range². Goldman Sachs projects cumulative hyperscaler capital expenditure of USD 1.15T over the 2025–2027 window, more than double the USD 477B deployed across the preceding three-year period. The annualized compound rate of growth in hyperscaler capital expenditure since the public release of the principal large language models in late 2022 stands in the order of 70%³: several analysts have characterised the current AI infrastructure cycle as one of the largest technology investment buildouts on record.

Concentration is also outstanding in the equity markets. According to Bank of America Global Research⁴, the "Big 10" technology equities currently account for roughly 48% of total U.S. market capitalization. Historical evidence suggests this threshold frequently acts as a terminal marker for asset bubbles; comparable concentration peaks materialized during the 1970s "Nifty-Fifty" era (40%), the late-1980s Japanese asset anomaly (44%), and the 1999 Technology, Media, and Telecommunications (TMT) boom (41%). While proponents argue this centralization is justified by unprecedented free cash flow generation, thereby differentiating it from the zero-revenue euphoria of the dot-com epoch, the erosion of market breadth leaves

¹ [Visual Capitalist](#), *Big Tech AI Spending Over Time* (using Epoch AI and SEC filings), 2026.

² [Futurum Group](#), *AI Capex 2026: The \$690B Infrastructure Sprint*, 2026; MUFG Americas, *Hyperscalers' Capex Above \$600Bn in 2026*, 2025.

³ [Epoch AI](#), *Hyperscaler Capex Trend* (data insight), 2026.

⁴ [BofA Global Investment Strategy](#), GFD Finaeon, Bloomberg, *The Flow Show*, 22 May 2026 (Japan measured as % of MSCI ACWI, all others as % of US stock market).

macroeconomic stability highly vulnerable to idiosyncratic shocks within a singular technological vector.




Figure 1 - Peak U.S. equity market concentration across speculative episodes. Leading group at its peak, as a share of total U.S. market capitalisation.

The economic significance of these magnitudes is twofold. First, the capital intensity now characteristic of the leading operators has risen to levels more typical of utility or heavy-industrial businesses than of software-led firms, and capital expenditure has in several cases begun to exceed internally generated cash flow, so that debt issuance has become a significant component of the financing mix. The financing of AI activity has thereby begun to migrate from the late-stage equity risk borne by private venture portfolios to the balance sheets of publicly traded firms that account for a substantial share of the major United States equity indices, an exposure that begins as sectoral but carries the potential to extend to the broader economy.

Second, the concentration of demand into a small number of input markets (semiconductors, memory, data-center capacity, power, and grid interconnection) has produced measurable structural adjustments in adjacent industries, broadening the cycle's systemic footprint well beyond the platform firms themselves. The detailed evidence on this value chain is presented in Chapter 3. For present purposes, the simultaneity of rising capital intensity at the platform layer and broadened exposure across input markets defines the economic context in which any assessment of valuation must now be situated.

1.2 Research Questions and Hypotheses Overview

The thesis is organized around one main research question:

To what extent do AI-related equities exhibit econometrically identifiable explosive price dynamics consistent with speculative bubble formation, and do those dynamics propagate systematically along the AI value chain in a manner robust to alternative detection methods?

Consistent with the diagnostic stance set out above, a bubble is treated here as an explosive, non-stationary departure of prices from a random-walk benchmark, identified from the statistical properties of the price series themselves agnostically from the estimate of the underlying fundamental value.

This question is decomposed into six sub-questions, each mapped to a component of the empirical design:

- **SQ1:** which AI-exposed equities exhibit statistically significant explosive episodes in the period following the public release of ChatGPT in November 2022?
- **SQ2:** whether explosive episodes appear in broad AI indices or only in individual securities;
- **SQ3:** how their incidence, timing, and duration vary across the segments of the AI value chain;
- **SQ4:** whether the alternative detection procedures agree on the existence and dating of episodes;
- **SQ5:** whether explosive behavior also reaches the commodity markets that supply AI infrastructure;
- **SQ6:** and whether the same test battery, applied to historical data, correctly recovers documented past bubbles.

The analysis is framed by six falsifiable hypotheses, each grounded in the literature of Chapter 2 and the industry account of Chapter 3:

- **H1:** that a subset of AI-exposed equities turned explosive after the public release of ChatGPT;
- **H2:** that explosive episodes concentrate in the supply-constrained upstream segments;
- **H3:** that broad indices show weaker explosivity than their individual constituents, through diversification dilution;
- **H4:** that the methods differ in selectivity, with the volatility-robust procedure identifying fewer episodes than its homoskedastic counterpart;
- **H5:** that the battery validates against known historical bubbles;
- **H6:** exploratorily, that speculative dynamics propagate into AI-relevant commodity markets more than into unrelated ones.

The operational tests, decision rules, and formal statement of each hypothesis are developed in Chapter 4.

1.3 Research Methodology Overview and Thesis Roadmap

The methodological architecture of this thesis relies on both qualitative and quantitative analysis. The qualitative component reconstructs the structural anatomy of the AI sector: its layered organization, its principal actors, the bilateral and multilateral commercial arrangements that link them, and the financing channels through which capital circulates. This component is derived principally from the synthesis of publicly disclosed financial statements, regulatory filings, announced commercial agreements, and a small set of secondary industry datasets, organized to make the circularities and dependencies of interest visible. The objective is to build an industry mapping precise enough to inform the empirical interpretation that follows.

The quantitative component applies the three stationarity tests ADF, Phillips-Perron and KPSS, the Supremum Augmented Dickey–Fuller test (SADF) of Phillips, Wu, and Yu (2011) and the Generalised Supremum ADF test (GSADF) of Phillips, Shi, and Yu (2015), and finally the Stochastic-Volatility Augmented Dickey Fuller (SV-ADF) of Sarkar and Wells (2026) to a set of AI-exposed equity series. These procedures extend the conventional unit-root testing framework to permit the detection of mildly explosive autoregressive behaviour through a recursive estimation window. The GSADF in particular admits the identification of multiple explosive episodes within a single series and provides a formal date-stamping strategy (Backward SADF) whose origination and termination dates can be compared directly to the structural timeline established in the qualitative component. The tests are applied at three levels of aggregation — broad equity indices, AI-exposed thematic baskets, and individual firms occupying distinct positions in the value chain — in order to assess whether the explosivity, where present, is a property of the sector as a whole, of subgroups within it, or of single firms; sensitivity of the principal results to window length and parameter specification construction is examined explicitly.

The integration of the qualitative and quantitative findings is performed in the closing analytical chapter. The objective is not to confirm or to refute the presence of a bubble through any single test statistic but to assess the consistency between the structural account and the empirical results. A formally identified explosive episode whose dating aligns with the intensification of circular financing constitutes substantially stronger evidence than either component would generate in isolation; conversely, a disjunction between the two is itself informative, signaling

either that the structural account omits a material driver or that the explosivity, where present, has origins outside the channels emphasized in this thesis.

The remainder of the dissertation is organized as follows. The next chapter reviews the relevant literature, distinguishing the theoretical foundations of bubble identification: rational bubble theory, behavioral and narrative approaches, and the literature on general-purpose technology cycles coming from empirical diagnostic tradition, with particular emphasis on the right-tailed unit-root framework on which the empirical analysis depends. The following Chapter 3 develops the structural account of the AI sector, beginning with its layered architecture and its principal actors, proceeding through its technical and economic constraints, and concluding with the documentary evidence on circular financing and leverage transfer. The subsequent Chapter 4 presents the research design and Chapter 5 the empirical analysis itself: construction of the AI-exposed universe, preprocessing of the price data, implementation of the testing procedures at the three aggregation levels, robustness checks, and complementary limitations analysis. The Chapter 6 interprets these results in light of the structural account, addresses the inferential limits of the tests, and discusses the implications for bubble diagnostics in sectors characterised by systemic innovation. The final Chapter 7 draws conclusions and recommendations.

2. Theoretical Foundations and Literature Review

2.1 Defining bubbles

The following section aims to outline all academically rigorous definitions of financial bubbles, both in the public and private markets. Academic definitions converge on one core idea: a financial bubble is a period in which asset prices significantly exceed their fundamental or intrinsic value, often accompanied by rapid price increases, speculation, and in most cases an eventual collapse. However, there is no single universally accepted definition, as noted repeatedly across the literature.

2.1.1 What is a bubble?

The definition of *financial bubble* is not unanimous in the academic literature. Definitions vary according to the theoretical framework adopted, and this has direct consequences for how bubbles are identified empirically.

Definitions based on price deviation from fundamental value

A financial bubble is most dominantly defined as a *sustained deviation of an asset's market price above its fundamental value*⁵ (Stiglitz, 1990; Shiller, 2000). Fama (1970) argues that such persistent asset mispricings are incompatible with the theory that markets are efficient, as rational arbitrageurs should theoretically be able to eliminate them. However, the literature distinguishes between two distinct drivers of such deviations: *rational* and *irrational bubbles*. Irrational bubbles arise when prices deviate from fundamentals due to behavioral biases and limits to arbitrage: they represent a fundamental failure of market efficiency. In contrast, rational bubbles are those that can exist even within a framework of rational expectations. Blanchard and Watson (1982) provided early mathematical foundations for this, showing that bubbles can be consistent with rational expectations as detailed in §2.2.1. Tirole (1985) later demonstrated that such deviations can in fact arise in overlapping-generations models (OGLs⁶), where an asset can be purchased at a price higher than its fundamental value if investors rationally expect to resell it at a higher price. This tension between the theoretical possibility of rational bubbles and the empirical evidence of market inefficiency remains a

⁵ The fundamental value of an asset is defined as the discounted present value of expected future cash flows.

⁶ Framework in which the economy is populated by successive cohorts of agents who are alive for a limited number of periods. At any point in time, multiple generations coexist: some earning, some spending, some investing. This structure creates a natural mechanism for assets to be passed from one generation to the next, which is what makes self-sustaining price bubbles theoretically possible.

central theme in finance. A significant practical limitation of this "fundamental-value" approach is that fundamental values are not directly observable and can only be estimated through assumptions about discount rates and growth expectations. This creates a "joint hypothesis problem," (Woodman, 2021) making this definition difficult to operationalize for real-time detection.

Definitions based on price dynamics

A second class of definitions characterizes bubbles according to the asset's price behavior over time. In this framework, a bubble is defined as "*a sharp, accelerating appreciation followed by abrupt collapse*" (Kindleberger & Aliber, 2005). This definition is primarily retrospective: a bubble can only be confirmed after the collapse has occurred, making this framework more suited to historical case studies than real-time diagnostics.

Definitions based on behavioral economics

Behavioral approaches define bubbles as sustained departures of market prices from fundamentals that are driven by investor psychology, social contagion, and self-reinforcing narratives. For Shiller (2000, 2015), prices can rise because enthusiasm spreads through markets and reinforces itself, even when fundamental justification is weak. De Long, Shleifer, Summers, and Waldmann (1990) provide the mathematical basis for why such mispricing persists: they argue that "noise trader risk" makes arbitrage costly and limits the ability of rational traders to correct deviations from fundamentals. Shiller's later concept of narrative economics extends this logic by treating investor stories as causal forces in economic fluctuations, a mechanism especially relevant in rapidly evolving technology markets.

Definitions based on econometric theory

A fourth definitional strand, which forms the methodological foundation of this thesis, is grounded in statistical and econometric theory. From this perspective, a bubble is defined as a *non-stationary*⁷, explosive deviation of asset prices from their fundamental benchmark (Phillips, Wu & Yu, 2011). In this view, while a bubble component remains conceptually tied to fundamentals, it can be classified as "explosive" independently of the unobservable

⁷ In statistics and stochastic process theory, a stochastic process is defined as *non-stationary* when its statistical properties such as mean, variance, and autocovariance structure are not constant over time. In contrast to (weakly) stationary processes, non-stationary processes exhibit time-dependent moments or evolving probability distributions, often due to trends, structural breaks, or changing underlying dynamics. (Hamilton, 1994; Brockwell & Davis, 2016; Box et al., 2015)

fundamental value. The econometric approach has the significant practical advantage of bypassing the need to estimate fundamental values directly. Rather than asking whether prices are *above* fundamentals, it asks whether price dynamics exhibit the statistical signature of explosive behavior, a property that can be tested using observable price series alone. This makes it the most operationally tractable definition for empirical research, and it has been applied across a wide range of asset classes, including equities, real estate, commodities, and cryptocurrencies.

2.1.2 Recurring Features of Speculative Bubbles: Historical Synthesis

The speculative excess often accompanying major technological and financial transitions is not a new phenomenon, and the academic literature has identified a relatively stable set of features that recur across episodes separated by centuries. Kindleberger and Aliber (2005), building on the earlier formalization by Minsky (1986), document a recurring phase structure — displacement, boom, euphoria, distress, revulsion — across episodes ranging from the South Sea Bubble of 1720 to the dot-com crash of 2000. As §2.1.3 discusses, Perez (2002) and Pástor and Veronesi (2009) provide theoretical accounts of why this recurrence is expected: speculative overvaluation is endogenous to the process by which economies absorb transformative technologies under uncertainty. The purpose of this section is not to rehearse each episode in full, but to extract the diagnostic features that the literature has identified as common across bubbles, and which will subsequently be used in Chapter 6 to assess the AI cycle.

2.1.2.1 Narrative-driven displacement and the role of new technology

A first recurrent feature is the presence of a credible *displacement*, typically a technological or institutional novelty, that leads investors to revise their expectations about future cash flows. Shiller (2000, 2015) emphasizes that bubbles are sustained by *narratives* that justify departures from prior valuation norms; these narratives are themselves amplified by media and social diffusion (Shiller, 2017). The South Sea episode of 1720 is the canonical early example: Temin and Voth (2004) show that even sophisticated investors such as Hoare's Bank rode the bubble while being aware of its speculative nature, indicating that the displacement narrative did not require collective irrationality to sustain prices. Frehen, Goetzmann, and Rouwenhorst (2013) extend this finding by showing that the 1720 bubbles in

Britain, France, and the Netherlands were concentrated in firms exposed to genuinely new business models (Atlantic trade, insurance, and finance) rather than distributed uniformly across listed equity. The displacement, in other words, was real; the mispricing concerned the magnitude and timing of its payoff, not its existence. As Chapter 3 will argue, AI shows displacement.

2.1.2.2 Genuine innovation combined with capital misallocation

A second feature, particularly emphasized by the literature on technological bubbles, is the coexistence of *real productive outcomes* and *speculative excess*. The railway mania of the 1840s illustrates this pattern with clarity. Campbell (2012) shows that investors were not naive: they priced railway securities consistently with short-term dividend expectations, exhibiting myopic but internally rational behavior. Campbell and Turner (2012) further document that the mania drew in a broad investor base, including first-time participants, and that the subsequent collapse was driven by a combination of capital calls and tightening credit conditions rather than a sudden revision of long-run fundamentals. Odlyzko (2010) emphasizes that contemporaries were aware of the speculative dimension well before the peak, yet continued to participate. This pattern is consistent with the rational-bubble formulations of Blanchard and Watson (1982) and with the limits-to-arbitrage arguments of Brunnermeier and Nagel (2004) on the dot-com episode. Crucially, the railway network generated substantial long-run economic value, consistent with Perez's (2002) installation–deployment sequence. The episode thus combined genuine speculative excess with real productive outcomes in the long run; a combination that complicates any simple diagnosis of irrationality and that reappears, as Chapter 3 will argue, in the current AI cycle.

2.1.2.3 Concentration, leverage, and the architecture of financing

A third feature concerns the *financial architecture* that sustains the boom. White (1990) and Galbraith (1954/2009) document how the 1929 episode was characterized by concentrated index leadership, a small number of firms associated with radio, automobiles, and electrification, and by the proliferation of leveraged investment vehicles, notably investment trusts with cross-shareholdings that magnified both gains and subsequent losses. Rappoport and White (1993, 1994) provide quantitative evidence that brokers' loans and margin requirements rose sharply ahead of the crash, indicating that intermediaries themselves anticipated rising tail risk. Papadovasilaki et al. (2022) document that the 1920s

boom was centered on a small number of transformative firms, generating a boom-and-crash cycle whose sectoral concentration anticipates patterns observed in later technology episodes. The Japanese asset price bubble of 1986–1991, analyzed by Hoshi and Kashyap (2004) and Okina, Shirakawa, and Shiratsuka (2001), adds a complementary feature: monetary accommodation and collateral-based lending mechanisms that allowed asset price increases to feed back into credit expansion. These contributions establish that bubble episodes are typically structured around identifiable points of leverage, concentration, and intermediary behavior, features that the comparative analysis will revisit in the context of hyperscaler capital expenditure, GPU-collateralized financing, and the concentration of US equity returns in a small number of AI-exposed firms.

2.1.2.4 The dot-com episode as the primary reference point

The dot-com episode of the late 1990s serves as the primary historical reference point for current AI valuations, both because of its proximity in time and because of its trivial similarity to the current cycle: a general-purpose technology, a wave of new listings, and an investor base spanning retail and institutional capital. Ofek and Richardson (2003) document the mechanics of the Nasdaq collapse with a finding that complicates purely behavioral interpretations: the collapse is there attributed in part to the expiration of lock-up agreements that had artificially constrained float, rather than to a sudden revision of fundamental expectations. Brunnermeier and Nagel (2004) show that sophisticated investors did not act as a stabilizing force: rather than betting against overvalued technology stocks, hedge funds increased their exposure to them through the run-up and reduced it only as individual stocks peaked. This the opposite of evidence of irrationality: it pays to ride a bubble as long as one expects to exit before the crowd. The behavior fits the model of Abreu and Brunnermeier (2003), in which rational arbitrageurs decline to correct a mispricing because they cannot coordinate the moment of selling; each waits to extract more gains, and the bubble persists precisely because the smart money has no incentive to be the first to leave. The implication for the present cycle is that the visible participation of professional capital in AI-linked equities cannot be read as a validation of those valuations; sophisticated investors may simply be timing an exit they expect to make ahead of everyone else.

Anderson, Brooks, and Katsaris (2010) find that the late-1990s bubble was not contained within technology: well over half of the S&P 500 by market capitalization, and seven of its

ten sectors, showed bubble-like behavior, extending into financials and industrials. The mispricing was therefore a broad market phenomenon rather than a sector-specific fad. Notably, information technology was the only sector in which most firms paid no dividends. Conventional fundamental anchors were weakest, which foreshadows a recurring difficulty: valuation is least disciplined precisely where a new technology's future cash flows are hardest to estimate. Anderson and Brooks (2014) extend this to the individual-stock level, showing that the size of a stock's bubble deviation is a priced risk factor distinct from the usual market, size, and value factors, and that much of the co-movement in returns normally attributed to market risk in fact reflects the co-movement of bubbles rather than of fundamentals. Both findings bear on the AI cycle, where the firms commanding the highest multiples are often those whose eventual earnings are most speculative, and where concentration means individual-stock mispricings can move the index as a whole. Goldfarb and Kirsch (2019) make a related point in their detailed history of the dot-com era: bubbles are best understood not through aggregate price indices alone but through the interaction of narrative, capital availability, and the institutional features of specific market segments.

2.1.2.5 Empirical evidence on the AI cycle and the limits of historical analogy

A small but growing empirical literature has begun to test whether the current AI cycle exhibits these same features as previous bubbles, with mixed results. Potrykus (2024) applies dynamic time warping to AI-sector stock price trajectories and finds that the AI market *as a whole* does not closely match the Nasdaq Composite during the dot-com euphoria phase of October 1998 to March 2000; the strongest correspondence instead appears in the post-COVID period, though several individual AI stocks do track the euphoria-phase pattern with high fidelity. Basele, Phillips, and Shi (2025) apply the real-time PSY bubble-detection methodology of Phillips, Shi, and Yu (2015) to monthly data from January 2017 to January 2025, time-stamping speculative episodes in both the Nasdaq index and the Magnificent Seven. They find evidence of intermittent bubble behaviour in the Nasdaq itself over two windows, approximately March 2017 to December 2021 and December 2022 to January 2025, and across some of the Mag-7 stocks individually, with the episodes persisting in six of the seven (Apple excepted) in the later period and Nvidia and Tesla exhibiting the fastest explosive dynamics. Maghyereh and Awartani (2025) document multiple synchronized AI bubble episodes during the pandemic exuberance of 2020–2022 and again post-2023,

attributing their formation to liquidity conditions, financial market stress, economic uncertainty, and investor sentiment. Taken together, these studies suggest that at least some AI-adjacent equity segments have exhibited bubble-consistent price dynamics in recent years, even as the aggregate picture resists a clean equivalence with the dot-com euphoria. The comparison is further complicated by structural differences. Teti and Maroni (2021) find that current technology valuations, while elevated, are accompanied by stronger accounting fundamentals relative to prices than was the case in 2000, although they also note a weakening relationship between market prices and accounting fundamentals after 2008. Leading AI-exposed firms generate substantial revenues and cash flows, unlike many dot-com era companies that collapsed under negligible earnings. Regarding dividends, today the picture is materially different: among the AI Big 10 leaders seven *do* pay (Apple, Microsoft, Alphabet, Nvidia, Meta, Broadcom and Micron), and several initiated dividends specifically in 2024–2026. So on this specific dimension, today's leaders look *more* mature than the dot-com cohort as in fact they are benefiting from consolidated revenue streams unrelated to AI.

2.1.3 Innovation and General-Purpose Technologies as natural drivers of exuberance

2.1.3.1 General-purpose technologies and the productivity paradox

Bresnahan and Trajtenberg (1995) define general-purpose technologies (GPTs) as innovations characterized by pervasiveness, inherent improvability, and the capacity to generate innovation spillovers across sectors. AI satisfies all three criteria. GPTs do not diffuse instantaneously: adoption requires large complementary investments in organizational capital, software, training, and process redesign, assets that are largely intangible and poorly measured in standard accounting frameworks. Brynjolfsson, Rock, and Syverson (2021) show that this creates a productivity J-curve: measured output initially stagnates or declines as resources are absorbed into intangible investment, before accelerating sharply once adoption or profitability reaches sufficient scale.




Figure 2 – Productivity J-curve by Brynjolfsson, Rock, and Syverson

The implication for bubble diagnostics is direct: the gap between high asset valuations and low measured productivity that typically motivates bubble concerns is, under GPT conditions, consistent with rational expectations about future productivity more than evidence of mispricing. This is sometimes called the "productivity paradox": the observation that major technological investments do not immediately show up in aggregate productivity statistics.

2.1.3.2 The Perez framework: frenzy is a structural phase of technological innovation

Perez (2002) offers the most comprehensive historical account of how technological revolutions interact with financial markets. Each major technological surge (canal and railway construction in the nineteenth century, electrification and steel in the early twentieth, ICT in the 1990s) follows a recurring pattern of phases: an installation phase driven by financial capital, a frenzy period of speculative excess, a turning point triggered by financial crisis, and a deployment phase in which productive capital takes over and broad diffusion occurs. The frenzy phase is not necessarily caused by mere hype: financial capital pours into new technologies faster than they can productively absorb it, inflating paper values that decouple from real output (Schot & Kanger, 2018; Courvisanos, 2009).




Figure 3 - Perez Technology surge cycle (2002). Source : ResearchGate

Kregel (2009) further argues that financial experimentation during paradigm transitions is itself a mechanism of discovery: speculative schemes test the boundaries of what the new technology can support. The turning point in Perez's framework typically coincides with a financial crisis (the collapse of the speculative phase) after which institutional and regulatory realignment enables the technology to diffuse productively. Examples include the railway mania crash of 1847, the crash of 1929, and the dot-com collapse of 2000–2001. On this reading, bubbles during technological revolutions are endogenous to the transition process and expected.

2.1.3.3 Rational overvaluation during technological transitions

In the model of Pástor and Veronesi (2009), the run-up and subsequent reversal of innovative firms' prices is driven not by absent fundamentals but by a shift in the *nature* of the risk attached to a genuinely promising technology: when new technology arrives, its average productivity is unknown, that uncertainty is largely idiosyncratic and investors price the technology optimistically; as adoption becomes likely, the uncertainty turns systematic, the required risk premium rises, and prices fall even though the underlying fundamentals were real throughout. High valuations resting on substantial revenues and cash flows are therefore not, on their own, evidence against a bubble-like trajectory. Their model matches empirical patterns in both the railway boom of the 1830s–1860s and the Nasdaq expansion of the 1990s. Johnson (2007) reaches a similar conclusion through an optimal learning framework: agents

rationally overinvest during early technology diffusion because investment is the primary means of resolving uncertainty about returns to scale. Gârleanu, Panageas, and Yu (2012) extend this to a long-run growth model, showing that large infrequent technological innovations endogenously generate boom-and-crash valuation cycles even among fully rational agents.

2.2 Empirical diagnostics of bubbles

The empirical literature on bubble detection is organized around a fundamental identification problem: asset price fundamentals are not directly observable and must be estimated under assumptions. The fundamental value of an asset is theoretically defined as the present discounted value of all future cash flows it will generate. The formula is:

$$P_t^* = \sum_{k=1}^{\infty} E_t \left[\frac{D_{t+k}}{(1+r)^k} \right]$$

where D_{t+k} are future dividends and r is the discount rate. The problem is that neither of these inputs is directly observable. Future dividends must be forecast under assumptions about growth rates. The discount rate must be estimated under assumptions about risk preferences and expected returns. Different reasonable assumptions about either input produce materially different estimates of P_t^* . This is the joint hypothesis problem (Gürkaynak, 2008): any test of whether prices deviate from fundamentals is jointly a test of the assumed pricing model and of the bubble hypothesis⁸.

Literature has produced two broad organizational frameworks for classifying diagnostic methods, which differ in their purpose and level of granularity. The first, proposed by Gürkaynak (2008) in a foundational survey of econometric bubble tests, distinguishes between *direct* and *indirect* methods. Direct methods compare market prices to independently estimated fundamental values (e.g., through discounted cash flow models or enterprise valuation) and identify a bubble as the residual gap between observed prices and estimated intrinsic value. Indirect methods, by contrast, do not require an estimate of fundamental value; instead, they detect non-fundamental price behavior through its statistical properties, such as

⁸ If the test finds no deviation, we cannot distinguish between two explanations: either there is no bubble, or the fundamental value model is wrong and the true fundamental is actually higher than estimated. If the test finds a deviation, we similarly cannot distinguish between the existence of bubble and the underestimates fundamentals by the model.

explosive dynamics, excess volatility, or structural breaks in the autoregressive process⁹ governing prices. Because estimating fundamental values introduces model dependence that can itself be the source of apparent deviations, indirect methods dominate modern empirical research. Wei, Li, and Wang (2022) confirm this distinction in the context of emissions trading markets, noting that indirect tests have become the operational standard precisely because they circumvent the unobservability of fundamentals.

The second framework, synthesized across the survey literature by Harsha and Ismail (2019), provides a more granular classification of indirect methods into six families: variance-bound tests, unit-root and explosive-root tests, regime-switching and nonlinear tests, log-periodic power law models, duration-based tests, and machine learning classifiers. These families differ in the specific statistical feature of bubble behavior they target, the assumptions they require about the data-generating process, and their replicability across different markets and sample sizes.

2.2.1 Rational Bubble Theory

In equilibrium, the price of an asset must satisfy:

$$P_t = \frac{1}{1+r} \mathbb{E}_t[P_{t+1} + D_{t+1}]$$

Rational bubble theory addresses a fundamental question: can asset prices deviate persistently from intrinsic value without any lapse in investor rationality? Blanchard and Watson (1982) showed that the answer is affirmative. Their framework embeds the bubble directly into the standard asset pricing equation.

Solving this equation forward by repeated substitution yields a general solution consisting of two components: the fundamental value (the expected present discounted value of all future dividends) and a residual term B_t , the bubble component. For B_t to be consistent with rational expectations, it must satisfy:

$$\mathbb{E}_t[B_{t+1}] = (1+r) \cdot B_t$$

This no-arbitrage condition requires the bubble to grow at exactly the required rate of return each period. If it grew more slowly, rational agents would sell the asset; if it grew faster, they

⁹ See Appendix – Autoregression for a clear explanation of the statistical concept.

would buy it until the deviation was eliminated. A bubble component satisfying this condition is therefore *self-sustaining*: rational investors hold the overpriced asset because they rationally expect it to appreciate further, making the elevated price path a self-fulfilling equilibrium.

Blanchard and Watson (1982) showed that such components are not ruled out by the equilibrium condition alone, because the asset pricing equation admits multiple solutions, of which the fundamental-value path is only one.

Rational bubbles can take several forms. A deterministic bubble grows at rate $(1+r)$ — continuously and without interruption. More relevant from an empirical perspective are stochastic bubbles, which survive with probability π each period and collapse with probability $1 - \pi$, yet still satisfy the no-arbitrage condition in expectation. Blanchard and Watson (1982) formalized this structure, showing that a bubble can periodically collapse and reinflate while remaining consistent with rational behavior throughout. The macroeconomic implications of these dynamics extend well beyond individual asset markets: as Martin and Ventura (2018) demonstrate, rational bubbles can be embedded into standard macroeconomic models to account for large swings in investment, capital flows, and output, all phenomena that purely fundamental-based models struggle to explain.

The theory nonetheless rests on assumptions that constrain its empirical reach. It requires that all investors share the same expectations, have access to the same information, and discount future payoffs at a common rate r . These conditions are rarely satisfied in real markets, where short-selling constraints prevent pessimistic investors from pushing prices down to fundamental value, and where liquidity frictions mean that even investors who recognize an overvaluation may be unable to act on it before the bubble expands further. In these settings, price deviations can persist and grow through mechanisms that do not require each individual investor to expect a capital gain at rate r , the strict condition rational bubble theory imposes. The theory therefore accounts for a narrower class of real-world deviations than its formal elegance might suggest.

Further internal limitations arise from the theory's own logic. The no-arbitrage condition $\mathbb{E}_t[B_{t+1}] = (1 + r) \cdot B_t$ implies that if the bubble is zero at any point in time, it must remain zero forever: a bubble component equal to zero, grown at any rate, remains zero. Diba and Grossman (1987) formalized this implication, proving that a positive rational bubble on a freely disposable asset cannot emerge at a finite point in time during the life of an asset: it must have been present from the very first moment of trading. This creates an obvious tension

with observed episodes in which prices appear to detach from fundamentals only after a period of stability, and then only in response to identifiable shocks or shifts in sentiment. If new bubble episodes cannot start, the theory offers no account of why they appear when they do. Despite these constraints, rational bubble theory provides the formal apparatus on which the first generation of econometric detection methods was built. The no-arbitrage growth condition implies that a bubble component introduces explosive dynamics into the price series; dynamics that have no counterpart in dividends or earnings and that can, in principle, be identified through the time-series tests examined in the following sections.

2.2.2 Price-Based Tests: From the Augmented Dickey-Fuller to the GSADF Framework

2.2.2.1 The baseline of bubble testing: unit-root testing and its failure

The earliest econometric attempts to detect bubbles relied on standard unit-root and cointegration¹⁰ tests. The logic behind these tests derives from rational bubble theory: if an asset price contains a bubble component, that component must grow at the rate of return required by investors, otherwise rational agents would arbitrage it away. A bubble component growing at this rate will cause the price series to exhibit explosive behavior, diverging from the stationary or random-walk process that prices follow in the absence of bubbles. Diba and Grossman (1988) operationalized this insight by applying left-tailed unit-root tests to stock prices and dividends, reasoning that if prices and dividends are cointegrated (they share a common stochastic trend) then no explosive bubble component can be present.

This approach, while theoretically sound, lacks empirical evidence as identified by Evans (1991). Periodically collapsing bubbles (bubbles that inflate and then burst before re-inflating) generate price series that, over a full sample, appear stationary or only mildly non-stationary even when genuine explosive episodes are present within the sample. Standard left-tailed unit-root tests therefore have low power against the most empirically plausible class of bubbles: those that expand and collapse repeatedly over time. The implication is that failing to reject a unit root does not constitute evidence against a bubble; it may simply reflect the inability of the test to detect episodic explosiveness that is obscured by subsequent collapses.

¹⁰ For a plain explanation of unit-root and cointegration and its role in bubble detection, see Appendices

2.2.2.2 The right-tailed approach: ADF as a bubble detector

The insight that bubbles should produce right-tailed explosive behavior rather than the left-tailed non-stationarity that standard unit-root tests are designed to detect, caused a reorientation of the testing strategy. Rather than asking whether a price series contains a unit root, the right-tailed approach asks whether the autoregressive root of the price process exceeds unity, which would be consistent with mildly explosive bubble dynamics. The standard vehicle for this test is a modified version of the Augmented Dickey-Fuller (ADF) regression.

The ADF regression, in its general form, models the change in a price series Δp_t as a function of the lagged price level p_{t-1} and lagged differences Δp_{t-j} :

$$\Delta p_t = \alpha + \delta p_{t-1} + \sum_{j=1}^k \phi_j \Delta p_{t-j} + \varepsilon_t$$

where α is a constant, δ is the coefficient of interest, k is the number of lagged differences included to control for serial correlation, and ε_t is a white noise error term. In the standard left-tailed test, the null hypothesis is $\delta = 0$ (a unit root, or random walk) against the alternative $\delta < 0$ (stationarity, or mean reversion). For bubble detection, the hypothesis is reversed: the null remains $\delta = 0$, but the alternative is $\delta > 0$, which implies that the autoregressive root $1 + \delta$ exceeds unity, an explosive process. This right-tailed ADF test produces a test statistic that, under the null, follows a non-standard distribution, and inference requires simulation-based critical values rather than standard normal or chi-squared tables. The right-tailed ADF thus directly tests for the statistical signature of rational bubble behavior: a self-reinforcing price acceleration that, in expectation, continues to grow. However, applied to the full time series, it inherits the same weakness as the Diba-Grossman approach: if the sample contains both bubble and non-bubble periods, the test is applied to a mixture of regimes, and the signal of explosiveness can be diluted by the non-bubble observations.

2.2.2.3 Phillipps Perron and KPSS tests

Two companion tests are conventionally used alongside the ADF to characterize the order of integration of a series before any explosive-root test is applied.

The Phillips–Perron (PP) test addresses the same unit-root null as the ADF but treats the nuisance structure of the errors differently. Real price-series errors are typically serially correlated and heteroskedastic, which distorts the ADF statistic unless corrected. Where the ADF absorbs this dependence parametrically, by adding lagged difference terms to the regression, the PP test runs the simple Dickey–Fuller regression and corrects the statistic afterwards, replacing the ordinary error variance with a long-run variance estimated non-parametrically. This long-run variance known as a Newey–West, or heteroskedasticity- and autocorrelation-consistent (HAC) estimator, sums each error’s own variance together with its covariances with neighboring errors, down-weighting more distant ones, so that the resulting statistic is valid without specifying a lag structure. The practical trade-off is that the ADF requires a choice of lag length whereas the PP requires a choice of bandwidth for the long-run variance; the two are therefore complementary, the ADF being sensitive to lag selection and the PP to bandwidth selection.

The KPSS test of Kwiatkowski et al. (1992) is complementary in a stronger sense: it reverses the hypotheses, taking stationarity as the null and a unit root as the alternative. Because the ADF and PP tests can have low power in finite samples, failing to reject a unit root even when the series is in fact stationary — pairing them with a test whose null is the opposite proposition provides confirmatory triangulation. When the ADF and PP tests fail to reject the unit-root null and the KPSS test simultaneously rejects its stationarity null, the three jointly point to a series that is integrated of order one, which is the maintained assumption of the right-tailed explosive-root tests.

2.2.2.4 The SADF: recursive detection with a fixed origin

Phillips, Wu, and Yu (2011) proposed the first systematic solution to the regime-mixing problem through a recursive testing strategy. Rather than applying a single ADF regression to the full sample, they proposed running a sequence of right-tailed ADF tests on subsamples that begin at a fixed point and expand forward through time, a procedure known as a forward recursive or expanding window approach. For each subsample ending at time t , they compute

the right-tailed ADF statistic, and the supremum of these statistics across all subsamples is taken as the test statistic. This is the Supremum ADF, or SADF, statistic:

$$\text{SADF}(r_0) = \sup_{r_2 \in [r_0, 1]} \text{ADF}_0^{r_2}$$

where r_0 denotes the minimum window size¹¹ as a fraction of the total sample, r_2 denotes the endpoint of each subsample, and $\text{ADF}_0^{r_2}$ is the right-tailed ADF statistic computed on the subsample from the beginning of the series to observation $r_2 \cdot T$. The null hypothesis is that the full series follows a random walk; the alternative is that at least one explosive subperiod exists. Rejection of the null hypothesis allows the researcher to date the origination and collapse of a bubble episode by identifying the subsamples in which the ADF statistic exceeds its critical value.

The SADF represented a major methodological advance: it was the first procedure capable of providing real-time date-stamps of bubble episodes rather than merely testing for their presence over the full sample. Phillips, Wu, and Yu (2011) applied it to the Nasdaq Composite index and identified an explosive episode beginning in mid-1995 and collapsing in early 2000, consistent with the dot-com narrative. The critical values for the SADF statistic are non-standard and are obtained via Monte Carlo simulation, with the distribution depending on the minimum window size r_0 , a parameter that must be chosen by the researcher and that affects the test's sensitivity to short-lived explosive episodes.

However, the SADF retains a structural limitation. Because the starting point of each expanding window is fixed at the beginning of the series, the test is designed to detect a single explosive episode. If the price series contains multiple bubble-and-collapse cycles, as is common in asset markets observed over long periods, the SADF may fail to identify later bubbles, because the observations from a prior collapsed bubble can obscure the signal of renewed explosiveness. Evans' (1991) critique thus applies partially to the SADF as well: it solves the problem for single-bubble series but not for multiple episodes.

¹¹ The minimum window size r_0 controls the shortest subsample on which an ADF regression is estimated. Too small a value allows estimation on very few observations, producing unreliable statistics; too large a value reduces the test's ability to detect short-lived explosive episodes.

2.2.2.5 The GSADF: flexible windows and multiple bubble detection

Phillips, Shi, and Yu (2015) extended the SADF framework to resolve the multiple-bubble problem through a generalization that allows both the start and the end of each test window to vary. Rather than fixing the origin at the beginning of the series, the Generalized SADF — or GSADF — takes the supremum of right-tailed ADF statistics computed over all subsamples defined by a flexible starting point r_1 and endpoint r_2 , subject to a minimum window size constraint:

$$\text{GSADF}(r_0) = \sup_{\substack{r_2 \in [r_0, 1] \\ r_1 \in [0, r_2 - r_0]}} \text{ADF}_{r_1}^{r_2}$$

where $\text{ADF}_{r_1}^{r_2}$ is the right-tailed ADF statistic computed on the subsample from observation $r_1 \cdot T$ to observation $r_2 \cdot T$. By allowing the start of the window to shift forward, the GSADF can isolate explosive behavior within a particular subperiod even when earlier or later observations are non-explosive. The global null hypothesis remains that the full series is a random walk; rejection implies the existence of at least one explosive subperiod anywhere in the sample.

The date-stamping procedure associated with the GSADF, known as the Backward SADF, or BSADF, extends this logic to every observation in the sample. For each observation r_2 , it computes:

$$\text{BSADF}_{r_2}(r_0) = \sup_{r_1 \in [0, r_2 - r_0]} \text{ADF}_{r_1}^{r_2}$$

A bubble episode is identified as originating at the first observation r_2 for which BSADF_{r_2} exceeds the corresponding critical value, and as collapsing at the first subsequent observation for which it falls below the critical value. This produces a sequence of identified bubble periods with explicit start and end dates, which can be plotted against the price series and compared to known market events.

The GSADF is the current methodological standard for empirical bubble detection in the academic literature. It has been applied across equity markets, real estate, currency markets,

commodities, and cryptocurrencies (Phillips & Shi, 2020; Basele, Phillips & Shi, 2025; Maghyreh & Awartani, 2025). Its key advantages over the SADF are its ability to detect multiple non-overlapping bubble episodes and its greater robustness to the presence of structural breaks or regime changes in the non-bubble portion of the series.

Critical values and inference

Because the distributions of both the SADF and GSADF statistics are non-standard under the null hypothesis — they depend on Wiener process functionals that do not reduce to familiar tabulated distributions — inference requires simulation-based critical values. These are generated by Monte Carlo methods: a large number of random walk series of the same length as the observed data are simulated, the SADF or GSADF statistic is computed for each, and the empirical distribution of these statistics across simulations provides the critical values at conventional significance levels (90%, 95%, 99%). The shape of this distribution, and therefore the critical values, depends on the minimum window size parameter r_0 and the sample size T . Researchers must therefore select these parameters before testing, and different choices can affect whether a given episode is classified as explosive or not, a source of specification sensitivity discussed further in Chapter 6.

Limitations

Despite its dominance in empirical literature, the GSADF framework carries several well-documented limitations that bear directly on its application to AI markets.

First, the test detects statistical explosiveness, not fundamental deviation. A price series can be explosive for reasons unrelated to speculative excesses, for instance, because of a genuine upward revision in expected fundamentals, a structural break in the discount rate, or simply because of a short-term liquidity shock. As Gürkaynak (2008) emphasizes, rejecting the null of a unit root in favor of explosiveness is not equivalent to establishing the presence of a bubble in any of the theoretical senses reviewed in §4.1.1; it establishes only that the price process is locally non-random-walk, which is consistent with but not proof of a bubble.

Second, the GSADF has limited power in short samples. When the explosive episode is brief relative to the total sample length, the window over which explosive behavior is detectable may contain too few observations to generate a statistic that reliably exceeds the critical value. This is particularly relevant for AI markets, where the most discussed episodes of apparent speculative excess have been concentrated in relatively short windows — the period

following the public release of ChatGPT in late 2022, for instance, spans only a few quarters of weekly data.

Third, as Homm and Breitung (2012) establish through systematic simulation comparisons, the GSADF's performance depends on the data-generating process assumed under the alternative hypothesis. When bubbles collapse periodically rather than growing monotonically to a single peak, the Chow-type break test they propose can outperform the GSADF in finite samples. This does not invalidate the GSADF but implies that its results should be triangulated against alternative procedures, particularly when the series being tested is likely to contain multiple short-lived episodes rather than a single sustained expansion.

Fourth, the GSADF is designed for publicly traded, continuously observable price series. It requires a sufficiently long time series with regular observations at a consistent frequency. For private market valuations like venture capital rounds, unicorn assessments, pre-IPO marks, neither condition is typically satisfied. This structural limitation, explored further in Chapter 6, is particularly consequential for the AI sector, in which a substantial share of economic activity and speculative investment occurs in markets that generate no continuous price signal.

2.2.2.6 Volatility-robust detection: the SV-ADF test

The recursive ADF tests above are built under homoskedasticity, or at most weak heteroskedasticity, in the innovation process. Financial price series, however, exhibit persistent and time-varying volatility, and a transient variance spike can mimic the right-tailed signature of explosiveness, inflating the ADF statistic in the absence of any genuine bubble. Sarkar and Wells (2026) address this by developing a stochastic-volatility-robust ADF (SV-ADF) test that extends the Phillips, Wu and Yu (2011) recursions to autoregressive models with strongly persistent mean and volatility dynamics, drawing on the moderate-deviation asymptotic theory of Sarkar and Wells (2025), and without imposing a strict parametric form on the variance process.

The test is based on the augmented autoregression

$$x_t = \mu + \delta x_{t-1} + \sum_{j=1}^L \phi_j \Delta x_{t-j} + u_t, \quad u_t = \sigma_t \varepsilon_t, \quad \varepsilon_t \sim \text{i. i. d. } \mathcal{N}(0, 1)$$

in which the innovation u_t is the product of a (possibly highly persistent) stochastic-volatility scale σ_t and a standardized shock ε_t . As in the right-tailed framework above, the unit-root null is tested against a mildly explosive alternative,

$$H_0 : \delta = 1 \quad \text{against} \quad H_1 : \delta > 1$$

On each expanding subsample ending at the fractional point $r \in [r_0, 1]$, with $\tau = \lfloor n r \rfloor$ observations, the coefficient δ is estimated and standardized by the locally estimated innovation variance, yielding the forward-recursive SV-ADF statistic

$$\text{SV-ADF}_r^\delta = \tau(\hat{\delta}(\tau) - 1), \quad \text{SV-ADF}_r^t = \left[\frac{\sum_{j=1}^{\tau} \tilde{X}_{j-1}^2}{\hat{\sigma}_\tau^2} \right]^{1/2} (\hat{\delta}(\tau) - 1)$$

where $\tilde{X}_{j-1}$ are the demeaned (and lag-adjusted) regressors on the subsample and $\hat{\sigma}_\tau^2$ is the corresponding estimate of the average innovation variance scale $m_{n,\tau}$. The variance normalization is the single feature that distinguishes the SV-ADF from the ordinary right-tailed ADF: it rescales the statistic by the volatility actually prevailing inside each window, so that a high-variance regime no longer mechanically enlarges it.

The appropriate point of comparison is therefore the single-window SADF test, not the GSADF. The GSADF is a multiple-bubble procedure: its double-supremum recursion is designed to detect and date-stamp several distinct explosive episodes within one sample. The SV-ADF, by construction, supports a single bubble episode per series — which is precisely why Sarkar and Wells benchmark it against the PWY/SADF framework and explicitly decline to compare it with the PSY/GSADF procedure. The SV-ADF should thus be read as the volatility-robust counterpart of the SADF, inheriting its forward-recursive, single-episode dating logic while replacing its homoskedastic calibration.

That calibration is reconsidered along two distinct margins. First, under the moderate-deviation asymptotics of Sarkar and Wells (2025), the limiting functionals coincide with the homoskedastic right-tailed case after normalization by the average volatility scale, so explosiveness is identified from the same standardized Brownian-functional limit; the finite-sample critical values, obtained by simulation, nevertheless differ from the constant-volatility benchmark. Second, and independently, the test abandons PWY's restrictive collapse-and-reinitialization mechanism, which forces the post-bubble process back to an $O_p(1)$ neighbourhood of the pre-bubble level — a restriction the authors show is implausible under mildly explosive dynamics, since the terminal bubble level grows exponentially in the span of

the episode. Origination and collapse are instead calibrated against separate thresholds, reflecting the distinct distributions governing entry into and exit from the explosive regime.

Date-stamping follows the forward-recursive logic of Phillips, Wu and Yu (2011). The estimated origination date is the first upward crossing of the right-tail origination boundary,

$$\hat{r}_e = \inf_{s \in [r_0, 1]} \{ s : \text{SV-ADF}_s^\delta > cv_{\beta_n, r}^\delta(s) \}$$

while, conditional on $\hat{r}_e$, the collapse date is the first subsequent downward crossing of the left-tail collapse boundary, after a short exclusion band,

$$\hat{r}_f = \inf_{s \in [\hat{r}_e + \log(n)/n, 1]} \{ s : \text{SV-ADF}_s^\delta < cv_{\beta_n, l}^\delta(s) \}$$

Crucially, the two boundaries are not the same. Calibration under the unit-root null is appropriate for origination, because the pre-bubble regime is non-explosive; for collapse the process has already passed through an explosive phase, so reusing the null-based cutoff is not theoretically justified. The thresholds are obtained by simulation under the relevant hypothesis and are well approximated, for the origination (upper) and collapse (lower) margins respectively, by

$$cv_{\beta_n, r}^\delta(s) \approx \frac{\log(n_s)}{10}, \quad cv_{\beta_n, l}^\delta(s) \approx \frac{\log(n_s)}{2}$$

where $n_s = \lfloor n s \rfloor$ is the size of the recursive subsample. Both critical sequences are nuisance-parameter free and diverge slowly, so the procedure remains directly implementable without knowledge of the explosiveness rate α .

The practical consequence for the present application is a reduction in volatility-induced false positives. Applied to AI-exposed equities, the homoskedastic recursions tend to flag explosive episodes during high-variance windows that the SV-ADF screens out, retaining only the persistent component of the price acceleration; this distinction is exploited in the empirical date-stamping of Chapter 5.

2.2.3 Nonlinear/dynamical models (Log-Periodic Power Law Models)

The explosive-root tests examined in §2.2.2 detect bubble episodes by identifying departures from a unit-root process, but they impose no structure on how a bubble grows or terminates. A distinct methodological tradition, rooted in statistical physics rather than econometrics, approaches bubble detection by examining the price dynamics that precede a

crash. The central claim of this literature is that speculative bubbles follow a characteristic trajectory (super-exponential price growth punctuated by accelerating oscillations) that reflects the self-reinforcing herding behavior of investors and, in principle, contains information about the timing of an eventual collapse.

The theoretical foundation of this approach was established by Johansen, Ledoit, and Sornette (2000), who proposed a rational expectations model in which crashes are driven by the progressive imitation among noise traders. As local herding intensifies, the probability of a crash per unit time, known as the hazard rate, increases. For investors to remain in the market despite this rising crash risk, expected returns must also increase, which requires prices to grow faster than exponentially. The interaction between this super-exponential trend and the discrete scale invariance of the underlying herding network produces log-periodic oscillations: cyclical deviations around the accelerating trend whose frequency increases as the market approaches the critical point. The model does not predict that a crash will occur at a specific date, but identifies the most probable crash time, the critical point t_c , as the moment at which the bubble dynamics become unsustainable.

These insights were formalized into the Log-Periodic Power Law Singularity (LPPLS) model, in which the price trajectory in the run-up to a crash takes the form:

$$\ln p(t) = A + B(t_c - t)^m + C(t_c - t)^m \cos(\omega \ln(t_c - t) + \phi)$$

where t_c is the critical time, $m \in (0,1)$ governs the rate of super-exponential acceleration, ω controls the frequency of log-periodic oscillations, and A , B , C , and ϕ are additional parameters. The term $(t_c - t)^m$ produces *faster-than-exponential* growth as $t \rightarrow t_c$, while the cosine term generates the oscillatory signature. Calibration requires nonlinear optimization over this parameter space, which is computationally demanding and sensitive to initial conditions. Geraskin and Fantazzini (2013) provide a systematic review of calibration methodologies and diagnostic tests, including applications to the gold bubble that collapsed in December 2009.

Three limitations of the LPPLS framework are relevant for the present thesis. First, the model requires a continuous and sufficiently long observable price series. Private AI company valuations, observed only at discrete funding rounds with irregular spacing, do not satisfy this requirement. Second, parameter estimation is fragile: small changes in the estimation window

can shift t_c substantially, and the critical time is subject to wide confidence intervals even in well-behaved cases. Third, the model identifies the most probable crash time but assigns a non-zero probability to the bubble reaching t_c without crashing; the price trajectory satisfying the LPPLS form is a necessary but not sufficient condition for an imminent collapse. For these reasons, and because this thesis is concerned with detecting and date-stamping explosive episodes rather than forecasting a specific crash time, the LPPLS framework is reviewed here for completeness but is not applied in the empirical analysis; the price-based explosiveness tests of the preceding sections are retained as the operational toolkit.

Despite these limitations, the LPPLS framework occupies a distinct and complementary role in the diagnostic toolkit. Unlike the GSADF, which tests for explosive behavior without specifying its source, the LPPLS model embeds a behavioral mechanism, herding under rational expectations, and produces a time-stamped signal that can be monitored in real time. Its application to publicly listed AI stocks and indices would be methodologically feasible, but lies outside the scope adopted here; its application to private markets would further require the construction of synthetic price proxies from funding round data, a methodological challenge noted in Chapter 5.

2.2.4 Fundamental-deviation and cross-sectional approaches

It is worth clarifying what practitioners mean by fundamental value in applied settings, since the academic literature inherits and formalizes this intuition. In equity markets, fundamental value is typically approximated through valuation multiples: ratios that relate a company's market price to an observable financial metric. The most widely used are the price-to-earnings ratio (P/E), which compares market capitalization to current or forward earnings; the enterprise value-to-EBITDA ratio (EV/EBITDA), which adjusts for capital structure and is preferred for capital-intensive firms; the price-to-sales ratio (EV/Revenue), which is applied when earnings are negative or uninformative, as is common in early-stage technology companies, and last the cyclically adjusted price-to-earnings ratio (CAPE), developed by Shiller (2000), which smooths earnings over a ten-year window to remove cyclical distortions. An asset is considered overvalued when its multiples deviate significantly from historical norms or from sector peers, and a bubble is informally diagnosed when those deviations are extreme and persistent. Academic methods formalize this intuition by constructing explicit present-value models of what prices should be (the variance-bounds,

cointegration, and option-based approaches reviewed below) and treating the residual between observed prices and model-implied fundamentals as the bubble component.

The methods examined in previous sections detect bubbles through the time-series behavior of prices, without directly measuring fundamental value. A distinct methodological family approaches the problem from the opposite direction: it estimates what the price *should* be and interprets the residual as a bubble component. These approaches are older but remain active in the applied literature, and they raise a diagnostic challenge that purely time-series methods avoid, which is the necessity of correctly specifying fundamentals. The earliest formal test in this tradition is the variance-bounds test of Shiller (1981). The intuition is that if stock prices equal the present discounted value of future dividends, then prices should not be more volatile than the dividends themselves: rational revisions to expected dividends should be the only source of price variation. Shiller showed that actual stock prices over the twentieth century were far more volatile than any rational dividend forecast could justify. This finding is consistent with systematic departures from fundamental value, though not uniquely attributable to a bubble rather than to time-varying discount rates or other misspecifications of the present-value model.

A methodologically distinct strand bypasses the need to specify fundamentals entirely by exploiting cross-sectional price discrepancies. If two assets carry identical cash flows but trade at different prices, the difference cannot reflect fundamentals: it must reflect a non-fundamental component in one or both prices. Pavlidis and Vasilopoulos (2020) apply this logic to Chinese stocks simultaneously listed on mainland and Hong Kong exchanges, where institutional barriers prevent arbitrage from equalizing prices. They show that deviations from the law of one price in these segmented markets exhibit explosive dynamics consistent with bubble behavior and carry predictive content for subsequent price movements, thus providing bubble evidence that requires no assumptions about the level of fundamental value.

A further approach that sidesteps fundamental specification entirely is the option-based framework of Jarrow and Kwok (2021). Under standard no-arbitrage conditions, the fundamental value of a stock can be inferred nonparametrically from cross-sections of European option prices: the option market reveals the market's own state-price distribution, from which the fundamental price can be bound without imposing any dividend or earnings model. The difference between the observed stock price and this option-implied fundamental

value constitutes a nonparametric bubble estimate. Jarrow and Kwok (2021) apply this procedure to individual equities and report statistically significant bubble components in several cases, providing interval estimates that are robust to model misspecification in a way that variance-bounds and cointegration tests are not.

Both the cross-sectional and option-based approaches share a structural advantage over time-series methods: they do not require a long historical price series and are not confounded by the Evans (1991) critique. Their limitation, directly relevant to this thesis, is one of data availability. Cross-listed Chinese equities and exchange-traded options exist for large publicly listed companies; neither instrument is available for private AI startups, where valuations are observed only at discrete funding rounds and no derivatives market exists. The diagnostic toolkit for private markets must therefore rely on adaptations of the time-series methods reviewed in Sections 4.2.1 and 4.2.3, with the approaches examined here serving as validation benchmarks for the publicly listed segment of the AI equity universe.

2.3 Leverage and financing structure during bubbles

2.3.1 Leverage cycles and the financial-instability hypothesis

The diagnostic methods surveyed in §2.2 operate on observable price series and are agnostic about the financing structures that generate or sustain explosive dynamics. A complementary strand of the literature argues that leverage and risk transfer are preconditions that shape both the likelihood and severity of speculative episodes: a dimension that price-based tests cannot capture and that motivates the structural analysis of AI financing conducted in Chapter 3.

The theoretical foundation for this view originates with Minsky (1986), whose financial instability hypothesis posits that periods of stability endogenously encourage leverage accumulation and the progressive erosion of lending standards, driving a transition from hedge to speculative to Ponzi financing structures. While developed before modern risk-transfer instruments existed, this framework has been extended and formalized in the subsequent literature. Brunnermeier and Oehmke (2013) show that leverage cycles interact with asset price dynamics through fire-sale externalities and balance-sheet contagion: when leveraged institutions face margin calls, forced asset sales depress prices further, amplifying

the initial shock into a systemic event. This mechanism is distinct from the behavioral and rational overvaluation dynamics discussed previously; it operates through financing structure instead of expectations or sentiment.

The empirical evidence for leverage as a forward-looking systemic risk indicator is well-established. Adrian and Brunnermeier (2016) demonstrate that leverage, size, and maturity mismatch significantly predict an institution's contribution to system-wide tail risk, measured through their CoVaR framework, and that this predictive relationship holds in the buildup phase rather than only during crises. Brunnermeier, Rother, and Schnabel (2020) extend this logic specifically to bubble episodes, showing across nearly thirty years of bank-level data that systemic risk rises substantially during a bubble's expansion phase and accelerates further during its bust, with the magnitude depending on both institution size and asset class.

2.3.2 Risk transfer and systemic opacity

Risk transfer instruments introduce a further complication. By redistributing tail risk away from originating institutions through credit derivatives, securitization, or structured lending these mechanisms reduce idiosyncratic exposure while potentially concentrating systemic risk in less transparent parts of the financial system. The same fire-sale and balance-sheet logic produces a central paradox at the system level: individually rational hedging at the institution level can produce collectively fragile outcomes, precisely because the distribution of risk becomes opaque and procyclical. Jarrow and Lamichhane (2021) formalize a related mechanism, showing that bubble size increases under conditions of high market liquidity, conditions that risk-transfer instruments can synthetically reproduce by enabling leveraged exposure without commensurate capital requirements. Taken together, this literature establishes that the severity of a bubble episode is not determined by price dynamics alone. Sectors characterized by high leverage, concentrated exposures, and significant risk redistribution may remain vulnerable to abrupt corrections even when price-based diagnostics remain inconclusive, or may amplify corrections that price tests do detect. This is the dimension the structural analysis of Chapter 3 supplies.

3. The AI Industry: Structure and Financing

The purpose of this chapter is to assemble, layer by layer, the elements of the AI industry whose interaction determines the relationship between the deployment of capital and the realization of revenue across the cycle. This chapter precedes the analysis documenting the features and dynamics that could make such an episode possible, amplify it, or render it harder to detect through standard price-based diagnostics. The research strategy of Chapter 4 is designed based on these elements, against which empirical results of Chapter 5 are interpreted.

The chapter is organized into three parts. §3.1 delimits the industry, decomposes it into functionally distinct layers, and identifies the principal actors at each layer. §3.2 considers the demand side of the cycle using forecasts, highlights physical constraints, assesses the economic sustainability of the announced expansion, and describes the operational and concentration risks that arise from the structure documented in §3.1. §3.3 turns to the financial dimension: the conceptual categories of circular financing, the documentary evidence on intra-ecosystem commercial flows, and the leverage and risk-transfer channels through which AI infrastructure is increasingly financed.

The exposition observes a strict source hierarchy. Where possible, claims are grounded in regulatory filings (10-K, 10-Q, 8-K, 6-K, 20-F), official statistical sources (IEA, EIA, BIS, IMF, OECD), and peer-reviewed work; commercial estimators (TrendForce, Synergy Research Group, Counterpoint Research, Mercury Research, Jon Peddie Research) are used with explicit methodological flagging; press reporting is used for context and deal tracking. Quantitative claims are always accompanied by an indication of whether they are directly observable, company-reported, regulator-reported, commercially estimated, survey-based, or forecast.

3.1 The AI Value Chain and Market Structure

3.1.1 Defining the AI industry

The expression “AI industry” denotes a set of interdependent activities rather than a single homogeneous sector. The industry has no single statistical-classification code: it spans elements

of NAICS¹² 334413 (semiconductor manufacturing), 513210 (software publishers), 518210 (data processing, hosting, and related services), and 541512 (computer systems design services), among others, and it does not fit cleanly into any single GICS¹³ sub-industry. Index providers, accordingly, define their AI universes differently: the MSCI ACWI Robotics & AI Index¹⁴ and the S&P Kensho AI Index¹⁵ rely on revenue-attribution methodologies, while the Nasdaq CTA Artificial Intelligence and Robotics Index¹⁶ uses a thematic-classification approach. Where index-level definitions diverge, the constituents differ materially in both number and weighting.

For the analytical purposes of this thesis, the AI industry is defined as the *set of activities that supply, integrate, or directly monetise the compute-intensive production and deployment of large-scale machine-learning systems*. This definition is layered. It includes, on the input side, the semiconductor equipment, chip design, memory, foundry, advanced packaging, and data-centre infrastructure that physically support large-scale training and inference; the cloud infrastructure and specialised compute providers that intermediate the supply of compute to model developers; and the data, foundation-model, and application layers at which AI capabilities are produced and delivered.

The definition excludes, for tractability, three adjacent categories governed by distinct mechanisms: defence-related AI applications, whose financing depends on government procurement rather than private capital markets; robotics and autonomous-vehicle systems, whose physical-product cycles differ structurally from data-centre-mediated AI; and the legal, ethical, and regulatory infrastructure that surrounds AI deployment but does not itself generate cyclical capital expenditure of the magnitude relevant to bubble analysis.

3.1.2 The AI stack: layers and economic functions

The AI industry is decomposed in the following section into three functional groupings that move from upstream physical infrastructure to downstream software, and within which individual layers exhibit distinctive economic properties. Companies have been grouped according to their position in the production process, and boundaries between layers are drawn

¹² NAICS is the U.S. statistical system used to classify business establishments by industry;

¹³ GICS is the MSCI-S&P global equity classification framework used to group companies into sectors, industries, and sub-industries.

¹⁴ [MSCI ACWI IMI Robotics & AI Index](#) methodology document.

¹⁵ [S&P Kensho Artificial Intelligence Index](#) methodology.

¹⁶ [Nasdaq CTA Artificial Intelligence and Robotics Index](#) methodology.

where production economics change. Each firm is assigned based on its primary value-capture function; firms with overlapping functions (for example, Google, which is both a Hyperscaler and a Foundation model developer) are explicitly noted. This taxonomy is an original analytical synthesis, grounded in the five-layered AI-supply-chain framing used in the BIS paper by Gambacorta and Shreeti (2025) and in OECD Competition in artificial intelligence infrastructure (which details the physical inputs of AI). Figure 4 summarizes this decomposition and indicates the direction of the two opposing flows that organize the cycle: the physical product flow from upstream to downstream, and the capital-commitment flow that propagates downstream demand signals back to upstream investment.




Figure 4 - The AI value chain. Source: author’s synthesis. Boxes represent layers; the upper bracketed groupings indicate the three economic functions (upstream, midstream, downstream).

The upstream grouping consists of activities whose output is the physical capacity to perform large-scale machine-learning computation. It comprises semiconductor equipment, chip design, memory (particularly high-bandwidth memory, HBM), advanced-node foundry services, advanced packaging, data-centre construction, and the power and cooling infrastructure that supports them. The economic signature of this grouping is high capital intensity, long lead times, and concentration: at the leading edge, each of these layers is dominated by between one and three firms, and entry requires both very large fixed-asset investment and access to scarce technological inputs that themselves originate from another concentrated layer. As documented in the next section, the principal bottleneck at any moment in the cycle tends to be located in this grouping; currently, the relevant constraints concern advanced packaging capacity (CoWoS at TSMC), HBM supply, and grid-interconnection timelines.

The midstream grouping intermediates compute. It comprises hyperscale cloud infrastructure (provided principally by AWS, Microsoft Azure, Google Cloud, and more recently Oracle

Cloud Infrastructure) and a smaller class of specialised AI-compute providers, often termed “neoclouds,” whose business model is the provision of access to GPU-dense clusters to model developers and large enterprise customers (CoreWeave, Lambda, Crusoe, Nebius). The hyperscale layer is oligopolistic and exhibits revenue growth that is partly driven by AI workloads and partly by the long-established commodification of general-purpose cloud computing; the neocloud layer is concentrated, smaller in absolute size, and operationally tied to upstream chip availability and to a small number of large customers. Both layers are capital-intensive but differ in financing structure: the hyperscalers finance buildout primarily from operating cash flow and corporate debt, while neoclouds rely on asset-backed and project-finance structures examined in §3.3.

The downstream grouping consists of activities whose output is software: foundation models, integration layers, and end-user applications. Foundation-model development (OpenAI, Anthropic, Google DeepMind, Meta, Mistral, xAI, and a small number of others) is capital-intensive in its training phase but software-like in its marginal cost structure; the principal cost item is compute purchased from the midstream layer, with the consequence that vertical integration, or contractual approximations to it, has become a defining feature of the sub-sector. The application and enterprise-integration layer is more fragmented and exhibits a wider distribution of business models, ranging from large incumbents adding AI features to existing software franchises (Microsoft Copilot, Salesforce Einstein, Adobe Firefly, ServiceNow Now Assist) to specialised application firms whose products depend on access to third-party foundation models (GitHub Copilot, Cursor, Windsurf, and the broader cohort of vertical AI software firms).

The capital cycle interacts asymmetrically with these three groupings. Capital expenditure is concentrated upstream and midstream, where the physical buildout is occurring; reported revenue growth is observed at every level, but the share originating from external (non-ecosystem) demand is most defensible at the application and enterprise-software layer, where it is also smallest in absolute terms. The asymmetry is structural: the largest dollar flows are in the layers least exposed to direct end-user demand signals, while the layer most exposed to such signals is the smallest. This is the geometry that any rigorous assessment of the cycle’s sustainability must take as its starting point.

Table 1 — The AI value chain: function, structure, and bottlenecks

<i>Layer</i>	<i>Function</i>	<i>Market structure</i>	<i>Principal actors</i>	<i>Principal bottleneck</i>	<i>Relevance to capital cycle</i>
<i>Semiconductor equipment</i>	Lithography, deposition, etch, metrology tools for chip fabrication	Oligopolistic; EUV monopolised	ASML, Applied Materials, Lam Research, Tokyo Electron, KLA	EUV throughput; High-NA ramp	Upstream bottleneck for all chip layers
<i>Chip design (AI accelerators)</i>	Architecture and design of GPUs, TPUs, custom AI silicon	Highly concentrated in training; more fragmented in inference	Nvidia (dominant), AMD, Intel; Broadcom & Marvell (custom silicon); Google TPU, AWS Trainium	Software ecosystem lock-in; performance per watt	Largest single value-capture node
<i>HBM (memory)</i>	High-bandwidth memory stacked on accelerators	Three-firm oligopoly	SK Hynix, Samsung, Micron	HBM3e / HBM4 capacity; advanced-process yields	Persistent supply constraint
<i>Foundry (leading-edge)</i>	Fabrication of advanced-node semiconductors	Quasi-monopoly at the leading edge	TSMC; Samsung Foundry; Intel Foundry; SMIC	Leading-edge wafer capacity	Single point of failure in chip supply
<i>Advanced packaging</i>	2.5D / 3D integration of accelerator + HBM	Highly concentrated (TSMC CoWoS)	TSMC (CoWoS-S, CoWoS-L); ASE; Amkor	CoWoS capacity ramp	Current binding constraint on accelerator supply
<i>Data-centre construction & power</i>	Physical buildout; electrical and cooling capacity	Fragmented in construction; concentrated in hyperscale-owned sites	Digital Realty, Equinix; hyperscaler self-build; specialised EPC firms	Grid interconnection; power availability	Becoming the binding constraint as chip supply normalises
<i>Hyperscale cloud</i>	General-purpose + AI compute services at scale	Oligopolistic (four-firm)	AWS, Microsoft Azure, Google Cloud, Oracle Cloud	Capacity expansion vs. AI workload demand	Largest capex deployers; ecosystem hub
<i>Specialised AI cloud (neoclouds)</i>	GPU-dense capacity, often single-tenant	Concentrated among 4–6 firms	CoreWeave, Lambda, Crusoe, Nebius	Chip allocation; customer concentration	Levered, asset-backed financing; principal risk-transfer locus
<i>Data services</i>	Data sourcing, labelling, vector databases, RAG infrastructure	Fragmented	Scale AI, Labelbox, Pinecone, Weaviate; in-house teams	Data quality, licensing	Smaller cost share but operationally critical
<i>Foundation models</i>	Training and provision of frontier models	Concentrated at the frontier	OpenAI, Anthropic, Google DeepMind, Meta, Mistral, xAI, DeepSeek	Compute access; model quality	Layer at which ecosystem-internal cash flows converge
<i>Applications & enterprise integration</i>	End-user products embedding AI capabilities	Fragmented at start-up layer; concentrated at incumbent layer	Microsoft (Copilot), Salesforce, ServiceNow, Adobe, GitHub; Cursor, Windsurf	Pricing power vs. inference cost	Layer most exposed to end-user demand; smallest in absolute revenue

Source: author's synthesis from company filings and publicly disclosed industry sources. The grouping reflects the analytical framing of this chapter; alternative groupings are possible.

3.1.3 Key actors and market structure

This section informs the previously described framework and identifies the principal firms at each layer of the value chain and the form of dominance (market share, technological control, infrastructure ownership, or financing relationships) that each exhibits. The discussion proceeds from upstream to downstream. Specific market-share figures are presented with explicit reference to source and methodology; where commercial estimators disagree, the divergence is noted.

Semiconductor equipment

ASML occupies a quasi-monopoly position at the leading edge: it is the sole producer of extreme-ultraviolet (EUV) lithography systems, the optical tools necessary to produce the logic chips below the seven-nanometre node and the densest DRAM and HBM stacks¹⁷.

Applied Materials, Lam Research, Tokyo Electron, and KLA dominate the remaining

<i>Metric</i>	ASML FY2025	segments of the semiconductor equipment market,
<i>Net sales</i>	€32.7bn	particularly deposition, etching, inspection, and
<i>Gross margin</i>	52.8%	metrology systems. Although none holds a monopoly
<i>Net income</i>	€9.6bn	comparable to ASML's position in EUV lithography,
<i>Year-end backlog</i>	€38.8bn	their competitive positions are reinforced by highly
<i>Q4 net bookings</i>	€13.2bn	specialised technologies, long customer qualification
<i>Q4 EUV bookings</i>	€7.4bn	processes, and substantial switching costs. ¹⁸
<i>2030 revenue guidance</i>	€44–60bn	

Chip design and AI accelerators

Nvidia's dominance is the most extreme single-firm concentration documented in the current cycle. In fiscal year 2026 (end of January 2026), the firm reported total revenue of \$216B, of which the Data Center segment accounted for approximately \$198B across the four quarters, roughly an order of magnitude above the next-largest competitor as in Figure 5¹⁹. Non-GAAP gross margins remained in the 73% - 75% range according to 8-K filings. Nvidia's revenue base is highly concentrated among a small number of hyperscale customers. The Q3 fiscal 2025 CFO


¹⁷ [ASML Holding NV](#), Form 6-K Q1 2025 press release (SEC filing).

¹⁸ SEMI, Worldwide Semiconductor Equipment Market Statistics Report, 2025; Lam Research, Form 10-K Annual Report, 2025, discussion of customer qualification processes and competitive conditions; KLA Corporation, Annual Report 2025, discussion of process-control integration and switching costs.

¹⁹ [Nvidia Corporation](#), Q4 and full-year fiscal 2026 results press release (25 February 2026).

commentary indicated that two direct customers alone accounted for approximately 39% of total revenue, while cloud-service providers represented roughly half of Data Center revenue. Nvidia’s Q3 fiscal 2026 10-Q discloses guarantee up to \$860M of a partner’s facility-lease obligations in exchange for warrants (rights to buy equity later)²⁰, a significant arrangement examined further in §3.3.

Figure 5 - Nvidia Vs AMD, data center segment revenue



Intel does not separately disclose material revenue contribution from its Gaudi accelerator platform within Data Center and AI segment reporting²². Custom silicon developed by hyperscalers (Google TPU, AWS Trainium and Inferentia, Microsoft Maia) is material in their internal allocation decisions but is not sold to third parties and is therefore not measured on a comparable revenue basis.

High-bandwidth memory

HBM is concentrated among three firms: SK Hynix, Samsung, and Micron. Industry reporting through late 2025 consistently identified SK Hynix as the leading supplier of advanced HBM products, particularly following early HBM3e qualification for Nvidia systems, while Micron expanded rapidly from a comparatively small base²³. TrendForce reports

Supplier	Approximate late-2025 share
SK Hynix	~55–60%
Samsung	~25–30%
Micron	~15–20%

²⁰ [Nvidia Corporation](#), Form 10-Q for the quarter ended 26 October 2025 (SEC filing).
²¹ [Advanced Micro Devices](#), Form 10-K (Annual Report) for fiscal 2025 (SEC filing).
²² Intel Corporation, *Form 10-K Annual Report*, 2025; Intel, *Q4 2024 and Full-Year Earnings Release*, 2025.
²³ TrendForce, *HBM Market Outlook*, 2025; Counterpoint Research estimates reported in *Chosun Biz*, Q2 2025; Micron Technology, FY2025 earnings commentary.

HBM demand growth of approximately 130% YoY in 2025, and capacity expansion is constrained by 18-to-36-month DRAM capital cycles²⁴.

Foundry

TSMC holds a near-monopoly in the leading-edge segment and a clear majority of the pure-play foundry market overall. Counterpoint Research data for Q3 2025 placed TSMC’s pure-play foundry share at 72%²⁵; the remaining shares are distributed as depicted in Figure 6. The economic relevance of this distribution is that the residual firms are not viable substitutes for TSMC at the leading edge. Nodes below

Pure-play foundry market share, Q3 2025




Figure 6 - Foundry segment – Pure play Market Share

Source: Counterpoint Research. TrendForce Q4 2025 places TSMC share at 70.4%.

seven nanometers represented approximately three-quarters of TSMC’s wafer revenue in Q3 2025, and 3-nanometre capacity was reported to be at full utilization through 2024 and 2025²⁶. TSMC’s gross margin in 2024 was 56.1 per cent; its 2024 capital expenditure exceeded \$30B²⁷. Samsung’s 3-nanometre yields have reportedly lagged TSMC’s, and Intel Foundry’s 18A node (1.8 nanometer-class technology) is still in testing phase.²⁸

Advanced packaging

Advanced packaging has emerged as a critical bottleneck within the AI semiconductor supply chain²⁹. CoWoS (Chip-on-Wafer-on-Substrate), TSMC’s dominant 2.5D packaging technology, is widely used to integrate AI accelerators with high-bandwidth memory (HBM). The technology is central to Nvidia’s accelerator architecture and became a major supply

²⁴ TrendForce, *HBM Market Outlook 2025*, 2025.
²⁵ [Counterpoint Research](#), Global Pure Foundry Market Share quarterly tracker.
²⁶ Taiwan Semiconductor Manufacturing Company, *Q3 2025 Earnings Presentation*.
²⁷ Taiwan Semiconductor Manufacturing Company, *Annual Report 2024*.
²⁸ Reuters, “Samsung struggles to close advanced-node yield gap with TSMC,” 2025; Intel Corporation, *Foundry Direct Connect Investor Materials*, 2025.
²⁹ Nvidia Corporation, *FY2025 Earnings Calls*; TSMC earnings commentary discussing CoWoS capacity constraints

constraint during the AI infrastructure expansion of 2024–2025. TSMC management repeatedly identified CoWoS capacity as a limiting factor on AI accelerator supply during earnings commentary throughout 2024 and 2025. Although exact production volumes are not publicly disclosed, industry analysts consistently described TSMC as the dominant provider of advanced AI packaging capacity. Industry reports further indicated that supply constraints were expected to ease gradually as TSMC expanded CoWoS production and outsourced portions of advanced packaging to partners including ASE and Amkor³⁰.

Hyperscale cloud

Cloud infrastructure market share, Q4 2025

Worldwide cloud infrastructure revenue share (%) — IaaS + PaaS + hosted private cloud




Figure 7 - Cloud infrastructure, Market Share

Synergy Research Group data for Q3 2025 placed the worldwide cloud infrastructure market at \$107B (quarterly), with the three largest providers holding the in Q4 the shares in Figure 7, with Oracle’s share rising into the mid-single digits on the strength of recent contract wins³¹. The aggregate share of the three leaders has remained narrowly within the 62–63% band over four consecutive quarters. Microsoft disclosed Intelligent Cloud revenue consistently with analyst estimates placing Azure annual revenue above \$75B in fiscal 2025³²; AWS recorded approximately \$130B in calendar-2025 revenue and approximately \$45B in operating profit, representing 57% of Amazon’s consolidated operating profit on 18% of consolidated revenue. The growth-rate dispersion within the Big Three is large: Azure expanded at approximately 33–40% YoY through 2025, Google Cloud at 28%–48% (with Q4 2025 quarterly growth reported at 48%), and AWS at 18–19%.

³⁰ Taiwan Semiconductor Manufacturing Co., *Quarterly Earnings Calls and Investor Presentations*, 2024–25; TrendForce, *Advanced Packaging Industry Analysis*, 2025; ASE Technology Holding investor materials, 2025

³¹ [Synergy Research Group](#), "Cloud Market Share Trends - Big Three Together Hold 63 %" (Q3-Nov 25 data)

³² Microsoft Corporation, Form 8-K Q4 fiscal 2025 [Press release](#) (30 July 2025) — Azure surpassed \$75 B in annual revenue

Specialised AI cloud (neoclouds)

This sub-sector is concentrated among four to six firms whose business model is the provision of GPU-dense capacity under long-term contractual commitments to a small number of large customers. CoreWeave is the most informative case because it is publicly listed and reports under SEC disclosure standards. The Form 10-K for fiscal 2025 reports total revenue of \$5.13B (up from \$1.92B in 2024 and \$229M in 2023); a persistent net loss of \$1.17B ; debt of approximately \$14.2B ; and a remaining performance obligation balance of \$60.7B as of 31 December 2025, more than four times the level at the start of the year³³. Customer concentration is structurally extreme: a single customer (identified in subsequent public reporting as Microsoft) accounted for 67% of full-year 2025 revenue³⁴; the top three customers (Microsoft, OpenAI, Meta) represent the majority of the contract backlog³⁵. Capital expenditure for 2026 is guided at \$30–35B , compared with \$10.3B in 2025. CoreWeave operated 43 data-centre sites with approximately 850 megawatts of active power capacity and 3.1 GW of contracted capacity at year-end 2025 (Figure 8). Nebius is the second-largest publicly traded neocloud (an Amsterdam-listed entity); Lambda and Crusoe remain privately held and are smaller in scale.

CoreWeave: revenue growth against persistent losses and customer concentration

Revenue and net loss (USD bn, FY2023–FY2025) - FY2025 revenue by customer




Figure 8 - CoreWeave: scale, profitability, customer concentration

³³ [CoreWeave](#) 10-K analysis (FY 2025): revenue \$5.13 B; net loss \$1.17 B; debt \$14.2 B; RPO \$60.7 B; Microsoft = 67 % of FY revenue.

³⁴ [CoreWeave, Inc.](#), Form 10-Q for the quarter ended 30 September 2025 (SEC filing).

³⁵ [Sacra](#), CoreWeave revenue and concentration analysis (Microsoft, OpenAI, Meta as top three customers).


Foundation models

Revenue at this layer is reported by privately held firms and is not subject to standard disclosure; the figures are drawn from secondary reporting (principally *The Information*, the *Financial Times*, Bloomberg, CNBC, and Epoch AI) and should be treated as indicative because they are not audited. The relative position of the two principal pure-play developers inverted during the sample window: having trailed OpenAI throughout 2024–2025, Anthropic crossed above it on reported run-rate in April 2026 and pulled sharply ahead thereafter (Figure 9), driven principally by enterprise adoption of its Claude Code agentic coding product. The crossover reflects Anthropic's disclosed acceleration more than decline at OpenAI, whose run-rate was last publicly dated at approximately \$25B in February–March 2026 and which has not published a comparable figure since; the gap shown in the later periods should therefore be read as a gap in OpenAI disclosure as much as a gap in revenue. The crossover that Epoch AI had projected for 2026–2027 under conservative assumptions therefore occurred near the early edge of that window. Both firms remain unprofitable on a free-cash-flow basis and carry valuations that are large multiples of run-rate revenue, and both are now reported to have filed or to be preparing for public listing during 2026. A material interpretive caveat attaches to the Anthropic figures specifically: the firm is reported to recognise cloud-reseller revenue on a gross basis (counting total end-customer spend as revenue and booking partner payouts as expense), which inflates its headline run-rate relative to net-reporting peers—so the reported crossover should be read with that accounting difference in mind. The projected loss trajectories are discussed in §3.2.3. Diversified developers (Google, Meta, Microsoft) pursue model development internally without disclosing model-level revenue; xAI, Mistral, and DeepSeek operate at materially smaller scale.

Figure 9 - OpenAI and Anthropic reported ARR trajectory

Reported annualised revenue run-rate: OpenAI and Anthropic

Press-reported run-rate (USD billions) — not company-confirmed under regulatory disclosure



All figures from press reporting and company blog posts (principally *The Information*, the *Financial Times*, Bloomberg, CNBC, Epoch AI) — not company-confirmed under regulatory disclosure. Anthropic reports cloud-reseller revenue on a gross basis, which inflates headline run-rate relative to non-reporting peers; the Apr-26 (\$30B) and May-26 (\$47B) figures should be interpreted with caution.

Table 2 — Foundation models: principal pure-play developers

<i>Metric</i>	<i>OpenAI</i>	<i>Anthropic</i>
<i>Latest reported run rate</i>	~\$25bn (Feb/Mar 2026)	~\$47bn (May 2026)
<i>Most recent disclosed valuation</i>	\$852bn (Mar 2026, \$122bn round)	~\$965bn (Series H, late May 2026)
<i>IPO status</i>	preparing	confidential S-1 filed; listing possibly Oct 2026

Application and enterprise integration

The downstream layer is more fragmented but is led, in revenue terms, by incumbent enterprise-software firms that have integrated AI features into existing product franchises. Microsoft is the most informative case in terms of disclosed magnitude: the firm's Q2 fiscal 2025 release (Jan 2025) disclosed an AI-business annualized revenue run rate of \$13B, +175% YoY³⁶; the subsequent disclosure in the Q3 fiscal 2026 results (30 April 2026) revised that figure to \$37B, +123% YoY³⁷. Microsoft's commercial RPO (remaining performance obligations) trajectory is portrayed in Figure 10³⁸, numbers of direct relevance to the circular-financing discussion in §3.3. The Microsoft trajectory is illustrated in Figure 10. Salesforce, ServiceNow, Adobe, and GitHub disclose AI features but generally do not segregate AI-specific revenue in regulatory filings, with the consequence that comparable cross-firm figures are unavailable. Specialized application-layer firms (Cursor, Windsurf, and the broader cohort of vertical AI software) are reported to be growing rapidly but operate at scales an order of magnitude below the incumbent layer.

³⁶ [Microsoft Corporation](#), Form 8-K Q2 fiscal 2025 press release (29 January 2025) — first disclosure of AI business \$13 B annualised run rate.

³⁷ [Microsoft](#) Q3 fiscal 2026 results (30 April 2026), reporting \$37 B AI run rate (up 123 % YoY).

³⁸ [Microsoft Corporation](#), Form 10-Q for the quarter ended 31 March 2026 (SEC filing) — \$627 B commercial RPO; \$633 B total RPO.




Figure 10 - Microsoft AI business and contracted backlog

Three features emerge from this layer-by-layer mapping. First, most layers of the value chain is characterized by oligopolistic or quasi-monopolistic concentration. Even the most fragmented layer (enterprise application integration) is in practice led, by absolute revenue, by a small number of incumbent software franchises. Second, the value-capture distribution is highly skewed toward two specific positions in the chain: the AI-accelerator design layer (Nvidia, with limited contestation by AMD) and the hyperscale cloud layer (the four-firm group). The combined revenue of these positions exceeded \$400B in 2025, with growth rates of 25–70% YoY. Third, customer concentration is always elevated: Nvidia’s Data Center segment depends on a small number of cloud customers for approximately half of its revenue; CoreWeave depends on a single customer for two-thirds of its revenue; Microsoft Azure’s contract backlog is attributed in significant part (approximately 45%) to a single external counterparty. These concentrations are the precondition for the circular-financing patterns that §3.3 will examine in detail.

3.2 Demand, Constraints, and Economic Sustainability


Section 3.1 mapped the supply structure of the AI industry: a value chain whose constituent layers are oligopolistic or quasi-monopolistic, whose value capture is concentrated in two specific positions (accelerator design and hyperscale cloud), and whose customer relationships are themselves structurally concentrated. The present section turns to the demand side. Three aspects are described: what magnitudes of capital expenditure are being committed and what magnitudes of end-user revenue are forecast to justify that expenditure (§3.2.1); what physical and technical bottlenecks may bind the buildout (§3.2.2); and at what point operational and concentration risks become first-order considerations for the sustainability of the cycle

(§3.2.3–§3.2.4). The analysis is descriptive: this section assembles the relevant magnitudes, identifies the principal divergences among forecasts, and documents the data on the constraints.

3.2.1 Demand, revenue, and capex projections

The capital expenditure being deployed in the current cycle has reached magnitudes without comparable precedent in the modern technology industry. Aggregate annual capital expenditure of the five largest US hyperscalers (Amazon, Alphabet, Microsoft, Meta, and Oracle) rose from approximately \$162 B in 2022 to \$408 B in 2025, a 2.5-

fold increase over three years. Epoch AI's tracking of quarterly cash property-plant-and-equipment and new finance leases from SEC filings places the annualised growth rate at approximately 72% between Q2 2023 and Q4 2025 (90% C.I.: 66–78%)³⁹.



Amazon, Alphabet, Microsoft, Meta, Oracle. 2026 ranges: Amazon \$200 B; Alphabet \$175–185 B; Microsoft \$110–120 B; Meta \$115–135 B; Sources: Visual Capitalist / Epoch AI tracker (SEC filings); Futurum; Data Center Frontier; company earnings calls.

Figure 11 - Aggregate annual capital expenditure of hyperscalers in Billions of USD dollars

Announced 2026 capital expenditure

plans from these same five firms aggregate to between approximately \$602B (the IEEE Communications Society estimate⁴⁰) and \$690B (the Futurum Group estimate⁴¹), a 50–70% year-on-year increase, with every one of the five firms raising its 2026 commitment well above its 2025 level. The firm-by-firm figures and the trajectory are shown in Figure 11; the relevant feature here is the uniformity and steepness of the increase across the group, not any single firm's number.

Two features of this trajectory are distinctive. First, capital expenditure as a share of revenue, also said capital intensity, has reached 45–57% for the largest hyperscalers, levels historically associated with utility and heavy-industry firms rather than technology firms⁴². Second, aggregate capital expenditure now exceeds operating cash flow at multiple firms in the group, with the consequence that hyperscalers have begun raising substantial debt to fund the buildout:

³⁹ [Epoch AI](#), "Hyperscaler capex has quadrupled since GPT-4's release" (February 2026) — 72 % annualised growth Q2 2023–Q4 2025.

⁴⁰ [IEEE Communications Society](#) Technology Blog, "Hyperscaler capex > \$600 bn in 2026 — a 36 % increase over 2025" (22 December 2025).

⁴¹ [Futurum Group](#), "AI Capex 2026: The \$690 B Infrastructure Sprint" (12 February 2026).

⁴² [Introl](#), "Hyperscaler CapEx Hits \$600 B in 2026" — Goldman Sachs \$1.15 T cumulative 2025–27 forecast; capital intensity 45–57 %; \$108 B debt raised in 2025.

approximately \$108B was raised by the Big Five in 2025 alone according to Introl-aggregated data, and Morgan Stanley research projects that approximately \$1T of technology-sector commitment to AI infrastructure will be required across 2025 and 2026⁴³, with Morgan Stanley separately estimating that total data-centre-related capital expenditure of approximately \$3T remains to be deployed, of which less than 20% had been spent at the time of its January 2026 outlook⁴⁴. Goldman Sachs equity research projects \$1.15 of cumulative hyperscaler capital expenditure across 2025–2027, more than double the \$477B spent in 2022–2024⁴⁵.

Longer-horizon forecasts are produced principally by three institutions whose methodologies differ in instructive ways. McKinsey & Company, in *The cost of compute: A \$7T race to scale data centers* (April 2025), projects \$6.7T of cumulative data-centre capital expenditure to 2030, of which \$5.2T is allocated to AI workloads and \$1.5T to non-AI workloads⁴⁶. McKinsey's base, low, and high scenarios span roughly \$3.7–7.9T to 2030, indexed to the GW of AI capacity assumed.

A central question raised by these magnitudes is whether end-user revenue can plausibly be expected to justify the capital deployed. The question was formalized in June 2024 by David Cahn of Sequoia Capital, who applied a four-step calculation: doubling Nvidia's annualised data-centre revenue to account for non-chip components of total cost of ownership (power, buildings, cooling, backup generation), and then doubling again to reflect a 50% gross margin requirement for the end-user of the GPU compute. Applied to Nvidia's Q4 fiscal 2024 data-centre run rate, the calculation produced a \$600B annual end-user revenue requirement against actual reported AI revenue of an order of magnitude lower⁴⁷. Updated to Nvidia's Q4 fiscal 2026 run rate, the same calculation implies a required end-user revenue of roughly \$1.0T against reported pure-play AI revenue of only a few tens of billions — a gap on the order of \$945B. The full revenue gap depends on how cloud-resident AI revenue and incumbent enterprise-software AI revenue (Microsoft Copilot, Google Workspace AI features, Salesforce,

⁴³ [Morgan Stanley Research](#), "Powering AI: Markets Race to Invest in AI Energy Solutions" (27 February 2026) — \$1 T 2025–26 commitment.

⁴⁴ [Morgan Stanley Research](#), "AI Market Trends 2026" — \$3 T data-centre capex estimate, less than 20 % deployed.

⁴⁵ [Introl](#), "Hyperscaler CapEx Hits \$600 B in 2026" — Goldman Sachs \$1.15 T cumulative 2025–27 forecast (continued reference for §3.2.1 closing claim).

⁴⁶ [McKinsey & Company](#), "The cost of compute: A \$7T race to scale data centers" (April 2025) by J. Noffsinger, M. Patel and P. Sachdeva.

⁴⁷ [David Cahn](#), "AI's \$600B Question", Sequoia Capital (20 June 2024).

ServiceNow) are counted; nevertheless, the order-of-magnitude difference is the point of the calculation.

The same question was raised, contemporaneously, by Goldman Sachs equity research in the June 2024 report *Gen AI: Too Much Spend, Too Little Benefit?*, which characterized the approximately \$1T of generative-AI capital expenditure committed at the time as having "so far little to show for it" and questioned whether the spend would pay off⁴⁸. The Goldman report incorporated explicitly divergent internal views: Head of Global Equity Research Jim Covello argued that current AI technology is not designed to solve the kinds of complex problems that would justify its costs, while sector analysts Kash Rangan and Eric Sheridan held more constructive positions on the returns to capital deployment. The MIT economist Daron Acemoglu, cited in the same report, projected approximately 0.5% total productivity uplift and approximately 0.9% GDP uplift over the following decade, figures materially below Goldman's own 6.1% GDP uplift estimate. The principal conclusion of the present section is that the gap between announced infrastructure capital expenditure and reported AI-attributable revenue (the magnitudes above) is not, currently, accompanied by an empirically demonstrated revenue trajectory of compatible magnitude.

3.2.2 Physical and technical constraints

The capital expenditure projections discussed in §3.2.1 are not unconstrained by physical reality. At every stage of 2024 and 2025, the build-out has been limited by one or more identifiable bottlenecks in the supply of inputs to AI data centers. The principal constraints have been, in approximate chronological order of their period of binding: high-bandwidth memory (binding through 2024); advanced packaging capacity (binding throughout 2024 and 2025, and projected to remain binding into 2026); leading-edge wafer capacity (recurrently tight at TSMC's 3-nanometre and incoming 2-nanometre nodes); and, increasingly visible in 2025 and projected to dominate from 2026 onward, electrical power and grid-interconnection capacity. This subsection documents the principal constraints and the dynamic of bottleneck succession that the data suggest.

⁴⁸ [Goldman Sachs Research](#), "Gen AI: Too Much Spend, Too Little Benefit?" — Top of Mind report (June 2024).

Advanced packaging

CoWoS (Chip-on-Wafer-on-Substrate), the 2.5D packaging technology used to integrate AI accelerators with high-bandwidth memory, has been described by TSMC management as the single tightest part of the AI semiconductor supply chain. TSMC's CoWoS capacity expanded from approximately 13–16 thousand wafers per month in 2023 to approximately 35–40 thousand by late 2024, and is projected to reach approximately 65–75 thousand wafers per month by end-2025 and 120–140 thousand by end-2026, a near-quadrupling of capacity in two years⁴⁹. On the Q3 2025 earnings call, TSMC CEO C.C. Wei stated that CoWoS capacity remained "very tight" and was "sold out through 2025 and into 2026"⁵⁰; the same Deloitte-tracked figures indicate that TSMC is allocating approximately 10–20% of its 2025 capital expenditure to advanced-packaging facilities⁵¹. Industry reporting indicates that Nvidia has secured between 50 and 70% of TSMC's available 2025–2026 CoWoS-L capacity, with bookings extending into 2026 and 2027⁵². The structural significance of this constraint is twofold: it is geographically concentrated in a single supplier in a single country (TSMC, Taiwan), and the principal customer is also the principal beneficiary of the constraint, which both validates and partially explains Nvidia's pricing power.

High-bandwidth memory and leading-edge wafer capacity

HBM supply was reported by all three suppliers (SK Hynix, Samsung, Micron) as substantially or fully allocated for 2026 by the time of their Q3 2025 earnings calls, with Samsung guiding to high-teens to low-twenties% price increases for HBM in 2026 contracts. TSMC has separately stated that advanced-node wafer demand exceeds available supply by approximately a factor of three; 3-nanometre capacity reached full utilization in 2024 and remained so through 2025, while 2-nanometre availability is structurally constrained as the node ramps. The relevant feature for the present analysis is not the existence of these constraints in isolation since capacity constraints are a recurring feature of the semiconductor cycle, but it is the simultaneity

⁴⁹ [SemiWiki](#), "CoWoS Capacity Set to Skyrocket by 2026" — TSMC capacity figures: 2023 ~15 K WPM; 2024 ~38 K WPM; 2025e ~65–75 K WPM.

⁵⁰ [Fusion WorldWide](#), "Inside the AI Bottleneck: CoWoS, HBM, and 2–3 nm Capacity Constraints Through 2027" (December 2025) — TSMC: CoWoS "sold out through 2025 and into 2026".

⁵¹ [Longyield](#) (Deloitte-cited), "TSMC at the Center of the AI Boom" — TSMC CEO C.C. Wei: CoWoS capacity "very tight"; 10–20 % of 2025 capex on advanced packaging.

⁵² [Introl](#), "CoWoS and Advanced Packaging" — Nvidia secured ~70 % of TSMC 2025 CoWoS-L capacity; bookings extending into 2026–27.

of binding constraints across multiple distinct supply chains, each with its own multi-year lead time. DRAM capital cycles run 18 to 36 months from investment decision to production output; leading-edge logic wafer capacity expansions run 24 to 36 months; advanced-packaging capacity expansions are similar. The binding constraint shifts from one input to another as each is partially relieved, but the aggregate constraint on the rate at which AI compute can be brought online does not relax materially within the planning horizon.

Electrical power and grid interconnection

The constraint that has become increasingly visible during 2025-26 is the availability of electricity and, more specifically, the timing of grid interconnection. The International Energy Agency's Energy and AI report (April 2025) places global data-centre electricity consumption at 415 terawatt-hours in 2024 (approximately 1.5% of global electricity consumption)⁵³. Its April 2026 follow-up reports double-digit growth in total data-centre electricity in 2025 and roughly 50% growth at AI-focused facilities specifically, the latter the fastest-growing component⁵⁴. On the IEA base case this roughly doubles by 2030 (a magnitude comparable to the entire electricity demand of Japan) with the United States, already the largest single consumer, accounting for close to half of the global total and rising fastest⁵⁵ (Figure 12, left panel).

The principal binding constraint at the US level is not aggregate generation capacity but the interconnection process itself. Lawrence Berkeley National Laboratory's Queued Up: 2025 Edition report (covering data through end-2024) documents approximately 10,300 projects active in US grid-interconnection queues, representing 1,400 GW of generation and 890 GW of storage, a total of approximately 2,290 GW, or nearly twice the existing US installed generating capacity⁵⁶. Two features of this dataset matter here. First, historically only a small fraction of queued capacity is ever built, on the order of one in seven, so the headline queue

⁵³ [International Energy Agency](#), Energy and AI (April 2025): data centres 415 TWh in 2024 (1.5 % of global electricity); base case projects 945 TWh by 2030.

⁵⁴ [IEA](#), Key Questions on Energy and AI (April 2026): data-centre electricity 485 TWh in 2025 (+17 % YoY); AI-focused data centres +50 % in 2025.

⁵⁵ [Pew Research Center](#), "What we know about energy use at U.S. data centers amid the AI boom" (October 2025) — US data centres 183 TWh in 2024 (4 %+ of US electricity); projected to grow 133 % to 426 TWh by 2030.

⁵⁶ [Lawrence Berkeley National Laboratory](#), Energy Markets & Policy: Queued Up: 2025 Edition — Characteristics of Power Plants Seeking Transmission Interconnection (data through end 2024).

figure overstates deliverable capacity by a wide margin⁵⁷. Second, interconnection wait times have lengthened to several years from request to operation. The implication for the AI infrastructure cycle is that even with announced capital available, the binding constraint on the pace at which incremental compute capacity can be deployed is the rate at which new generation can be interconnected to the grid (Figure 12, right panel).




Figure 12 - Data center electricity demand and US grid interconnection queue

The shifting bottleneck

A pattern is observable in the succession of binding constraints across the cycle to date. From late 2022 to early 2024, the binding constraint was GPU supply itself, mediated by TSMC wafer allocation. From mid-2024 to 2025, the binding constraint shifted to advanced packaging (CoWoS), and within this constraint, the principal beneficiary of the rationing has been the principal customer. The constraint is projected to ease through 2026 as TSMC's capacity expansion comes on line. From 2026 onward, multiple sources (TSMC, Morgan Stanley, the IEA, LBNL, and the hyperscalers themselves on their earnings calls) converge on the view that the principal binding constraint will be electrical power and grid interconnection. The hyperscalers' response to this anticipated bottleneck has been to enter directly into long-term power-purchase agreements: the IEA reports that the technology sector accounted for approximately 40% of all corporate renewables power-purchase agreements signed in 2025, and that the pipeline of conditional offtake agreements between data-centre operators and small-modular-reactor nuclear projects rose from approximately 25 GW at end-2024 to approximately

⁵⁷ [Lawrence Berkeley Lab](#) 2025 Interconnection Queue Report — 2,290 GW total active capacity; average completion 55 months; historical build-out 14 % of 2000–18 queue capacity.

45 GW by April 2026⁵⁸. The pivot from chip-bound to power-bound reflects the maturation of the cycle: the upstream chip supply chain has expanded faster than the surrounding physical infrastructure, with the consequence that the marginal infrastructure dollar in 2026 must be deployed against constraints that are slower to relieve.

Two implications follow for the analysis of subsequent sections. First, the magnitude and persistence of the constraints documented here are inconsistent with a market in which supply can equilibrate rapidly to demand. The forecasts of \$5–7T of cumulative capital expenditure to 2030 require sustained physical buildout at a rate that, on the evidence of 2024–2025, is itself the binding constraint on the cycle’s realisation. Second, the substitution of one binding constraint for another, from chip availability to power availability, produces a buildout in which the geographic and counterparty concentration shifts but does not diminish: the central role of TSMC in 2024–2025 is being matched, prospectively, by the central role of a small number of grid operators, regulated utilities, and independent power producers in 2026 and beyond. The economic-sustainability and concentration-risk questions to which Sections 3.2.3 and 3.2.4 now turn are inseparable from this geometry of constraints.

3.2.3 Economic sustainability of the AI model

The preceding subsections have documented two features of the cycle: the magnitudes of capital being committed (§3.2.1), and the physical constraints that bind the rate at which capital can be deployed against computational capacity (§3.2.2). The present subsection turns to the question that follows: is the economic model under which this capital is being committed sustainable on the revenue it can reasonably be expected to generate? The question is approached through three sources of evidence: the financial trajectories of the principal pure-play AI firms, the trajectory of unit economics (inference costs), and the trajectory of end-user adoption and demonstrated productivity, jointly delineating the boundaries of the empirical question.

Pure-play AI firm financials

A useful and directly observable indicator of economic sustainability at the foundation-model layer is the projected loss trajectory of the two largest pure-play firms to date. According to OpenAI financial documents shared with investors in mid-2025 and reported by the Wall Street

⁵⁸ [IEA](#), Key Questions on Energy and AI (April 2026) — technology sector ~40 % of 2025 corporate renewables PPAs; SMR conditional offtake pipeline 25→45 GW between Q4 2024 and Q2 2026.

Journal in November 2025, OpenAI expected a 2025 net loss of approximately \$9B on \$13B of revenue (total spending of roughly \$22B, equivalent to approximately \$1.69 spent per dollar of revenue), with losses rising to a peak of approximately \$74B in 2028 (roughly three-quarters of that year's projected revenue, driven principally by compute costs) before a transition to profitability beginning in 2030⁵⁹. On the same documents, OpenAI's cash burn was projected to remain at approximately 57% of revenue in 2026 and 2027. Subsequent reporting in February and April 2026 revised the picture in two directions: the 2030 revenue target was raised to approximately \$280B (from approximately \$200B) and cumulative cash burn through 2030 to approximately \$665B, while the 2028 loss estimate was revised upward to approximately \$85B and 2028 compute spend to approximately \$121B⁶⁰. These figures are not strictly comparable across vintages, as outlets report variously operating loss, net loss, and net cash burn; the consistent feature is the direction and magnitude, not the precise line item.

The same Wall Street Journal documents indicate a fundamentally different trajectory for Anthropic: cash burn falling to approximately one-third of revenue by 2026 and approximately 9% by 2027, with break-even projected for 2028, predicated on substantially lower training costs (reported at approximately one-quarter of OpenAI's)⁶¹. Both projections are conditional on substantial continued revenue growth from a low base. The cumulative pre-profit cash absorption of the two principal firms, OpenAI approximately fourteen times that of Anthropic before break-even, on the WSJ figures, is itself an indication that the foundation-model layer cannot self-finance and is therefore structurally dependent on external capital throughout the projected period. Figure 13 summarizes the projected trajectories.

⁵⁹ [Fortune](#), "OpenAI says it plans to report stunning annual losses through 2028—and then turn wildly profitable just two years later" (12 November 2025), reporting financial documents obtained by [The Wall Street Journal](#) — projected \$9B net loss on \$13B revenue in 2025; ~\$74B operating loss in 2028 (~three-quarters of revenue); profitability from 2030.

⁶⁰ On the upward revision: Fortune/Bloomberg, "OpenAI forecasts its revenue will top \$280 billion in 2030" (20 February 2026; CNBC first to report the revenue estimate); and Wall Street Journal reporting (Berber Jin, April 2026), as summarized in [Sherwood News](#), "Report: OpenAI on track to burn \$85 billion in 2028, expects profitability by 2030" (6 April 2026) — 2028 loss revised to ~\$85B; 2028 compute spend ~\$121B.

⁶¹ [Cybernews](#), "OpenAI burns cash but Anthropic is about to break even soon" (18 November 2025), reporting the WSJ figures — Anthropic projected cash burn ~33% of revenue in 2026, ~9% in 2027, break-even in 2028




Figure 13 - OpenAI and Anthropic projected financial trajectories

Inference cost trajectory

The principal mechanism through which the cycle is, in principle, expected to bend the cost curve is the rapid decline in the cost of running a trained model at a given level of performance. Stanford's AI Index 2025 reports that the cost of inference at GPT-3.5-level performance fell by a factor of approximately 280 between November 2022 (approximately \$20 per M tokens) and October 2024 (approximately \$0.07 per M tokens), a magnitude with few comparable precedents in modern technology pricing⁶². The same source reports that hardware costs have declined at approximately 30% annually, and that energy efficiency has improved at approximately 40% annually. Figure 14 presents the cost trajectory on a logarithmic scale.




Figure 14 - Inference cost reduction over time (as \$ per M tokens)

The implication is bidirectional: falling inference costs open previously unprofitable use cases, expanding the addressable market, but they also reduce revenue captured per query, so the margin available to repay fixed infrastructure costs compresses unless volume grows at least as fast as price falls. The relevant question is therefore not whether inference costs are falling — they unambiguously are — but whether total revenue

⁶² [Stanford HAI](#), The 2025 AI Index Report — inference cost dropped 280-fold for GPT-3.5-equivalent performance between November 2022 and October 2024; hardware costs declining ~30 % annually; energy efficiency improving ~40 % annually; organizational AI use 55 → 78 % (2023→24).

(price × volume) is consistent with the financing requirements of §3.2.1. The pure-play-firm financials of the preceding paragraph indicate that at the level of the two principal model developers, the answer is at best contingent on continued external capital provision through 2030. Inference-cost decline of the observed magnitude is a necessary but not sufficient condition for the cycle to clear on its own resources.

Enterprise adoption and productivity

The end-user demand pull required to validate the capital expenditure of §3.2.1 must, in the medium term, manifest as enterprise adoption and demonstrable productivity gains. Two distinct sets of survey evidence document the adoption trajectory. Stanford's AI Index 2025 reports that organizational AI use among surveyed firms rose from 55% in 2023 to 78% in 2024, and generative-AI-specific use from 33 to 71% over the same single year⁶³. McKinsey's 2025 State of AI survey, covering more than 1,500 organisations, places the figures higher: 88% of organisations regularly use AI in at least one function, and 79% use generative AI⁶⁴. Panel data on US firms' paid AI subscriptions (covering more than 45,000 businesses) report paid-AI penetration rising from approximately 5% in early 2023 to approximately 44% in late 2025, with average annual contract values rising from approximately \$39,000 to approximately \$530,000 over the same period.

The translation of reported adoption into measured productivity remains, however, partial. Brynjolfsson, Li and Raymond analysis of 5,172 customer-support agents at a US software firm, published in the *Quarterly Journal of Economics* (May 2025), documents a 14–15% average increase in issues resolved per hour following the staggered introduction of generative-AI assistance⁶⁵. The same study reports a 34% productivity gain among novice workers and a minimal to slightly negative effect among the most experienced. At the aggregate level the picture is less clear: the macroeconomic productivity estimates diverge by an order of magnitude (the Acemoglu–Goldman split discussed in §3.2.1). The reconciliation of these estimates depends on judgements about the share of tasks that will be cost-effectively

⁶³ [Stanford HAI](#), The 2025 AI Index Report.

⁶⁴ [McKinsey & Company](#), The State of AI in 2025

⁶⁵ Erik Brynjolfsson, Danielle Li and Lindsey R. Raymond, "Generative AI at Work", [Quarterly Journal of Economics](#) 140(2), 2025: 889–942 (NBER Working Paper 31161) — 14 % average productivity gain in customer support; 34 % for novices.

automated, the elasticity of demand for automated services, and the rate at which complementary investments in organizational adaptation occur.

Three observations follow for the sustainability question: inference economics are improving fast enough to widen the market but not to self-finance the foundation-model layer; enterprise adoption is mainstreaming, yet aggregate productivity through 2025 still lags the rates the capital-expenditure forecasts imply; and the most rigorous micro evidence shows gains that are heterogeneous across skill and task. Taken together, this means that the cycle is economically sustainable on a horizon of approximately five to ten years if and only if enterprise adoption translates into commensurate revenue at the application and incumbent-enterprise-software layers which, in turn, depends on whether the productivity gains documented at the micro level scale to the macro level.

3.3 Vendor Financing in the AI Capital Cycle

The parallel between the AI expansion and the dot-com bust has one of its strongest claims is the partial circularity of the cash flows that nominally validate the expansion. Over 2024 and 2025 the largest cloud-platform operators took equity stakes or signed multi-year preferred-supplier agreements with the model developers, whose principal cost item is cloud compute purchased from those same operators; chip vendors in turn invested in both: a web of reciprocal equity-and-procurement commitments documented in detail in §3.3.2. The transactions are real and the goods delivered; the arrangements are not, strictly, accounting irregularities: for this reason, the dynamic described in the following section is not referred to as round-tripping⁶⁶. Their economic effect, however, is that a portion of the revenue recorded at each node originates as capital injected by an adjacent node, so that sectoral revenue growth is not independent of the rate of capital deployment by the same firms. Two consequences follow: the growth metrics on which conventional valuation rests become partly endogenous to the investment programme they are meant to justify, producing a feedback loop in which deployment generates apparent demand that justifies further deployment irrespective of external end-user demand; and the risk of a demand shortfall is borne unevenly, concentrated in the platform and application layers whose values rest on growth rates that may not survive a

⁶⁶ “Round-tripping” (also termed circular trading or, colloquially, “Lazy Susans”) denotes reciprocal sales in which a company sells an asset while agreeing to repurchase it at substantially the same price, producing reported revenue with no net economic substance — a violation of the substance-over-form principle. The label was popularised in reporting on the energy-trading and telecom scandals of 2000–2002 (notably Enron and AOL Time Warner); The Wall Street Journal’s characterisation of the practice is the source most often credited with fixing the term in general usage.

contraction in upstream spending. Because these dependencies are invisible to price-based diagnostics applied in isolation, the value chain must first be mapped. This chapter builds that map; the empirical analysis that follows interrogates the price series against it.

The evidence is assembled in sequence: the conceptual framework for circular and vendor-financed structures, anchored in the 1999–2001 telecommunications precedent (§3.3.1); the empirical record of commercial flows between principal participants (§3.3.2); and the leverage and risk-transfer instruments funding the buildout (§3.3.3).

3.3.1 Conceptual definitions and the telecommunications precedent

Vendor financing, the practice by which a supplier extends credit or capital to its customer to enable the customer to purchase the supplier's products, is a recurrent feature of capital-goods cycles. It serves a coordinative function during the initial growth phase of an industry: a supplier with available capacity provides funding to customers whose own access to capital is insufficient, or whose creditworthiness has not yet been established by external lenders. When the cycle is sustained by genuine end-user demand, vendor-financed receivables are subsequently repaid out of customer revenue. When the cycle is not sustained, hence when end-user demand fails to materialize at the level anticipated, vendor-financed receivables become bad debts on the supplier's balance sheet, and the supplier's revenue contracts simultaneously through the loss of the customer. The revenue-recognition standards (IFRS 15 / ASC 606), reformed after the post-2001 cycle, require revenue to reflect the substance of a customer contract and impose additional disclosure where related parties transact, but do not exclude supplier-funded purchases from recognition. The analytical implication is that reported revenue figures in the AI cycle cannot be interpreted at face value where the supplier of the relevant good or service is simultaneously providing capital to the customer that funds the purchase; requiring examination of the substance of the transaction rather than its legal form.

J. G. Sidak's analysis of the WorldCom fraud and the broader collapse of American telecommunications after deregulation, published in the *Yale Journal on Regulation* in 2003 is a rigorous treatment of vendor financing as a systemic phenomenon. Sidak's framework distinguishes between vendor financing as a means of bridging short-term liquidity gaps in customer firms (an analytically defensible practice) and vendor financing as a means of artificially sustaining demand for the supplier's products at quantities that exceed underlying end-user demand (a practice that, in the limit, produces the cyclical instability documented in the 1999–2001 telecommunications crash). The diagnostic criterion is the relationship between

vendor-financing magnitudes and the underlying revenue of the customers being financed: when the former approaches or exceeds the latter, the analytical interpretation of the supplier's recognized revenue requires explicit consideration of the self-funding component.

The principal telecommunications-equipment suppliers had provided vendor financing in magnitudes consistent with the second interpretation by the peak of the cycle. By fiscal-year-end 2000 the principal equipment suppliers Lucent, Nortel, and Cisco each had extended customer vendor-financing commitments running into the billions, with Lucent the most exposed. When the cycle reversed, the consequences were of corresponding magnitude: Lucent's revenue more than halved in two years, it swung to a multi-billion-dollar net loss, and its equity was all but wiped out⁶⁷. A large share of the vendor-financed receivables became bad debt, and Nortel was subsequently charged by the SEC with accounting fraud spanning fiscal 2000–2003⁶⁸; the company entered bankruptcy in 2009. Direct comparative analyses of the telecommunications vendor-financing structure to the present AI cycle have been advanced by several practitioner-oriented analysts⁶⁹. The magnitudes of the telecom precedent are summarized in Figure 15.




Figure 15 - Vendor Financing during the dot-com boom

⁶⁷ [American Affairs Journal](#), "Who Lost Lucent?: The Decline of America's Telecom Equipment Industry" (Summer 2020) — Lucent revenue declined from \$30 B (1999) to \$12 B (2001); 2001 net loss \$16.1 B.

⁶⁸ [U.S. Securities and Exchange Commission](#), "Nortel Networks Pays \$35M to Settle Financial Fraud Charges" (Press release 2007-217, 15 October 2007) — civil fraud charges relating to fiscal years 2000-2003; revenue overstatement of approximately \$1.4 B in Q4 2000.

⁶⁹ [Tomasz Tunguz](#), "Circular Financing: Does Nvidia's \$110 B Bet Echo the Telecom Bubble?" (2025) — comparative analysis of Lucent (\$8.1 B), Nortel (\$3.1 B / \$1.4 B outstanding) and Cisco (\$2.4 B) vendor-financing magnitudes against current AI-era commitments.

Three diagnostic features of the telecommunications cycle are relevant for the analysis of the AI cycle. First, the supplier of capital goods to a growing industry is simultaneously the principal financier of that industry's customers, creating a feedback loop in which customer purchases and supplier revenue are not independent. Second, the revenue recognized by the supplier on customer purchases is, in part, funded by the supplier's own capital provision, so that the supplier's reported revenue overstates the underlying end-user demand. Third, when end-user demand fails to validate the cycle's growth, the resulting losses are concentrated at the supplier, whose simultaneous exposure as creditor and as revenue-recipient transmits the customer's distress to the supplier's solvency. These three concepts define the lens through which the empirical evidence of §3.3.2 and §3.3.3 is examined.

3.3.2 Evidence on circular flows in the AI industry

The Figure 16 - "A guide to the Circular Deals Underpinning the AI Boom" by Bloombers News by Bloomberg, which has circulated widely in recent commentary on the AI boom, has become a viral visual shorthand for the dense network of financial and strategic ties now shaping the sector.





Figure 16 - "A guide to the Circular Deals Underpinning the AI Boom" by Bloombers News

To simplify the analysis and put it in parallel to the value-chain layers identified in §3.1, a small record of commercial flows between principal AI-cycle participants from 2023 to early 2026 is documented in Table 3 - Principal announced commercial flows between AI-cycle participants, 2023–2026., which lists a set of the largest announced transactions. A more schematic and layer-aggregated diagram (Figure 17 **Error! Reference source not found.**) was also built to facilitate the understanding of the graph and abstract the transactions to the value-chain. The relevant feature for the present analysis is not the magnitude of any individual transaction but the pattern of inter-firm flows that the matrix describes. Six categories of flow are identified and illustrated in Figure 17.

Table 3 - Principal announced commercial flows between AI-cycle participants, 2023–2026.

<i>Parties (flow direction)</i>	<i>Flow type</i>	<i>Announced magnitude</i>	<i>Date</i>
<i>Nvidia → OpenAI</i>	Equity investment (chips paid in cash; equity for non-controlling shares)	Up to \$100bn (intent; tranches "one at a time")	Sep 2025
<i>Nvidia → Anthropic</i>	Equity investment + GPU purchase	Up to \$10bn equity	Nov 2025
<i>OpenAI → Oracle</i>	Multi-year compute commitment	\$300bn / 5yrs (4.5 GW)	Sep 2025
<i>OpenAI → Microsoft</i>	Incremental Azure capacity	~\$250bn	Oct 2025
<i>OpenAI → Amazon</i>	Compute commitment (7yr)	\$38bn	Nov 2025
<i>Anthropic → Microsoft</i>	Cloud commitment + \$5bn MSFT equity in Anthropic	\$30bn	Nov 2025
<i>AMD → OpenAI</i>	Customer warrant (160M sh @ \$0.01) vs. purchase milestones	~\$90bn hardware-revenue potential (6 GW)	Oct 2025
<i>AMD → Meta</i>	Customer warrant (160M sh), MI450 inference	6 GW	Feb 2026
<i>Nvidia → neoclouds (CoreWeave, Crusoe, Lambda, Nebius, Nscale)</i>	Minority equity + GPU-allocation priority	Minority stakes	2023–ongoing
<i>Stargate consortium (OpenAI–Oracle–SoftBank)</i>	Infrastructure target (aggregates above)	\$500bn headline	Jan 2025

Circular financial flows in the AI value chain (selected major flows, 2024–2026)



Sources: company press releases (Nvidia, AMD, Oracle, Microsoft, Amazon, OpenAI, Anthropic, CoreWeave); SEC filings (10-K, 10-Q, 6-K); Wall Street Journal; Reuters; The Information.

Figure 17 - Circular Financing by layer of the AI value-chain, author's synthesis

(a) Chip-vendor equity flows to the model developer:

The most direct expression of this category is the September 2025 letter of intent between Nvidia and OpenAI, under which Nvidia "intends to invest up to \$100B in OpenAI" as 10 GW of Nvidia systems are progressively deployed for OpenAI's infrastructure⁷⁰. The structure has two intertwined transactions: OpenAI pays cash to Nvidia for chips, and Nvidia invests in OpenAI for non-controlling shares; the first \$10B of Nvidia's equity investment is reportedly to be released upon definitive agreement for the chip purchase⁷¹. Subsequent statements by Nvidia CEO Jensen Huang and CFO Colette Kress (December 2025) clarified that the \$100B figure "was never a commitment" and that Nvidia would consider any funding tranches "one at a time"; Nvidia's separately disclosed \$500B of chip bookings through 2026 does not include

⁷⁰ Nvidia Corporation and OpenAI joint press release, "[OpenAI and NVIDIA Announce Strategic Partnership to Deploy 10 Gigawatts of NVIDIA Systems](#)" (22 September 2025).

⁷¹ [CNN Business](#), "Nvidia to invest up to \$100B in OpenAI" (22 September 2025) — two intertwined transactions: OpenAI pays cash for chips, Nvidia invests for non-controlling shares; first \$10 B begins on definitive purchase agreement.

OpenAI orders⁷². A parallel structure was announced for Nvidia's relationship with Anthropic in November 2025: Nvidia committed up to \$10B of equity investment, paired with substantial GPU compute purchases by Anthropic from Nvidia's hardware roadmap.

(b) End-customer compute commitments to cloud providers

OpenAI has, since mid-2024, accumulated a series of multi-year compute-purchase obligations to cloud providers whose announced cumulative magnitude is on the order of \$1.4T over the eight-year period 2025–2032⁷³. The largest individual element is the September 2025 announcement that OpenAI will purchase \$300B of compute from Oracle over five years from 2027 (approximately \$60B per year, 4.5 GW of capacity), reportedly the largest cloud-computing contract in history⁷⁴. Microsoft committed approximately \$250B in incremental Azure capacity to OpenAI in October 2025 (over and above the existing arrangement under which Microsoft was OpenAI's principal cloud provider); Amazon signed a \$38B seven-year compute commitment with OpenAI in November 2025. These commitments are aggregated within the OpenAI–Oracle–SoftBank "Stargate" consortium announced in January 2025 with a headline \$500B infrastructure target, of which the Abilene, Texas flagship (1.2 GW) is reported as live and operational⁷⁵. Anthropic announced, also in November 2025, a \$30B multi-year cloud commitment with Microsoft, paired with \$5B of Microsoft equity investment and \$10B of Nvidia equity investment in Anthropic.

(c) Cloud providers as the principal source of chip-vendor revenue

The third edge of the circular flow is the empirical pattern documented in §3.1.3: Nvidia's Data Center revenue, which is structurally concentrated on hyperscale cloud providers, is the revenue base that is being substantially recapitalized by the equity flows of category (a) and the compute commitments of category (b). The closing of the loop is the relevant feature for the present analysis: chip-vendor equity capital invested in OpenAI funds OpenAI's compute commitments

⁷² [Bloomberg/Yahoo Finance](#), "Pledge to Invest \$100B in OpenAI Was 'Never a Commitment,' Says Nvidia's Huang" (December 2025) — UBS Global Technology and AI Conference commentary by CFO Colette Kress; Nvidia's \$500 B disclosed chip bookings do not include OpenAI orders.

⁷³ [Tomasz Tunguz](#), "OpenAI's \$1T Infrastructure Spend" (November 2025) — aggregates announced OpenAI compute commitments: Oracle \$300 B, Microsoft \$250 B incremental, AMD ~\$90 B, Amazon \$38 B, Nvidia (up to) \$100 B.

⁷⁴ [Data Center Dynamics](#), "OpenAI signs \$300 bn cloud deal with Oracle - report" — citing Wall Street Journal (September 2025); 4.5 GW data-centre capacity; \$60 B per year from 2027; part of Stargate consortium.

⁷⁵ [IntuitionLabs](#), "Oracle & OpenAI's \$300 B Deal: AI Infrastructure Analysis" (December 2025, updated April 2026) — Stargate Abilene 1.2 GW facility live; ~7 GW Stargate planned capacity; >\$400 B in investment; Oracle backlog ~\$523 B.

to Oracle and to Microsoft; Oracle and Microsoft purchase chips from the chip vendor to fulfil those commitments; the chip vendor recognizes revenue on those chip sales. The economic substance is that the chip vendor is partially funding its own revenue.

(f) Customer warrants in exchange for purchase commitments

The most explicit case in which a chip vendor's equity is exchanged for purchase commitments is AMD's October 2025 agreement with OpenAI: AMD issued OpenAI a warrant for up to 160M AMD common shares (approximately 10% of AMD outstanding shares) at an exercise price of \$0.01 per share, with vesting milestones tied both to deployment progress (the first tranche on the initial 1 GW deployment, scaling to a 6 GW total) and to AMD share-price targets (up to approximately \$600 per share). The announcement explicitly characterises the agreement as having approximately \$90B of "cumulative hardware revenue potential"⁷⁶. An identical structure was replicated in AMD's February 2026 agreement with Meta — also 6 GW, also a warrant for 160M shares covering inference workloads on MI450-based custom chips⁷⁷. The economic substance of these arrangements is that the customer receives a contingent equity claim in the supplier whose magnitude depends on the customer's purchase activity, a configuration that aligns the parties' incentives in a manner that may not be neutral with respect to the question of underlying end-user demand.

(d) Chip-vendor minority equity stakes in cloud subcontractors

Nvidia has, since 2023, accumulated minority equity stakes in a number of specialised cloud providers (CoreWeave, Crusoe, Lambda, Nebius, and Nscale among them) paired with GPU-allocation priority that has enabled these firms to scale at a pace that would have been impractical absent the priority allocation. CoreWeave's Form 10-K analysis (cited in §3.1.3) documents that 67% of its FY 2025 revenue (and 70% of its nine-month revenue) came from Microsoft, with the top three customers (Microsoft, OpenAI, Meta) representing the substantial majority of the \$60.7B remaining-performance-obligation balance. The structural significance of these arrangements is that the neoclouds are simultaneously customers of the chip vendor

⁷⁶ [AMD Investor Relations](#) press release, "AMD and OpenAI Announce Strategic Partnership to Deploy 6 Gigawatts of AMD GPUs" (6 October 2025) — warrant for up to 160 M shares (~10 % of AMD outstanding) at \$0.01 strike, milestone-vesting; first tranche on initial 1 GW deployment.

⁷⁷ [The Decoder](#), "AMD basically copy-pasted its OpenAI deal for Meta, six gigawatts and ten percent equity included" (24 February 2026).

(purchasing GPUs) and suppliers to the model developers (selling cloud capacity), with the chip vendor holding minority equity in both ends of the pipeline.

Across the matrix, two observations are necessary. First, no single transaction in Figure 17 is, in isolation, an example of vendor financing in the strict sense documented for telecommunications in 1999–2001: no chip vendor has, to date, extended direct credit to its customer to enable a chip purchase that the customer could not otherwise finance. Second, however, the aggregate pattern of equity investments, warrants, multi-year compute commitments, and minority cross-holdings does produce a structural configuration in which the same parties appear repeatedly on both sides of the principal commercial flows of the cycle. The empirical question of whether this configuration is materially different from vendor financing, or is a contemporary instantiation of it, depends on the analytical weight given to the equity-investment leg of the bilateral transactions and to the inter-firm dependence that the aggregate matrix produces. The §3.3.1 suggests that the relevant test is whether the chip-vendor revenue recognized on customer purchases is, in substantive economic terms, independent of the chip-vendor's own capital provision to those customers; the present subsection establishes that, in the aggregate, this independence cannot be presumed.

3.3.3 Evidence on leverage and risk transfer

The financial structure of the cycle is further characterized by a sharp reported increase in the rate at which the principal hyperscaler firms are raising external debt against the projected revenue streams of the AI buildout. According to Bank of America Securities, the Big Five US hyperscaler firms (Amazon, Alphabet, Meta, Microsoft, Oracle) issued approximately \$121B of US investment-grade corporate bonds in 2025, against an average of approximately \$28B per year over the 2020–2024 period⁷⁸, 90B of which concentrated in the final three months of the year⁷⁹. Specific issuances and their trajectory are summarized in Figure 18. Forecasts for 2026 vary by methodology: BofA Global Research revised its hyperscaler-specific forecast upward from \$140B to \$175B following Amazon's deal⁸⁰; JPMorgan analysts estimate that AI-related infrastructure could drive on the order of \$1.5T of additional investment-grade bond

⁷⁸ [Bank of America Securities](#), January 2026 report on hyperscaler corporate-bond issuance

⁷⁹ [Mellon Investments Corporation](#), "Record-Breaking AI-Related Debt Issuance in 2025" (December 2025) — quantifies \$121 B 2025 issuance with \$90 B+ in the final three months alone.

⁸⁰ [Reuters](#), "Analysts revise AI hyperscaler debt forecasts after Amazon bond sale" (17 March 2026) — Amazon's near-record \$54 B investment-grade bond sale; BofA Global Research revised 2026 hyperscaler issuance forecast to \$175 B (from \$140 B); Barclays projects total US IG issuance >\$2 T in 2026.

financing over roughly five years; UBS raised its 2026 forecast for US investment-grade technology bond issuance to about \$360B (from \$300B)⁸¹.




Figure 18 - Hyperscaler corporate bond issuance and forecast

Microsoft is the principal exception to the pattern. Mellon Investments' December 2025 commentary notes that Microsoft *"remains the only hyperscaler not tapping debt markets recently"* for direct corporate-bond issuance, with the consequence that Microsoft's AI infrastructure exposure is reflected through joint-venture and partnership structures rather than through direct balance-sheet leverage: Microsoft–BlackRock AI infrastructure partnership and its participation in the October 2025 \$40B Aligned Data Centers acquisition is the largest data-centre transaction on record.

A second mechanism for funding the AI buildout has been the rapid expansion of off-balance-sheet financing structures, in which special-purpose vehicles (SPVs) raise debt against AI infrastructure assets that are subsequently leased back to the hyperscaler that effectively controls them. Industry analysis estimates that approximately \$120B of AI-infrastructure debt had been shifted off the principal hyperscaler balance sheets by year-end 2025 through such structures⁸². The mechanism is consistent across the principal participants: the most prominent case is Meta's October 2025 "Hyperion" agreement with Blue Owl, in which roughly \$27B of

⁸¹ [CNBC](#), "Dust to data centers: The year AI tech giants, andB s in debt, began remaking the American landscape" (31 December 2025) — Morgan Stanley and JPMorgan estimate AI infrastructure could drive up to \$1.5 T in additional tech-sector borrowing; UBS analysts forecast as much as \$900 B in new issuance in 2026 alone.

⁸² [AI Daily News](#), "Tech Giants Shift \$120B in AI Debt Off Balance Sheets" (December 2025) — Meta-Blue Owl Hyperion SPV "Beignet Investor" \$30 B; Oracle SPV deals totalling \$69 B (Abilene \$13 B / Texas-Wisconsin \$38 B / New Mexico \$18 B); xAI and CoreWeave participation.

SPV debt funds a data centre that Meta then leases back with the debt held off its consolidated balance sheet, and Oracle, xAI, and CoreWeave have adopted comparable structures across the Stargate buildout, financed by a recurring group of private-credit providers (Blue Owl, JPMorgan, Pimco, BlackRock, Apollo). The analytically relevant feature is the structure (debt raised against AI assets and kept off the controlling firm's balance sheet) not the magnitude of any single facility.

FINANCING STRUCTURE	MAIN EXAMPLES
SPV LEASE-BACK STRUCTURES	Meta–Blue Owl Capital (Oct 2025): \$30B Hyperion data center in Louisiana (4M ft ²), SPV “Beignet Investor.” Oracle Stargate SPVs: \$13B Abilene, TX (of which \$10B debt), \$38B Texas + Wisconsin, \$18B New Mexico.xAI and CoreWeave also using similar structures.
JOINT-VENTURE STRUCTURES	Microsoft–BlackRock–MGX AI infrastructure partnership. \$40B acquisition of Aligned Data Centers (Oct 2025).
GPU-COLLATERALISED PRIVATE CREDIT	CoreWeave debt facilities: Aug 2023: \$2.3B (Magnetar + Blackstone), with Nvidia H100 chips as collateral.May 2024: \$7.5B (Blackstone + Magnetar + Coatue, Carlyle, CDPQ, DigitalBridge Credit, BlackRock, PIMCO).Dec 2025: \$2.25B convertible at 1.75% coupon.

A third mechanism, particularly significant for the neoclouds, is direct private-credit financing collateralised by GPU hardware. CoreWeave's August 2023 facility \$2.3B led by Magnetar Capital and Blackstone Tactical Opportunities was the first major transaction in which Nvidia H100 chips were used as collateral for a debt facility⁸³. This was followed in May 2024 by a \$7.5B debt facility led by Blackstone and Magnetar (with strategic participation from Coatue, Carlyle, CDPQ, DigitalBridge Credit, BlackRock, Eldridge Industries, and Great Elm Capital Corp), described by Blackstone as "one of the largest private credit financings in history"⁸⁴. By year-end 2025, CoreWeave's total indebtedness had reached approximately \$14.2B, including a December 2025 \$2.25B convertible bond at a 1.75% coupon level made possible by the equity-upside component of the convertible structure⁸⁵. The structural significance of GPU-collateralised lending is twofold: it permits asset-light cloud providers to scale at a pace that would be incompatible with their underlying equity base, and it transfers the credit risk of those

⁸³ [Blackstone press release](#), "CoreWeave Secures \$2.3B Debt Financing Facility led by Magnetar Capital and Blackstone" (3 August 2023) — debt collateralised by Nvidia H100 chips; participants include Coatue, DigitalBridge Credit, BlackRock, PIMCO, Carlyle.

⁸⁴ [Blackstone press release](#), "CoreWeave Secures \$7.5B Debt Financing Facility led by Blackstone and Magnetar" (17 May 2024) — "one of the largest private credit financings in history";

⁸⁵ [Energy Now](#) / Bloomberg, "The \$3T AI Data Center Build-Out Becomes All-Consuming For Debt Markets" (February 2026) — JPMorgan projects approximately \$300 B per year of AI- and data-centre-related deals; Oracle CDS spreads widened; CoreWeave December 2025 convertible bond at 1.75 % coupon.

providers from their corporate balance sheets to the private-credit funds that finance them, with the consequence that the systemic risk of the buildout is being progressively redistributed from listed equity into private credit. The same investor names (Blackstone, Magnetar, Pimco, BlackRock, Carlyle, Apollo) appear repeatedly across the principal off-balance-sheet and private-credit structures, with the consequence that the diversification ordinarily available to the buy-side in a multi-sector private-credit allocation is, in this cycle, limited by the concentration of AI-related exposure within a small number of large alternative-asset managers.

The aggregate financial-structure picture across the three mechanisms is that the cycle is being financed by a combination of: balance-sheet bond issuance (concentrated among Alphabet, Amazon, Meta, and Oracle, with cumulative 2025–2026 magnitude in the range of \$200–300B); off-balance-sheet SPV debt (cumulative approximately \$120B at end-2025, projected to grow substantially); and GPU-collateralised private credit (most prominently CoreWeave’s \$14B -plus, with similar structures at multiple smaller neoclouds). These flows sit within a far larger capital programme. Goldman Sachs equity research projects \$1.15T of cumulative hyperscaler capex over 2025–2027 (§3.2.1); on a longer horizon, JPMorgan estimates roughly \$300B per year of AI- and data-centre-related deal activity, implying cumulative deal volume approaching \$1.5T over five years. Credit-default-swap spreads on Oracle widened materially at the end of 2025, with the cost and volume of insurance contracts against Oracle debt rising in close proximity to the announcement of the \$300B OpenAI commitment, reflecting market assessment of the supplier most exposed to AI-revenue-recognition risk among the investment-grade hyperscalers.

3.3.4 Synthesis: the structural pattern and the empirical hypothesis

Three features emerge from the evidence of §3.3.2 and §3.3.3. First, the equity-investment and compute-commitment flows between the principal participants in the AI cycle (Nvidia and AMD on the chip-vendor side; OpenAI and Anthropic on the model-developer side; Microsoft, Oracle, Amazon, Google, and the neoclouds on the cloud-provider side) produce a configuration in which the same parties appear repeatedly across the matrix, with the consequence that the demand for the chip vendor’s product is, in substantive economic terms, partially funded by the chip vendor itself. Second, the financing of the buildout has shifted decisively from operating cash flow to external capital, with hyperscaler corporate debt issuance running at approximately four times the pre-cycle annual rate and a further \$120B -plus of debt held off the principal balance sheets through SPV structures. Third, the credit risk

of the asset-light cloud providers that participate most directly in serving the model developers has been progressively transferred to private credit, where the same investors (Blackstone, Magnetar, Pimco, BlackRock, Carlyle, Apollo) appear repeatedly as financiers of the lower tiers of the value chain. The same firms, in short, appear simultaneously as supplier, customer, equity holder, and creditor; the diversification ordinarily available to participants in a complex value chain is, in this cycle, structurally limited.

The comparison with the telecommunications precedent of 1999–2001 is instructive both for what is similar and for what is different. The similarities are: the supplier of capital goods to a rapidly growing industry is also a significant capital provider to that industry’s customers; end-user revenue from the supplier’s products has not yet, at the time of the principal capital commitments, been validated at the magnitude that the commitments imply; and the cycle’s financial logic depends on the eventual realization of revenue at multiples of the level observed at the time of analysis (in the AI cycle, the revenue-gap calculation of §3.2.1 implies a gap on the order of \$945B between observed and required revenue, applying the Sequoia methodology of Cahn (2024) to Nvidia’s Q4 fiscal 2026 run rate). The differences are: no direct customer-loan vendor financing has been documented in the AI cycle to date (the chip-vendor capital provision is through equity investment and warrants rather than receivable extensions); the principal model developers are private rather than public companies, with the consequence that the public-market exposure of the cycle is concentrated in the chip-vendor, cloud-provider, and neocloud equity values rather than in the model-developer valuations; and the scale of capital expenditure being committed is of a magnitude (approximately \$5–7T cumulative to 2030 in the McKinsey range) that has no comparable precedent in modern industrial history. The aggregate effect is that the cycle exhibits structural features that are recognizably continuous with the telecommunications precedent (the cross-holding of equity between supplier and customer, the off-balance-sheet treatment of buildout debt, the dependence on continued external capital provision to maintain growth) while also exhibiting features that are unprecedented in scale.

The combination of the demand-magnitude question (§3.2.1: the revenue gap discussed above), the sustainability question (§3.2.3: the principal model developers project cumulative pre-profitability cash burn of approximately \$115B through 2029), the concentration-risk question (§3.2.4: the principal customers and the principal counterparties are interconnected through equity holdings, off-take agreements, and shared supply chains), and the financial-structure question (the present section: an unprecedented increase in external debt issuance and a

documented pattern of circular commercial flows) describes a cycle whose economic features collectively warrant the application of formal econometric tests for asset-price explosivity.

The documentation of the present chapter provides the conceptual framework against which the test results will be interpreted: an econometric finding of explosivity in AI-related asset prices would, conditional on the structural evidence of this chapter, support a characterization of the AI capital cycle as exhibiting the features of a financial bubble; an econometric finding of no explosivity, would inform the absence of inflated valuation in the public markets, but would not invalidate the trigger warnings derived from this industry analysis.

4. Research Design: Hypothesis and Methodology

This chapter operationalizes the research question outlined in the first chapter of this thesis. whether the price dynamics of AI-exposed equities exhibit the statistical signature of speculative bubbles, and how any such behavior is distributed across the layers of the AI value chain. It bridges the descriptive account of the AI industry developed in Chapter 3 and the survey of bubble-detection methods presented in Chapter 2 by applying a battery of validated econometric tests to a purpose-built panel of AI-related securities. The chapter is organised in four parts. §4.1 states the empirical research question, its sub-questions, and the hypotheses to be tested. §4.2 describes the construction of the AI-exposed universe, the commodity-market extension, the price data, and the preprocessing steps. §4.3 sets out the methodology, including the full battery of unit-root and explosive-root tests, the Monte Carlo procedure used to obtain critical values, the volatility-robust SV-ADF test, the episode-dating rules, the historical validation design, and the computational pipeline built to ensure reproducibility.

4.1 Research question and hypotheses

4.1.1 Research question and sub-questions

The Introduction frames the thesis around one research question, which can be stated in operational form as follows:

To what extent do AI-related equities exhibit econometrically identifiable explosive price dynamics that are consistent with speculative bubble formation and do those dynamics appear along the AI value chain in a pattern that is robust to alternative detection methodologies?

This formulation deliberately follows the econometric definition of a bubble adopted in §2.1: a bubble is treated as a non-stationary, explosive deviation of prices from a random-walk benchmark, identified through the statistical properties of the price series rather than through an estimate of fundamental value. The main question is decomposed into six sub-questions, each of which corresponds to a distinct part of the empirical design:

- **SQ1 — Explosive behavior in AI-related equity prices:** Which ones, if any, of the AI-exposed equities exhibit statistically significant explosive episodes in the period following the public release of ChatGPT in November 2022?

- **SQ2 — Indices versus individual securities.** Do broad AI-tracking indices and sector benchmarks display explosive episodes, or is explosive behavior confined to individual firms?
- **SQ3 — Value-chain heterogeneity.** How do the incidence, timing, and duration of explosive episodes vary across the value-chain segments?
- **SQ4 — Method agreement.** Do the SADF, GSADF, and SV-ADF procedures agree on the existence and dating of episodes, and what does any disagreement reveal about the sensitivity of each method to specification choices?
- **SQ5 — Propagation to real-resource markets.** Does explosive behavior also appear in the commodity markets that supply AI infrastructure, notably the energy and industrial-metal markets?
- **SQ6 — Validation.** When the same test battery is applied to historical data, does it correctly recover well-documented past bubbles, in particular the dot-com episode and the pre-2008 cycle?

4.1.2 Hypotheses

The empirical analysis is structured around six hypotheses. Each is grounded in the literature reviewed in Chapter 2 or in the features of the AI industry documented in Chapter 3, and each is falsifiable with the data and methods described in the remainder of this chapter.

- **H1 (Post-ChatGPT explosiveness):** At least a subset of AI-exposed equities exhibit statistically significant explosive episodes in the period following the public release of ChatGPT in November 2022. This hypothesis follows directly from Basele, Phillips and Shi (2025) and Maghyereh and Awartani (2025), who already report bubble-consistent episodes in AI-adjacent equity segments.

Operational test: GSADF applied to each equity over 2019–2026, with the BSADF date-stamping algorithm and the 21-day minimum-duration filter.

Decision rule: supported when the GSADF statistic exceeds the Monte Carlo 95% critical value and the BSADF dates at least one episode whose origination lies after 1 November 2022.

- **H2 (Value-chain heterogeneity):** Explosive episodes are concentrated in the upstream, supply-constrained segments of the value chain (compute, foundry, and memory) rather than being distributed uniformly across all layers. The rationale, developed in Chapter 3, is that capacity in these segments is inelastic in the short run, so that revisions to AI demand expectations are capitalized more directly and more abruptly into prices than in the more diversified hyperscaler layer.

Operational test: segment-level episode counts from the BSADF dating procedure.

Decision rule: supported when the per-firm episode rate in {compute, foundry, memory} strictly exceeds that of {hyperscalers, REITs, software platforms}.

- **H3 (Index versus single-stock dilution):** Broad indices exhibit fewer and weaker explosive episodes than individual AI-exposed equities, because diversification across non-explosive constituents dilutes the explosive signal. This is consistent with the regime-mixing logic of Evans (1991) and with the observation that aggregate indices average over heterogeneous firm-level dynamics.

Operational test: GSADF and the SV-ADF applied separately to the Nasdaq Composite, the PHLX Semiconductor Index, and the S&P 500.

Decision rule: supported when the three indices fail to reject the GSADF null while the majority of their AI-exposed constituents reject it.

- **H4 (Method selectivity):** Detection of explosivity is highly method sensitive. Applied to the same series, the SADF null is rejected more often than the SV-ADF null, reflecting the effects of allowing for persistent time-varying volatility in the SV-ADF procedure. This hypothesis is examined directly in the cross-method comparison of §5.2.3.

Operational test: episode-detection comparison restricted to the post-ChatGPT window common to both methods.

Decision rule: supported when the SV-ADF episode count is strictly less than the SADF episode count on the common sample for each security.

- **H5 (Historical validation):** When applied to historical index data, the test battery correctly detects periods of explosivity manifested in the equity markets: the documented dot-com episode (origination in the second half of the 1990s, collapse in

2000–2001) and the bitcoin rallies of 2016–2017 and 2020–2021, confirming that the implementation detects bubbles that are independently known to have occurred.

Operational test: GSADF and SV-ADF applied to the Nasdaq Composite over 1990–2001, the S&P 500 over 1995–2001, and the S&P 500 over 2003–2009.

Decision rule: supported when the BSADF episode dates fall within the windows reported by Phillips, Wu and Yu (2011) and Phillips, Shi and Yu (2015).

- **H6 (Commodity propagation — exploratory):** If AI-infrastructure investment is propagating speculative dynamics into real-resource markets, explosive episodes should be more readily detectable in AI-relevant commodities, notably copper and energy carriers linked to data-centre power demand, than in commodities with no plausible AI linkage, such as agricultural and livestock products. Given the indirect nature of this channel, H6 is treated as exploratory rather than confirmatory.

Operational test: GSADF applied to the commodity panel and comparison of episode-date distributions between AI-relevant and placebo groups.

Decision rule: supported in conditional form when the post-Q4-2023 fraction of total episode-days is strictly higher in the AI-relevant groups than in the placebo groups; the literal strong form (“placebos display no episodes at all”) is rejected by the data and is not used as the operational test.

4.2 Data

4.2.1 Construction of the AI-exposed universe: indices, equities, commodities

The empirical analysis requires a panel of securities whose prices can plausibly be expected to reflect expectations about the AI industry. Selecting firms on the basis of recent price performance would introduce a mechanical bias towards securities that have already risen; instead, the universe is constructed by mapping a set of candidate firms resulting from the industry overview to the value-chain taxonomy developed in Chapter 3. The universe is deliberately extensible: securities can be added or removed by editing a single configuration file (config/settings.py), and the computational pipeline recomputes only the quantities affected by the change.

Each firm is assigned to exactly one of thirteen segments according to its primary economic exposure. A firm that operates across several layers, as many large technology companies do, is classified by the layer that dominates its revenue and its public investment narrative. Three criteria govern inclusion: the security must be publicly listed in the US market, and continuously traded over the sample window; it must have sufficient liquidity for daily closing prices to be economically meaningful; and it must have a clearly identifiable primary exposure to one segment of the AI value chain. The resulting universe comprises thirty-seven individual equities across the thirteen segments, organized into three layers (upstream, midstream, and downstream) summarised in Table 4 - The AI-exposed equity universe, organised by value-chain layer and segment.

Within the upstream layer, for instance, semiconductor equipment (lithography and process tools: ASML, AMAT, LRCX) is separated from chip design (NVDA, AMD, INTC, AVGO, MRVL), from high-bandwidth memory (MU, the only US-listed HBM producer; SK Hynix and Samsung are Korean-listed and excluded because they don't respected the three criteria cited above), and from leading-edge foundry (TSM, which also proxies for the advanced-packaging layer via its CoWoS technology). Segment 5, storage hardware (WDC, STX), occupies a boundary position: HDD and NAND storage are not treated in Chapter 3 as part of the core AI capital stack, so this segment functions as a within-universe control and negative reference against which the upstream-to-downstream gradient (H2) is read. Tesla (TSLA) is retained as an AI-narrative sentinel despite robotics and autonomous systems being formally outside the AI-industry scope as defined in Chapter 3 §3.1.

Table 4 - The AI-exposed equity universe, organised by value-chain layer and segment.

<i>Layer</i>	<i>Segment</i>	<i>Tickers</i>	<i>Role in the AI value chain</i>
<i>Upstream</i>	1. Semiconductor equipment	ASML, AMAT, LRCX, KLAC	EUV lithography monopoly, wafer-fab equipment
	2. Chip design (AI accelerators)	NVDA, AMD, INTC, AVGO, MRVL	GPU, accelerator, and custom-ASIC design
	3. Memory — HBM-relevant	MU	DRAM and high-bandwidth memory for training clusters
	4. Foundry and advanced packaging	TSM, ASX	Advanced-node fab + CoWoS (TSM); outsourced packaging (ASX)
	5. Storage — non-HBM control	WDC, STX	HDD/NAND storage; within-universe negative reference

Midstream	6. Data-centre networking & optical	ANET, CSCO, NOK, CIEN	DC switching, optical interconnect
	7. Data-centre REITs	EQIX, DLR	Physical colocation infrastructure
	8. Power & grid	CEG, VST	Generation tied to DC power demand
	9. Hyperscale cloud	MSFT, AMZN, GOOGL, META, ORCL	Cloud AI platforms integrating compute and model development
	10. Specialised AI cloud	CRWV, NBIS	GPU-dense neocloud capacity
Downstream	11. Enterprise AI / hybrid cloud	IBM	Enterprise AI software and hybrid-cloud incumbent
	12. Application & AI platform software	PLTR, SNOW, DDOG, AI	AI-native application platforms and observability tools
	13. Edge AI & devices	QCOM, AAPL, TSLA	On-device silicon, consumer devices, autonomy narrative

The commodity-market extension

The empirical design is extended beyond equities to a panel of commodity markets. The rationale is that the AI capital cycle does not operate exclusively through financial markets. The construction of AI infrastructure generates large and concentrated demand for physical inputs: copper for data-centre power distribution and interconnects, natural gas and other energy carriers for data-centre electricity consumption, and a range of metals used in electronic components. If speculative dynamics are propagating from AI equities into the markets for the resources that AI infrastructure consumes, then explosive behavior should become detectable in AI-relevant commodity prices. Testing this proposition connects the thesis to the broader macroeconomic debate on whether the AI build-out constitutes a genuine investment supercycle, and it provides a natural placebo test: commodities with no plausible AI linkage (agricultural and livestock products) should display no such dynamics.

The sixteen-instrument commodity panel is organized into four analytically distinct groups, as summarized in Table 5. Because continuously traded futures contracts require contract-month roll conventions that complicate the construction of consistent long price series, the panel is built primarily from exchange-traded funds and notes that track the underlying commodity markets as continuous adjusted-price series. One instrument, JJC (iShares Copper ETN), ceased

active trading in mid-2023, introducing a structural break in that series; this limitation is noted in §4.2.1 and alternative copper instruments are identified as outstanding work.

Precious metals (GLD, SLV) are treated as a separate *confound* group rather than being classified as AI-relevant metals. Gold and silver have material monetary and safe-haven drivers that operate independently of AI-infrastructure demand, making it difficult to attribute any explosive dynamics in those prices to the AI channel. They are included to allow their behaviour to be compared with genuinely AI-linked metals, but no AI-propagation hypothesis is tested on them and their results are interpreted with this caveat.

<i>Group</i>	<i>Instruments</i>	<i>Analytical role</i>
<i>AI-relevant: Energy</i>	USO, UNG, BNO	Primary test of propagation via data-centre electricity demand (H6)
<i>AI-relevant: Industrial metals</i>	JJC, PPLT, PALL	Copper and platinum-group metals used in hardware manufacturing (H6)
<i>Confound: Precious metals</i>	GLD, SLV	Gold and silver; monetary and safe-haven dynamics confound the AI-demand channel
<i>Placebo: Agriculture</i>	DBA, WEAT, CORN, SOYB	No plausible AI linkage; included as null-hypothesis controls
<i>Placebo: Livestock</i>	COW	Second placebo group
<i>Benchmark: Broad indices</i>	DBC, GSG, PDBC	Aggregate commodity-market references

Table 5- The sixteen-instrument commodity panel. The energy and industrial-metals groups carry the H6 propagation test. Precious metals are separated as a confound. Agriculture and livestock serve as placebos for which no explosive behaviour is expected.

Benchmark indices and controls

Three market indices are included as controls: the Nasdaq Composite (^IXIC), the PHLX Semiconductor Index (^SOX), and the S&P 500 (^GSPC). The indices serve two purposes. First, they allow the index-versus-individual-security comparison required by sub-question SQ1 and hypothesis H3: if explosive behavior is detectable in individual AI equities but not in the indices that contain them, this constitutes evidence that index diversification dilutes the

signal, which has material implications for how practitioners using index products are exposed to any nascent bubble. Second, they provide a reference against which the firm-level results can be read: the Nasdaq Composite is the index to which the dot-com explosive-root literature has most often been applied and constitutes the primary out-of-sample validation benchmark.

4.2.2 Price data, source, and sample window

Daily price data are obtained from Yahoo Finance through the `yfinance` interface. For every security the series used is the adjusted closing price, which incorporates corrections for stock splits and dividend distributions and therefore provides a consistent total-return-adjusted price level. The explosivity tests are applied to the adjusted price level directly. The price-dividend ratio specification favoured in parts of the classical literature is not used, for two reasons: a substantial share of the firms in the universe pay no dividend, so a price-dividend ratio would be undefined or uninformative for those securities; and the recent literature on technology-sector bubbles, for example Basele, Phillips and Shi (2025), applies explosive-root tests directly to price levels.

For the explosivity diagnostics (SADF, and GSADF) a single fixed sample window is used, running from *1 January 2020* to *1 May 2026*. This window satisfies three requirements:

1. it begins well before the AI-investment surge, so that a pre-boom baseline is included and the tests are not applied to a sample that consists only of the episode under investigation;
2. it is long enough (approximately 1,842 trading days for a continuously listed security) for the recursive procedures to have adequate power, since the Phillips-Shi-Yu minimum-window constraint consumes approximately 5% of the sample before the expanding recursion begins;
3. it includes the COVID-19 market dislocation of March 2020, which provides a demanding test of the episode-dating filters: a well-specified procedure should distinguish a short-lived volatility shock from a sustained AI-related explosive episode and not conflate the two.

The SV-ADF test is treated differently from the other methods. Because its expanding recursion accumulates evidence from the start of the chosen window, a very long window can cause early explosive episodes to dominate the test statistic and reduce its sensitivity to more recent dynamics. To ensure comparability between the tests, two complementary windows are defined

and run in parallel. *Window 1 (W1)*, the primary post-ChatGPT window, runs from *1 November 2021* to *1 May 2026*. This start date is set roughly one year before the ChatGPT launch, capturing the pre-announcement anticipation phase while keeping the window narrow enough for the recursion to remain sensitive to the post-2022 dynamics. *Window 2 (W2)*, the pre-ChatGPT comparison window, runs from *1 January 2020* to *31 October 2022*. Running both windows on the same panel allows a direct temporal comparison: if the SV-ADF episode rate is materially higher in W1 than in W2, this constitutes independent corroboration of H1 (post-ChatGPT explosiveness) that is not subject to the full-sample confounding present in the GSADF results.

4.2.3 Preprocessing

Preprocessing is deliberately minimal, so that the tests operate on data as close to the raw observed series as possible. The individual security series are aligned to a common trading calendar and assembled into a single wide table indexed by date. Missing values arising from non-synchronous trading or occasional data gaps are forward-filled, so that a security carries its last observed price across a gap rather avoiding artificial discontinuity. No detrending, smoothing, or outlier removal is applied: the explosive-root tests are designed to operate on the price level and any smoothing would risk removing precisely the dynamics the tests are intended to detect. The assembled price table is stored once as a single cached file and reused by every subsequent stage of the computational pipeline, so that all methods operate on an identical input and any difference in results across methods cannot be attributed to differences in the underlying data.

4.3 Methodology

4.3.1 The multi-method strategy

As established in §2.2.2, no single bubble-detection test dominates the others: their relative performance depends on the data-generating process under the alternative, and rejecting a unit root in favor of explosiveness establishes only a local departure from a random walk, which is necessary but not sufficient for a bubble in the operational sense of this thesis, and neither necessary nor sufficient for one in the theoretical, fundamental-value sense. For both reasons, the empirical strategy of this thesis is explicitly multi-method: the same price series are subjected to a suite of tests, and the analysis concentrates on where the methods agree and on what their disagreements reveal. The suite comprises four components: a set of stationarity

pre-tests (ADF, PP, KPSS); the SADF of Phillips, Wu and Yu (2011); the GSADF and its associated date-stamping procedure (Phillips, Shi and Yu, 2015); and the volatility-robust SV-ADF of Sarkar and Wells (2026). The remainder of this section describes each in turn, together with the Monte Carlo procedure used to obtain critical values, the episode-dating rules, the historical validation design, and the computational pipeline.

4.3.2 Stationarity pre-tests: ADF, Phillips–Perron and KPSS

Before the explosive-root tests are applied, each price series is examined with three classical stationarity tests. Their role is preliminary diagnostic: they establish the order of integration of the series and confirm that the price levels behave, in the absence of an explosive episode, as the non-stationary processes that the right-tailed tests assume as their null. The series must be non-stationary in a unit-root sense for the explosive-root test to be meaningful.

The Augmented Dickey–Fuller (ADF) test estimates the regression

$$\Delta p_t = \alpha + \delta p_{t-1} + \sum_{j=1}^k \varphi_j \Delta p_{t-j} + \varepsilon_t$$

and tests the null of a unit root ($\delta = 0$) against the left-tailed alternative of stationarity ($\delta < 0$). The number of lagged differences k is selected by the Akaike Information Criterion⁸⁶ (AIC). ADF was implemented in Python with `statsmodels.tsa.stattools.adfuller`.

The Phillips–Perron test (§2.2.2) targets the same unit-root null as the ADF. Lacking a dedicated Python library, it was implemented manually by estimating an AR(1) regression in `statsmodels`, then applying a Newey–West long-run-variance correction to compute the PP Z-t statistic. The developed method is in the Appendix 8.3.1.

The KPSS test (§2.2.2), whose null is stationarity, was implemented with `statsmodels.tsa.stattools.kpss`, using the chosen lag option and comparing the test statistic with the reported critical values.

For the price levels examined here, the expected and economically sensible result is that the ADF and Phillips–Perron tests fail to reject the unit-root null while the KPSS test rejects its stationarity null, jointly indicating series that are integrated of order one, the maintained assumption of the explosive-root tests that follow.

⁸⁶ Akaike Information Criterion (AIC) follows the standard practice of selecting the specification that best balances goodness of fit with model parsimony and mitigates residual autocorrelation in the ADF regression

4.3.3 The Supremum ADF test (SADF)

The SADF test of Phillips, Wu and Yu (2011), introduced in §2.2.2 is implemented here on the expanding-window recursion described there. In the present implementation the SADF is evaluated against the asymptotic right-tailed 5% critical value of 1.49 reported by Phillips, Wu and Yu (2011); the Python implementation is given in Appendix 8.3.2. As noted in §2.2.2, the fixed window origin restricts the SADF to a single explosive episode, a limitation the GSADF removes.

4.3.4 The Generalised SADF and BSADF date-stamping

The GSADF test of Phillips, Shi and Yu (2015), constructed in §2.2.2, is implemented here with the flexible-window recursion and BSADF date-stamping defined there. Two implementation choices are recorded below.

First, the minimum window fraction follows the Phillips, Shi and Yu (2015) rule of thumb $r_0 = 0.01 + 1.8/\sqrt{T}$; for the sample used here ($T \approx 1,850$ trading days) this gives $r_0 \approx 0.052$, a minimum window of about 95–96 observations, balancing power against the risk of spurious explosiveness in very short subsamples.

Second, the GSADF and BSADF supremum must be evaluated over a doubly indexed set of admissible subsamples, making this the most demanding part of the pipeline; the implementation described in §4.3.9 accelerates the inner double loop substantially. The full code is available in Appendix 8.3.3.

4.3.5 Monte Carlo critical values

The distributions of the SADF and GSADF statistics under the null hypothesis are non-standard: they depend on functionals of Wiener processes that do not reduce to tabulated distributions, and their shape depends on both the sample size and the minimum window size. Inference therefore relies on simulation. The entire simulation is implemented in Python: for each combination of sample length and minimum window present in the universe, a custom function, available in the Appendix 8.3.4, draws a large number of independent random-walk price paths of the corresponding length using NumPy's random-number generator, computes the GSADF and BSADF statistics for each simulated path via the same recursive least-squares routine applied to the real data, and takes the empirical quantiles of the resulting distribution as the critical values. The analysis uses $MC = 999$ Monte Carlo replications, a fixed random seed

(set via `numpy.random.seed`) for full reproducibility, and the 95th percentile of the simulated distribution as the 5% critical value sequence used in the date-stamping rule.

Because the Monte Carlo step is computationally expensive, the simulated critical values are written to a persistent on-disk cache, implemented as a Python dictionary serialized to a JSON file, indexed by a tuple key comprising the sample length, minimum window, number of replications, and random seed. The cache has two consequences that matter for the practical conduct of the analysis. First, it makes the analysis incremental: adding a new security to the universe triggers a fresh simulation only if that security's sample length and minimum window do not already appear in the cache, and an interrupted simulation run can be resumed without repeated work. Second, it makes the choice of replication count an explicit and recorded parameter: results computed with different replication counts coexist in the cache and can be compared directly, which is used in the robustness analysis of §**Error! Reference source not found.**

4.3.6 Episode detection and the spurious-bubble problem

Translating the output of the GSADF into a set of dated episodes requires a rule for converting the sequence of BSADF statistics and critical values into discrete intervals. The natural rule for which an episode is any maximal run of observations for which the BSADF statistic exceeds its critical value has a well-known practical weakness. Monte Carlo noise in the simulated critical-value sequence, together with the volatility of the statistic itself, produces a large number of very short crossings, so that the raw output contains many spurious micro-episodes lasting only a handful of days. Phillips, Shi and Yu (2015) do not specify a minimum episode duration; their framework imposes only a minimum window in the recursion itself, and uses a minimum-duration heuristic of the order of the logarithm of the sample size, which for samples of this length is only a few days and does not remove the problem.

Two filtering rules are therefore applied to the raw BSADF output, and they are signaled explicitly because they are an addition to the original methodology rather than part of it. First, a minimum-duration filter discards any candidate episode shorter than twenty-one trading days, approximately one calendar month. Second, a merge rule joins two candidate episodes separated by a gap of 21 trading days or fewer into a single episode, on the grounds that a brief, sub-monthly interruption does not economically constitute the end of one bubble and the start of another. The two thresholds are deliberately aligned with the 21-day collapse filter of the SV-ADF, so that the two methods apply a comparable notion of minimum episode duration.

Both thresholds are exposed as parameters of the pipeline and their values are reported alongside every set of results, so that the sensitivity of the episode count to the filtering choice can be assessed.

4.3.7 The volatility-robust SV-ADF test

The fourth component of the battery is the SV-ADF test of Sarkar and Wells (2026), a recent procedure for bubble date-stamping designed to remain valid in the presence of the time-varying volatility that characterizes equity returns.

The mechanics of the SV-ADF and the rationale for its volatility-robust, asymmetric calibration are set out in §2.2.2.6; only the implementation choices are recorded here. The origination and collapse thresholds are applied in closed form as $\ln(n)/10$ and $\ln(n)/2$ respectively, so the procedure requires no Monte Carlo step. Episode dating is filtered by calendar-duration rules: an origination is recognised only when the statistic stays above the origination threshold for at least forty-two consecutive trading days, and a collapse only when it stays below the collapse threshold for at least twenty-one.

Dr. Sarkar's principal methodological observation concerned the sensitivity of the SV-ADF to the choice of sample window. Because the test's recursion expands from a fixed origin, a window that is too broad causes the procedure to lock onto the earliest or most intense explosive episode in the sample, so that later episodes are underweighted. The practical recommendation that follows, and the reason the SV-ADF window is configurable rather than fixed, as noted in §5.2.3, is to apply the SV-ADF over windows narrow enough to isolate the period of interest. The default post-ChatGPT window reflects this advice.

The SV-ADF dates at most one origination–collapse pair per window (the single-episode constraint explained in §2.2.2.6). The practical implication for this thesis is that for firms with multiple GSADF-dated episodes, the SV-ADF will identify only the single most statistically sustained explosive phase within the chosen window. The two windows W1 and W2 are therefore not designed to count episodes but to isolate distinct market regimes; their episode-rate comparison is a test of whether explosive dynamics became more prevalent after the ChatGPT launch, not a test of how many separate cycles occurred within that window.

The collapse classification in the output distinguishes between two cases that carry different economic interpretations. An episode classified as "end-of-sample" means the SV-ADF statistic remained above the origination threshold throughout the window and no collapse was recorded

before 1 May 2026: the explosive dynamics, as measured by this procedure, had not abated at the end of the observation period. An episode classified as "post-bridge" means the statistic crossed the origination threshold for the required minimum run and subsequently fell below the collapse threshold $\log(n)/2$ within or just after the bridge window. The bridge window is a configurable interval following the estimated origination date within which a transient retreat below the collapse threshold is treated as a consolidation rather than a genuine collapse; only a sustained breach below the threshold lasting beyond the bridge terminates the episode. The rationale for the bridge is the same economic logic behind the minimum-run filter: a drawdown that reverses within weeks is more likely to reflect intra-bubble volatility than a structural change in the autoregressive dynamics. The post-bridge classification therefore indicates a completed cycle (origination and collapse both dated) while end-of-sample indicates an ongoing or unresolved episode.

4.3.8 Validation on historical bubbles

Before the test battery is applied to AI-exposed securities, GSADF and SV-ADF are validated against historical episodes date-stamped by the academic literature for each one of them. For GSADF, the validation is carried out on three historical episodes whose existence and approximate timing are independently documented. The first is the dot-com episode in the Nasdaq Composite, examined over a sample beginning in 1990; Phillips, Wu and Yu (2011) date the explosive episode in this index as originating in the mid-1990s and collapsing in 2000. The second is the dot-com episode as it appears in the broader S&P 500, examined over a sample beginning in 1995. The third is the pre-2008 cycle in the S&P 500, examined over a sample beginning in 2003. For SV-ADF, the only historical episode date-stamped by Sarkar (2026) among the mentioned is the Nasdaq Composite over the 1982-2012 window. The validation exercise serves a specific purpose: it does not establish that the methods are correct in any absolute sense, but it confirms that the implementation, as built, recovers episodes that are known to have occurred. A failure to recover these episodes would indicate an implementation error; their successful recovery provides the evidence base for hypothesis H5 and licenses the application of the same code to the AI universe.

4.3.9 Computational implementation and reproducibility

The empirical analysis is carried out with a purpose-built computational pipeline written in Python. Because reproducibility itself is a methodological commitment of this thesis, the design of the pipeline is described here and further explored in the Appendix.

The pipeline is organized into separate layers with distinct responsibilities. A methods layer implements each statistical test as a self-contained function. A pipeline layer orchestrates the application of each method to each security and writes the results to disk. A plotting layer reads results from disk and produces figures and tables, and is fully separated from the computation, so that the presentation of results can be revised without recomputing anything. This separation means that each of the four method families can be run independently: the SV-ADF can be re-estimated, for example, without re-running the GSADF or repeating the Monte Carlo simulation.

Three features of the implementation are specifically directed at reproducibility and computational tractability. The GSADF inner double loop, which is the performance bottleneck, is compiled to machine code so that the Monte Carlo simulation completes in a practical amount of time, and the compiled kernel is itself cached between runs. All stochastic components use a single fixed random seed, so that the Monte Carlo critical values, and therefore every downstream result, are exactly reproducible. The pipeline also records, alongside each result, the parameter values used to produce it, including the Monte Carlo replication count and the episode-filtering thresholds. Finally, the dependency versions are pinned, so that the computational environment can be reconstructed exactly. The architecture is deliberately structured so that the same result files can later be served by an interactive front-end, allowing the analysis to be explored selectively without re-execution.

5. Empirical results

The following chapter contains a synthesis of the most relevant findings of the multi-method explosivity testing across the AI value chain. The analysis has indeed produced a much richer set of visualizations diagrams and comparison tables that are available in the output folder of the repository: (link) **Github Master's Thesis repository**.

5.1 Preliminary Analysis

5.1.1 Stationarity results

Three complementary unit-root diagnostics (the Augmented Dickey-Fuller (ADF), Phillips-Perron (PP) and KPSS tests) are applied to row price levels of all 56 securities over the full sample. ADF and PP share the null hypothesis of a unit root, so failure to reject is evidence of non-stationarity (I(1)); KPSS inverts this logic, testing level stationarity under the null, so rejection signals I(1). The three tests therefore form a cross-validating battery whose unanimous agreement constitutes strong evidence on the order of integration.

Table 5.3 — Stationarity test results. Decisions are reported at the five-per-cent significance level.

Summary Table — for the main text:

	<i>ADF (H_0: unit root)</i>	<i>PP (H_0: unit root)</i>	<i>KPSS (H_0: level stationary)</i>
<i>Fail to reject H_0</i>	55 / 56 (98.2%)	56 / 56 (100%)	1 / 56 (1.8%)
<i>Reject H_0</i>	1 / 56 (1.8%)	0 / 56 (0%)	55 / 56 (98.2%)
<i>5% critical value</i>	≈ -2.86	-2.86	0.463
<i>Typical stat. range</i>	-2.8 to $+9.6$	-2.8 to $+9.6$	0.26 to 6.15

Across the full universe, the verdict is unambiguous: 54 of 56 securities (96.4%) receive a unanimous I(1) classification, with ADF and PP both failing to reject the unit-root null and KPSS strongly rejecting level stationarity: KPSS statistics typically range from 1.1 to 6.2, roughly 2 to 13 times the 5% critical value of 0.463. The two partial exceptions, C.AI and

CRWV, receive an I(1)* designation reflecting a 2-of-3 majority: for C.AI the ADF statistic reaches -5.46 (rejected) while the PP and KPSS results still point to non-stationarity; for CRWV, the very short available history ($T = 274$ observations since IPO) renders the KPSS test low-powered and its failure to reject is not diagnostic. Both series are therefore retained as I(1) for all subsequent analysis. The unanimous I(1) finding confirms the prerequisite for the GSADF and SV-ADF frameworks: all price series contain a stochastic trend against which a transient explosive component can be identified. Full test statistics, critical values and p-values for each ticker are reported in the Appendix 8.1.2.2.

Note: I(1) cases are C.AI (ADF rejects at -5.46 , PP and KPSS still indicate I(1); $T = 1,353$) and CRWV (KPSS fails to reject but $T = 274$ since IPO; ADF and PP fail to reject). Both are retained as I(1) in the main analysis. Read: for ADF and PP, rejection means stationary (I(0) signal); for KPSS, rejection means non-stationary (I(1) signal). Unanimous I(1) therefore

I(1) — unanimous	54	96.4%	ADF & PP fail to reject; KPSS rejects — all three agree
I(1)* — (2/3)	2	3.6%	Mixed signals; majority still non-stationary (see note)
I(0) — any	0	0%	—
Total	56	100%	All treated as I(1) for bubble detection

requires ADF fail-to-reject + PP fail-to-reject + KPSS reject.

Table 6 - ADF, PP, KPSS summary results

5.2 Explosivity analysis

5.2.1 SADF and GSADF results on indices and benchmarks

The GSADF test is first applied to the three benchmark indices: the Nasdaq Composite (^IXIC), the PHLX Semiconductor Index (^SOX), and the S&P 500 (^GSPC). The test is carried out over the full sample window of 1 January 2020 to 1 May 2026. For each index, the GSADF statistic is compared against its Monte Carlo 95% critical value and the BSADF date-stamping procedure is applied to identify any explosive sub-periods. The results are summarised in Table 7.

Table 7 - GSADF results for the benchmark indices over the 2020–2026 sample window

Index	GSADF statistic	MC 95% critical value	Global null rejected?	Dated episodes
Nasdaq Composite (^IXIC)	1.04	2.39	No	0
PHLX Semiconductor (^SOX)	2.21	2.39	No	0
S&P 500 (^GSPC)	1.35	2.39	No	0




Figure 19 - GSADF results, Index/control

None of the three indices rejects the global null hypothesis of no explosive dynamics at the 95% significance level (no red shading in either of the three rows in the Figure 19). The PHLX Semiconductor Index returns the highest GSADF statistic of the three (2.21), which falls 0.18 units below the simulated 95% critical value of 2.39; the Nasdaq Composite and the S&P 500 return statistics of 1.04 and 1.35 respectively, both well within the non-rejection region. The BSADF date-stamping procedure identifies no episode satisfying the minimum-duration filter for any of the three series.

The uniform failure to reject across all three indices constitutes direct empirical support for hypothesis H3. The Nasdaq Composite and the PHLX Semiconductor Index contain, as constituents, several firms whose individual GSADF statistics range from 2.9 to 8.5 and who individually exhibit multiple dated explosive episodes as described later in this section. That the index-level statistics are substantially lower than even the weakest constituent rejection reflects the regime-mixing mechanism identified by Evans (1991): aggregating across firms whose explosive episodes do not coincide in calendar time, and across firms whose price series exhibit no explosive behavior at all, dilutes the aggregate signal to the point of

statistical non-detection. The PHLX Semiconductor Index comes closest to the rejection threshold, consistent with its relatively high concentration of AI-relevant constituents; the gap of 0.18 between its GSADF statistic and the 95% critical value nonetheless prevents rejection under the baseline parameterization, and the non-rejection is confirmed at the 90% threshold as well.

This finding carries a direct practical implication. An investor or analyst relying on index-level price behavior to assess whether speculative dynamics are present in AI-exposed markets would, on the basis of the GSADF test applied to these three indices, find no statistical basis for such a characterization over the sample window. The same test applied to the constituent equities at the firm level, as reported in §5.1.3, yields a majority-rejection rate of 67.7%. Index-level absence of explosive signals is indeed a positive signal if analyzed through the lens of the dot-com bubble: the Nasdaq over that period failed to reject the null hypothesis under the same parameters. Indexes remain less informative than firm-level diagnostics when speculative behavior is concentrated in a subset of constituents, consistent with the empirical findings of Basele, Phillips and Shi (2025) for technology-sector indices over comparable periods.

5.2.2 GSADF results on individual equities

This paragraph presents the firm-level GSADF results, organized by value-chain segment so that the cross-sectional pattern can be read directly against the taxonomy of Table 4. For each security the GSADF statistics, its critical value, the reject decision, and the dated episodes are reported in the Appendix Table 13, with the episode-filtering thresholds stated alongside. The analysis is then visually aggregated to the segment level, comparing the incidence, timing and duration of episodes across the eight segments. This is the central evidence for sub-question SQ2 and hypothesis H2: it is where the claim that explosive behavior is concentrated in the upstream, supply-constrained segments is tested. A segment-level summary table reports, for each segment, the number of constituent firms, the number with at least one dated episode, and the total episode count.

Of the 37 securities in the equity universe, 26 (70.3%) return GSADF statistics that exceed their Monte Carlo 95% critical values, rejecting the null of no explosive dynamics. A further eight securities (ASML, AMAT, AMZN, CSCO, DDOG, META, NBIS, QCOM) fall marginally below the formal rejection threshold yet produce at least one dated BSADF episode, because the BSADF episode-dating procedure applies a time-varying point-in-time critical value that

differs from the single global supremum threshold. When both the formal rejection and the presence of at least one dated episode are considered together, 34 of the 37 tickers (91.9%) exhibit explosive behavior. The three securities with neither a formal rejection nor a dated episode are EQIX, DLR, and SNOW. The segment-level aggregation is presented in

Table 8.

<i>Layer</i>	<i>Segment</i>	<i>N</i>	<i>Reject (n/%)</i>	<i>≥1 episode (n/%)</i>	<i>Episodes</i>	<i>Timing</i>
Up	1. Semiconductor Equipment	4	2/50%	4/100%	7	2020–21; 2025–26
Up	2. Chip Design	5	5/100%	3/60%	8	2023–25
Up	3. Memory (HBM)	1	1/100%	1/100%	3	2020–21; 2024–26
Up	4. Foundry + Adv. Packaging	2	2/100%	2/100%	7	Recurrent 2020–26
Up	5. Storage (control)	2	2/100%	2/100%	4	2020–21; 2024–25
Mid	6. DC Networking & Optical	4	3/75%	3/75%	4	2024–26
Mid	7. DC REITs	2	0/0%	0/0%	0	—
Mid	8. Power & Grid	2	2/100%	2/100%	4	2023–25
Mid	9. Hyperscale Cloud	5	3/60%	4/80%	10	2021; 2023–25
Mid	10. Specialised AI Cloud	2	1/50%	1/50%	1	2025
Down	11. Enterprise AI	1	1/100%	1/100%	2	2023–25
Down	12. App & AI Platform SW	4	2/50%	2/50%	2	2024–26
Down	13. Edge AI & Devices	3	2/67%	3/100%	4	2019–21

Table 8 -Segment-level GSADF episode summary, used to assess hypothesis H2

Note: "GSADF null rejected" uses the global 95% Monte Carlo critical value. "Firms with ≥ 1 episode" counts tickers with at least one BSADF episode passing the 60-day minimum-duration and 42-day merge-gap filters. These two counts differ for tickers where the global supremum statistic falls below the 95% CV but the time-varying BSADF path crosses its point-in-time critical value for a sustained period. Segment 5 (Storage) is included as an intra-universe control rather than a hypothesis-bearing segment.

The directional ordering predicted by hypothesis H2 is supported: the upstream segment rejection rate of 85.7% (12/14) exceeds the downstream rate of 62.5% (5/8), a wider margin than in the original 31-ticker panel (80.0% vs 62.5%) and the mean episode count per firm is 1.75 upstream versus 1.0 downstream.

➤ Segment 1: Semiconductor Equipment (ASML, AMAT, LRCX, KLAC)

All four firms produce at least one dated episode, and two formally reject the null: Lam Research (LRCX, GSADF statistic 4.52) and KLA Corporation (KLAC, 2.86); ASML (2.00) and Applied Materials (AMAT, 2.22) fall marginally below the 95% critical value yet each generate dateable episodes through the BSADF rolling procedure. The episodes cluster in two phases: a wave in 2020–2021 coinciding with the broad semiconductor supercycle, and a renewed wave in late 2025 and ongoing in 2026 as wafer-fab equipment orders track the announced TSMC capacity expansion. The tight sequencing of this second wave across three of the four names is consistent with the long procurement cycles of semiconductor equipment and with the segment's position at the furthest upstream point of the capital-deployment sequence.





Figure 20 - BSADF datestamping - Semiconductor Equipment

➤ Segment 2: Chip Design (NVDA, AMD, INTC, AVGO, MRVL)

All five firms formally reject the global null, the only segment in the universe with a 100% formal rejection rate. NVIDIA's BSADF path remains above its critical-value sequence for 704 consecutive calendar days from 20 March 2023 to 21 February 2025 — the longest episode in the dataset — spanning the full period of the company's data-centre revenue acceleration from approximately USD 5B to USD 50B per quarter. Broadcom (AVGO) produces three

overlapping post-ChatGPT episodes covering the period from May 2023 to the end of the sample. AMD generates one episode originating in July 2025, reflecting the delayed uptake of its Instinct GPU line. Intel (INTC) and Marvell (MRVL) reject the global null statistically but produce no dated episode passing the minimum-duration filter, indicating explosive dynamics that are real but too brief or fragmented to constitute a sustained episode; INTC's detection against a backdrop of declining market share and negative operating momentum is noted as an episodic rather than structural finding.


Figure 21 - BSADF datestamping - Chip Design



➤ Segment 3: Memory, HBM-relevant (MU)

Micron Technology is the sole representative of the HBM layer and returns a GSADF statistic of 7.86, rejecting the null with a wide margin. Three dated episodes span the sample: one in November 2020–March 2021 (the COVID-era DRAM shortage), one in March–July 2024 (the first wave of HBM3e demand from Nvidia), and one from September 2025 onward. The lag between NVDA's March 2023 episode origination and MU's first post-ChatGPT episode in March 2024 is consistent with the 18-to-36-month HBM capacity cycle documented in §3.2.2 and reflects the time required for AI accelerator demand to translate into memory procurement.

Figure 22 - BSADF datestamping - Memory



➤ Segment 4: Foundry and Advanced Packaging (TSM, ASX)

TSMC (TSM) produces five dated episodes across the full sample, the highest episode count of any single security in the upstream universe, reflecting its position as the dominant advanced-node foundry and, via CoWoS, the principal advanced-packaging supplier for the leading accelerator architecture: each intensification of AI infrastructure demand registers as an episodic repricing of TSMC's capacity utilization and pricing power, with post-ChatGPT episodes originating in January 2024, October 2024, and September 2025. ASE Technology (ASX, stat 6.13), the outsourced advanced-packaging subcontractor added to this segment, formally rejects the null and produces two episodes: one in the 2020-2021 semiconductor boom (December 2020-March 2021) and one from October 2025 onward, still running at the end of the sample. ASX's second episode is consistent with the gradual diversification of advanced-packaging capacity away from TSMC's CoWoS monopoly toward outsourced subcontractors, a development that entered mainstream supply-chain reporting in late 2025; the two-year gap between NVDA's March 2023 episode origination and ASX's second post-ChatGPT episode reflects the long lead times of packaging-capacity expansion documented in §3.2.2. The segment's episode timing, TSM's one-to-three-quarter lag behind the chip-design episodes, ASX's confirmation arriving still later, is consistent with the downstream transmission of accelerator demand first into wafer orders and only subsequently into outsourced packaging capacity.




Figure 23 - BSADF datestamping - Foundry and Advanced Packaging

➤ Segment 5: Storage (WDC, STX)

Both Western Digital (WDC, GSADF statistic 8.49) and Seagate (STX, 7.94) formally reject the null of no explosive dynamics and produce two dated episodes each, including episodes that remain active toward the end of the sample. These results indicate that explosive price behaviour was not confined to the accelerator-centred semiconductor segments but also extended to the persistent-storage infrastructure used by cloud and data-centre operators. The dated episodes nevertheless cannot be attributed exclusively to AI demand. The earlier episodes overlap with the post-pandemic technology and storage cycle, while the episodes beginning in 2024–2025 coincide with the expansion of hyperscale infrastructure and the growing storage requirements associated with AI-generated and AI-consumed data. Other potentially relevant drivers include industry supply discipline, changing storage prices, cloud procurement cycles, product transitions and broader revisions in expected data-centre demand. In Western Digital’s case, interpretation is further complicated by the separation of its flash business in 2025, meaning that the earlier price history reflects a broader HDD-and-flash business perimeter than the company represented toward the end of the sample. The evidence should consequently be read as showing that explosive dynamics reached a complementary data-centre infrastructure segment, without establishing either that AI demand was their sole cause or that the storage episodes formed a causal continuation of the semiconductor propagation sequence




Figure 24 - BSADF datestamping - Storage

➤ Segment 6: Data-Centre Networking and Optical (ANET, CSCO, NOK, CIEN)

Three of four firms formally reject the null (75%), with Cisco (CSCO, stat 1.37) the exception. The segment shows a clear division between data-centre-native players and legacy telecoms equipment vendors: Arista Networks (ANET, stat 3.04) produces two episodes totalling 525 days, with its first episode beginning in December 2023 as data-centre fabric order volumes became publicly visible; Ciena (CIEN, stat 7.94) generates one episode from September 2025 onward, reflecting optical interconnect demand tied to next-generation AI cluster buildouts. Nokia (NOK, stat 5.14) formally rejects but produces no dated episode passing the minimum-duration filter, suggesting that the detected explosivity is only fragmented. Nokia is, however, an imperfect segment representative because its exposure is primarily to broader telecommunications infrastructure rather than specifically to data-centre networking; its result is therefore interpreted as that of a legacy networking comparator. All episodes in this segment are post-ChatGPT, making it one of the latest-originating of the upstream and midstream layers, consistent with the capital cycle's progression from silicon to physical infrastructure.

GSADF episodes — 6. Data-Centre Networking & Optical




Figure 25 - BSADF datingstamping - DC Networking & Optical

➤ Segment 7: Data-Centre REITs (EQIX, DLR)

Neither EQIX nor DLR produces a formal rejection or a dated episode. This is the most striking negative finding in the equity panel, given that both firms operate physical infrastructure directly exposed to AI buildout demand. The most plausible explanation is the interaction between the REIT financing structure and the Federal Reserve tightening cycle of 2022–2023: rising discount rates mechanically penalise REIT valuations through the yield-sensitive component of their pricing, partially offsetting any AI-demand signal in the price series. This confound is specific to the REIT structure and does not apply to the non-REIT names in adjacent segments; the result should be read as a product of macroeconomic environment rather than as evidence that physical data-centre capacity is unaffected by speculative dynamics.

GSADF episodes — 7. Data-Centre REITs (colocation & build-out)




Figure 26 - BSADF datingstamping - Data-Centre REITs

➤ Segment 8: Power & Grid (CEG, VST)

Vistra (VST, stat 8.49) and Constellation Energy (CEG, stat 3.05) both formally reject the null and each produce two dated episodes—the strongest rejection rate (100%) of any midstream segment. VST’s first episode begins in December 2023, contemporaneously with the emergence of data-centre power demand as a visible constraint in industry reporting (§3.2.2); its second episode, from June to October 2025, coincides with the broader midstream repricing wave also visible in Oracle and Ciena. CEG’s two episodes, from June to November 2023 and from February to July 2024, occur during the period in which hyperscaler agreements increased the market visibility of nuclear generation as a source of data-centre power. These temporal correspondences are consistent with, but do not establish, an AI-demand explanation, since both companies’ valuations are also influenced by electricity prices, generation mix, regulatory developments and broader power-market conditions. Overall, the results support treating electricity generation as a distinct component of midstream compute infrastructure: both firms exhibit strong explosive dynamics during the period in which power availability became an increasingly salient constraint on data-centre expansion.




Figure 27 - BSADF datestamping - Power and Grid

➤ Segment 9: Hyperscale Cloud (MSFT, AMZN, GOOGL, META, ORCL)

Three of five formally reject the random-walk null, and four of five produce dated episodes; with ten total episodes, this is the highest aggregate episode count of any segment in the universe. Oracle (ORCL, stat 3.49) exhibits four episodes including one originating in September 2024, approximately contemporaneous with the Stargate consortium announcement and Oracle's USD 300B OpenAI compute commitment. Google (GOOGL, stat 2.75) produces two episodes, with the second ongoing from September 2025. Microsoft (MSFT, stat 2.57) produces a marginal rejection; Meta (META, stat 2.14) and Amazon (AMZN, stat 2.32) fall below the formal rejection threshold despite each generating episodes, consistent with the attenuation of their AI-specific price signal by the breadth of their diversified revenue bases.

The hyperscale cloud segment therefore exhibits a heterogeneous cross-sectional pattern that is more informative at the firm level, particularly for Oracle, than at the segment average.




Figure 28 - BSADF datestamping - Hyperscale Cloud

➤ Segment 10: Specialised AI Cloud

This segment has the shortest average trading history in the universe, and its results should therefore be interpreted cautiously. CoreWeave (CRWV) rejects the global null with a GSADF statistic of 3.31 against a critical value of 1.84, calibrated to its 274-observation post-IPO sample. However, the BSADF procedure identifies no episode that passes the 60-day minimum-duration filter; the global rejection therefore indicates the presence of explosivity somewhere within the available sample, rather than a continuously explosive path from the March 2025 IPO. Nebius (NBIS, stat 1.86) falls below its critical value of 2.18 but generates a 170-day dated episode from May to October 2025, reflecting the difference between the global GSADF decision and the time-varying BSADF dating criterion. With one of two firms formally rejecting the null, the segment records a 50% rejection rate and only one sustained dated episode. The findings provide preliminary evidence of explosive dynamics in specialised AI-cloud equities, but the limited and non-comparable trading histories prevent a strong segment-level conclusion.




Figure 29 - BSADF datestamping - Neoclouds

➤ Segment 11: Enterprise AI / Hybrid Cloud (IBM)

IBM is the sole representative of this segment and formally rejects the null with a statistic of 2.69. Its two dated episodes, November 2023–April 2024 (133 days) and September 2024 onward (510 days and ongoing), are both post-ChatGPT, with the second being the fourth-longest episode in the equity dataset. The first episode coincides with the period in which enterprise AI pilot deployments began generating reportable commercial revenue; the second corresponds to consecutive quarters of disclosed AI-attributable growth in IBM's Consulting and Software segments. The segment is too small for cross-sectional inference, but the result illustrates that explosive dynamics are not confined to pure-play infrastructure firms and extend to established incumbents successfully repositioning into enterprise AI.




Figure 30 - BSADF datestamping - 9 Enterprise AI / Hybrid Cloud

➤ Segment 12: Application and AI Platform Software (PLTR, SNOW, DDOG, AI)

The segment (Figure 31) shows the widest divergence of any in the universe: two firms formally reject the null (50%) and two produce no episode at all (SNOW, DDOG). Palantir (PLTR, stat 6.72) produces the third-longest episode in the dataset, originating in September 2024 and running for 505 days to the end of the sample; its origination date is substantially later than all upstream episodes, consistent with speculative dynamics reaching the application layer only after the AI infrastructure narrative had been sustained for multiple quarters. C3.ai (AI, stat 4.33) formally rejects but produces no dated episode, a pattern consistent with multiple short-

lived price accelerations around earnings and partnership announcements rather than a sustained explosive phase. Snowflake (SNOW) and Datadog (DDOG) exhibit no explosive dynamics in the post-ChatGPT window, likely reflecting the market's more sceptical assessment of their AI monetisation timeline relative to infrastructure-layer names.




Figure 31 - BSADF datestamping - Application & AI Platform Software

➤ Segment 13: Edge AI and Devices (QCOM, AAPL, TSLA)

All three firms produce dated episodes, but every episode in this segment predates the ChatGPT launch by more than a year. Tesla's two episodes (October 2019–March 2020, June 2020–April 2021), Qualcomm's single episode (November 2020–February 2021), and Apple's episode (October 2019–February 2020) all originate in the general technology and growth-stock rally of 2019–2021. TSLA formally rejects with a high statistic of 7.83, driven by the magnitude of its COVID-era re-rating; AAPL marginally rejects (2.43 vs CV 2.39); QCOM falls below the formal rejection threshold (2.14). The pre-ChatGPT sub-period robustness check (§5.1.7, A1) confirms this: the edge devices segment's explosive dynamics are attributable to the 2020–2021 technology bull market rather than to the AI capital cycle, and contribute to the overall equity rejection rate without bearing on any of the AI-specific hypotheses.

GSADF episodes — 11. Edge AI & Devices




Figure 32 - BSADF datestamping - Edge AI & Devices




Figure 33 - Post ChatGPT BSADF paths – explosivity propagation evidence

The temporal sequence of episode origination dates, visualized across all segments in Figure 33, reveals a downstream pattern that mirrors the capital deployment sequence of Chapter 3. Explosive dynamics originate in chip design (Q1–Q2 2023), propagate to the power and grid and foundry layers (Q3–Q4 2023 and Q1 2024 respectively), advance to memory and cloud (Q1 and Q3–Q4 2024), and reach the application software layer last (Q4 2024). The edge devices segment's pre-ChatGPT episodes stand apart from this sequence entirely. This propagation structure, read against the account of Chapter 3, provides evidence in line with the price-market signature of the AI capital cycle, although it should be interpreted as evidence of alignment rather than proof of causality.

5.2.3 SV-ADF results and cross-method comparison

The SV-ADF procedure of Sarkar and Wells (2026), whose technical specification is set out in §4.3.7, is designed to remain valid under the persistent, time-varying volatility that characterizes AI equity returns, a feature that standard homoskedastic tests, including the GSADF in its baseline form, do not explicitly accommodate. The practical consequence is that the SV-ADF is substantially more conservative: it requires explosive dynamics to be sustained across a minimum run length and to survive a volatility-aware threshold calibration before an episode is recorded, and it deliberately declines to flag high-volatility periods as explosive unless the recursive coefficient path remains durably above the origination threshold. Where the GSADF and the SV-ADF agree, the finding carries greater weight than either method alone; where they disagree, the divergence itself carries diagnostic content about whether the explosive dynamics are sustained or confounded with volatility clustering.

The SV-ADF is applied here in two complementary windows as described in §4.3.7. Running both windows on the same panel allows a within-methodology test of hypothesis H1: if the SV-ADF episode rate is materially higher in W1 than in W2, this constitutes independent corroboration, through a procedure orthogonal to the GSADF, that explosive dynamics became concentrated after the ChatGPT launch rather than being a continuous feature of the sample.

Figure 2c — Top 8 tickers by SV-ADF explosivity. * = detected in both windows (W1 post-ChatGPT orange, W2 pre-ChatGPT teal). Dashed vertical = formal collapse date.




Figure 34 - Top 8 tickers by SV-ADF explosivity

Applied over W1, the SV-ADF identifies episodes in 14 of the 37 equity tickers, an overall rate of 37.8%, against a W2 rate of 0.0% (no equity in the universe generates a W2 episode). The upstream segment rate in W1 is 57.1% (8/14), the midstream rate is 33.3% (5/15), and the downstream rate is 12.5% (1/8). These aggregate figures confirm the core directional

predictions of H1 and H2: the post-ChatGPT window shows far more explosive activity than the pre-ChatGPT window, and the upstream layer leads the downstream layer by a wide margin.

Within the upstream chip design segment, NVIDIA and Broadcom are the two firms for which the SV-ADF returns W1 episodes, with origination dates of 10 January 2024 and 8 January 2024 respectively, within two trading days of each other. Both are classified as post-bridge collapses, meaning the explosive phase met the sustained-run criterion, subsequently crossed the collapse threshold within the bridge period, and was dated as a completed episode. The temporal clustering of these two originations in the opening days of 2024 is consistent with a market-wide re-rating event: the period immediately following Nvidia's November 2023 earnings release, at what point data-center revenue guidance began to confirm AI infrastructure demand at a scale that materially revised consensus forecasts upward. AMD, Intel, and Marvell produce no W1 episode.




Figure 35 - SV-ADF date stamping - 2 Chip Design

For AMD (Figure 35), who's single GSADF episode runs for 107 days from July 2025 (Figure 21), the explosive sub-period is too fragmented to satisfy the SV-ADF's sustained-run requirement. Intel and Marvell represent the two cases where a formal GSADF rejection (statistics of 3.89 and 3.60 against a 95% critical value of 2.39) coexists with no dated episode in the BSADF output, an apparent contradiction that has a precise methodological explanation. The GSADF statistics takes the highest value of the rolling BSADF statistics over all valid

sub-samples in the full dataset. In practice, this means that even a brief moment when the BSADF path rises above its point-in-time critical value is enough to push the overall GSADF statistic above the global threshold and lead to rejection of the null hypothesis. The 60-day minimum-duration filter applied at the episode-dating stage then discards any crossing that does not last at least 60 consecutive trading days, so that no episode appears in the figure even though the null is formally rejected. For INTC and MRVL, the BSADF path reached its peak threshold crossing too briefly to satisfy the duration filter producing a rejection without a dateable episode. The SV-ADF confirms this reading (Figure 35): neither firm generates a W1 episode, consistent with explosive dynamics that are statistically detectable in the supremum but too transient to constitute a sustained explosive phase under either the duration filter or the SV-ADF's sustained-run criterion.

In the memory and foundry-and-packaging segments, all three constituents generate a W1 episode. Micron (MU, origination 1 October 2025) is end-of-sample, unabated at the close of the window. TSMC (TSM, origination 9 September 2025) is classified post-bridge, having fallen back below the collapse threshold by December 2025. ASE Technology (ASX) produces its own end-of-sample episode from December 2025. The late origination of all three names, relative to the chip-design segment, is consistent with the lag structure identified in §5.1.3, with ASX's arrival a further quarter behind TSM consistent with outsourced packaging capacity absorbing overflow demand only after TSMC's own CoWoS lines are already saturated.




Figure 36 - SV-ADF datestamping - Foundry and memory

Within the semiconductor-equipment segment, only Lam Research (LRCX) produces a W1 episode among the four constituents, with origination dated 24 November 2025, classified end-of-sample. ASML, Applied Materials, and KLA Corporation (KLAC) produce none: ASML's explosive phase is concentrated in the 2020-2021 window that falls outside W1, and KLAC's single GSADF episode (October 2025 onward) is too recent and too short to satisfy the SV-ADF's sustained-run filter.

Ciena (CIEN), Western Digital (WDC), and Seagate (STX) all generate end-of-sample episodes with originations concentrated in the second half of 2025 (CIEN September 2025, WDC August 2025, STX June 2025). ASE Technology (ASX) shares the same late-2025 timing (December 2025) despite sitting in a different segment, reflecting contemporaneous but distinct drivers: optical-interconnect repricing (CIEN), NAND/HDD repricing linked to inference-serving infrastructure scale-up (WDC, STX), and outsourced packaging capacity absorbing overflow demand (ASX). The SV-ADF's confirmation of all four as genuine sustained explosive episodes, rather than volatility artifacts, adds credibility to what might otherwise be read as AI-adjacent. noise.




Figure 37 - SV-ADF datestamping - 5 Storage




Figure 38 - SV-ADF datestamping - 6 Data Center Networking & Optical

In the hyperscale cloud segment, the SV-ADF is notably selective, detecting episodes only in Alphabet (GOOGL, origination 13 October 2025, post-bridge) and Oracle (ORCL, origination 12 September 2024, post-bridge). Microsoft, Amazon, and Meta produce no W1 episode. Oracle's September 2024 origination falls approximately ten trading days after the first major public reporting of the Oracle–OpenAI compute commitment, making it the most structurally event-aligned episode date in the hyperscale group. Alphabet's October 2025 origination coincides with Google Cloud's AI-attributed revenue growth accelerating above consensus. The non-detection of Microsoft, Amazon, and Meta narrows the hyperscale explosive-dynamics story from a broad platform-layer finding to two firm-specific narratives: a compute-commitment-driven episode at Oracle and an AI-cloud-revenue-driven episode at Alphabet.

Neither CoreWeave (CRWV) nor Nebius (NBIS) produces a W1 SV-ADF episode, in contrast to their GSADF results, where CRWV formally rejects the null and NBIS generates a dated episode. This is most plausibly a sample-length effect: CRWV's 274-observation history leaves little room for the SV-ADF's 42-consecutive-step origination filter to be satisfied even where the coefficient path is locally explosive.




Figure 39 - SV-ADF datestamping - 8 Hyperscale Cloud

In the application and platform software segment, Palantir (origination 3 October 2024, post-bridge) is the sole W1 episode in the entire downstream layer, confirming a finding from §5.1.3. The absence of SV-ADF episodes for IBM, C3.ai, Snowflake, and Datadog is informative: none of these names sustains explosive dynamics that survive the heteroskedasticity assumption, suggesting that whatever explosive sub-periods the GSADF detects in IBM and C3.ai reflect shorter or more volatile repricing phases rather than the kind of durable explosive dynamic the volatility-robust procedure requires.

In the power-and-grid segment, Vistra (VST, origination 25 January 2024, post-bridge) and Constellation Energy (CEG, origination 27 February 2024, post-bridge) both originate within days of the NVDA and AVGO chip-design episodes, consistent with §3.2.2, where power and grid constraints are documented as a binding input from early 2024 onward.




Figure 40 - SV-ADF datestamping - Power & Grid (planned)

Window W2 results

Applied over the pre-ChatGPT window, the SV-ADF detects episodes in 35.5% of the equity panel. These episodes are, with minor exceptions, concentrated in the commodity benchmarks: DBC and PDBC produce W2 episodes originating in January 2022, attributable to the Ukraine-driven commodity price shock; SOYB produces an episode beginning in November 2020. No equity in the main AI panel generates a SV-ADF W2 episode. The directional contrast with W1 — 48.4% versus 35.5%, with the entire W2 equity component at zero — confirms, through a methodology that does not share a critical-value family with the GSADF, that the explosive dynamics in the AI equity universe are a post-ChatGPT phenomenon.

External validation against Sarkar and Wells (2026)

All twelve firms analysed by Sarkar and Wells (2026) — the Magnificent Seven plus NVDA, TSMC, AVGO, ASML, MU, and PLTR — are present in the equity universe of this thesis, allowing a direct comparison across independent implementations. For the post-2022 AI period, the agreement with the formal classifications reported in Table 2 of their paper is complete: our W1 results and theirs agree on all twelve firms, confirming positive detection for NVDA, GOOGL, TSMC, AVGO, PLTR, and MU, and non-detection for AAPL, MSFT, AMZN, META, TSLA, and ASML. The origination dates for the six positive detections agree within approximately one calendar month across implementations using different sample windows, different panels, and independently developed code: NVDA (January 2024 in both), TSMC (September 2025 in both), Micron (October 2025 in both), Broadcom (December 2023 – January 2024 in both), Palantir (September – October 2024 in both), and Alphabet (October 2025 in both). The only divergences arise in the pre-ChatGPT period: Sarkar and Wells (2026) detect explosive episodes in ASML (December 2020 – March 2021) and Tesla (June 2020 – March 2021) in their longer 2018–2026 window, while neither appears in our W2. Both

divergences are most plausibly attributable to the shorter initial recursion available in W2 (which begins January 2020 rather than 2018) and the specification difference documented in §4.3.6. The complete post-2022 agreement across two independent implementations provides strong evidence that the SV-ADF results presented in this section are not artefacts of the specific computational pipeline.

The SV-ADF results refine rather than contradict the GSADF findings, allowing us to distinguish explosivity in prices from explosivity in volatility. The upstream-over-downstream pattern is preserved (57.1% versus 12.5% in W1), while the external agreement with Sarkar and Wells (2026) provides additional validation of the implementation. The set of firms confirmed by both methods numbers fourteen: LRCX, NVDA, AVGO, MU, TSM, ASX, WDC, STX, CIEN, CEG, VST, GOOGL, ORCL, and PLTR. Every ticker with a W1 SV-ADF episode also satisfies the GSADF reject-or-episode criterion, meaning that the list of firms supported by both methods coincides exactly with the SV-ADF W1 set. As such, this constitutes the most methodologically conservative bubble-candidate list supported by the dataset.

5.2.4 Commodity-market results

The commodity panel of sixteen instruments, described in §4.2.1, is organized into four analytically distinct groups: AI-relevant energy (USO, UNG, BNO), AI-relevant industrial metals (JJC, PPLT, PALL), precious metals as a confound (GLD, SLV), and two placebo groups — agriculture (DBA, WEAT, CORN, SOYB) and livestock (COW), for which no AI-demand channel exists. The broad benchmark instruments (DBC, GSG, PDBC) serve as aggregate reference points. The test of hypothesis H6 is organized around two criteria. The strong form asks whether the GSADF rejection rate is higher in the AI-relevant groups than in the placebos. The conditional form, stated in §5.1.6, asks whether the fraction of total episode-days falling after Q4-2023, the point at which AI infrastructure demand first appeared in commodity-market reporting, is higher in the AI-relevant groups than in the placebos. Both criteria are assessed against the GSADF results; the SV-ADF provides a supplementary filter. The results are summarized in Figure 41.




Figure 41 - Commodity Market evidence for AI-demand propagation

5.2.4.1 GSADF results: AI-relevant energy

All three energy instruments reject the null of no explosive dynamics: USO (stat 4.79, CV 2.44), BNO (stat 5.11), and UNG (stat 3.02). Four episodes are dated in total. USO produces one episode from 28 February to 26 June 2020 (119 days) and a second from 4 February to 16 June 2022 (132 days). BNO produces one episode from 4 February to 16 June 2022 (132 days), precisely contemporaneous with the USO second episode. UNG produces one episode from 29 June to 25 October 2021 (118 days). The dating is unambiguous: the February–June 2022 episodes in USO and BNO correspond to the Brent crude price spike following the Russian invasion of Ukraine; the UNG 2021 episode corresponds to the European natural gas shortage of that autumn; and the USO February–June 2020 episode corresponds to the COVID-related oil market dislocation. Every dated episode in the AI-relevant energy group originates in a well-identified geopolitical or macroeconomic event that has no plausible AI-demand channel. Not a single energy episode occurs after November 2022.




Figure 42 - GSADF results, commodity sector, energy

5.2.4.2 GSADF results: AI-relevant industrial metals

The industrial metals group presents a more differentiated picture. JJC (copper ETN) fails to reject the global null with a GSADF statistic of 2.04 against a critical value of 2.44, though it produces one BSADF episode from 18 February to 26 July 2021 (158 days), the post-COVID industrial recovery copper rally, via the same mechanism of brief point-in-time threshold crossing without global-null rejection discussed for INTC in §5.1.4. Palladium (PALL) rejects the null with a statistic of 3.70 but produces no dated episode, indicating explosive dynamics that are too fragmented to survive the minimum-duration filter. The one potentially AI-relevant finding in this group is platinum (PPLT): a GSADF rejection at 5.41 paired with a single episode running from 6 June 2025 to 27 February 2026 (266 days). This episode originates in the post-ChatGPT window and falls within the period when platinum-group metals attracted speculative attention amid their use in advanced electronics manufacturing and hydrogen fuel-cell applications. It is, however, an isolated and ambiguous signal: copper, the more direct AI-infrastructure proxy for power distribution and data-center interconnects, produces no episode in the same period; palladium produces none; and the SV-ADF W1 does not confirm a PPLT episode. The platinum finding is noted as a tentative observation rather than evidence for H6.




Figure 43 - GSADF results, commodities, Industrial metals

5.2.4.3 GSADF results: agriculture and livestock placebos

The agriculture group rejects the null for all four instruments: DBA (stat 2.94), WEAT (stat 4.52), CORN (stat 3.75), and SOYB (stat 3.82). Eight episodes are dated in total across the group. The dominant pattern is a wave of explosiveness in late 2020 through mid-2021 — SOYB from September 2020 to August 2021 (344 days), DBA and CORN from December 2020, all attributable to the post-COVID global agricultural demand recovery and supply disruptions — followed by a second wave in February–June 2022 across all four instruments, corresponding to the Ukraine-driven food price shock. A single episode in DBA from March to June 2024 (91 days) falls in the post-ChatGPT window but is attributable to agricultural supply conditions rather than to any AI-demand channel. The livestock instrument (COW) rejects the null with a statistic of 2.95 but produces no dated episode. The 100% GSADF rejection rate in the agriculture group matches the 100% rejection rate in the AI-relevant energy group exactly, which is the most direct falsification of the strong form of H6: a test that cannot distinguish between crude oil, wheat, and soybeans offers no differential diagnostic of AI-infrastructure demand.




Figure 44 - GSADF results, commodities, placebo (Agriculture and Livestock)

5.2.4.4 GSADF results: precious metals and benchmarks

Gold (GLD) and silver (SLV) both reject the null and produce the most recent episodes in the commodity panel. GLD generates a GSADF episode from 26 August 2024 to the end of the sample (612 days), the longest continuously running episode in the entire commodity universe. SLV generates an episode from 25 September to 10 March 2026 (166 days). Both episodes originate in the second half of 2024, contemporaneous with the AI-equity explosive wave, but their attribution is straightforwardly macroeconomic: the gold rally of 2024–2025 is widely attributed to central bank reserve accumulation, geopolitical risk premia (Middle East, Ukraine), US fiscal trajectory concerns, and trade-policy uncertainty following tariff announcements in early 2025. The broad commodity benchmarks (DBC, GSG, PDBC) each produce two or three episodes concentrated in the 2021–2022 commodity supercycle and, for all three, a further episode originating in late January 2026.




Figure 45 . GSADF results, commodities, metals and broad commodity indices

5.2.4.5 SV-ADF results: aggregated commodity panels

The SV-ADF analysis over W1 (post-ChatGPT window) detects no episodes in the AI-relevant energy group, no episodes in the AI-relevant industrial metals group, and no episodes in either placebo group. The only W1 detections in the commodity panel are in the precious metals confound: GLD (origination 28 January 2025, post-bridge, 99 days) and SLV (origination 19 September 2025, post-bridge, 164 days). Applied over W2, the SV-ADF detects a 90-day episode in SOYB (origination 10 November 2020, post-bridge), consistent with the soybean price surge of that period, and 90-day episodes in the broad benchmarks DBC and PDBC both originating on 18 January 2022 and classified as post-bridge, consistent with the early-2022 Russia-Ukraine commodity shock. The SV-ADF results confirm the GSADF reading with the additional precision that the volatility-robust procedure makes explicit: the commodity episodes that survive the sustained-run filter in the post-ChatGPT window are confined to precious metals, which is the group with the strongest macroeconomic confounds and the weakest theoretical connection to AI infrastructure demand.



Figure 46 . SV-ADF results, commodities indices

5.2.4.6 Assessment of hypothesis H6

Neither form of H6 is supported by the data. The strong form — that AI-relevant commodities show higher explosive-episode rates than placebos — is directly contradicted:

both groups achieve 100% GSADF rejection, and the episode timing in both is dominated by the same 2021–2022 macroeconomic event sequence. The conditional form — that the post-Q4-2023 fraction of episode-days is higher in AI-relevant groups than in placebos — is also not supported: the AI-relevant energy group has zero post-Q4-2023 episodes, against one in the agriculture placebo (DBA, March–June 2024). The only post-Q4-2023 signal in an AI-relevant instrument is PPLT (platinum, June 2025), which is unconfirmed by the SV-ADF, unaccompanied by contemporaneous copper or palladium episodes, and has multiple competing attributions. The precious metals group, which does produce SV-ADF-confirmed post-Q4-2023 explosive episodes, is the confound group by construction, and its dynamics are more plausibly attributed to macroeconomic uncertainty than to AI infrastructure demand. The result is therefore treated as an informative non-finding: the commodity-market data, as represented by the available tradeable instruments over this sample window, do not provide affirmative evidence that speculative dynamics from AI equity markets propagated into real-resource markets. Whether this reflects a genuinely absent transmission channel, a channel that is too early-stage to have generated detectable price dynamics, or a limitation of the available commodity proxies cannot be resolved within the present framework.

5.2.5 Validation results

To assess the empirical credibility of the GSADF procedure before applying it to the AI equity universe, two historically documented bubble episodes and one deliberate negative control are examined in Figure 47. The validation logic follows Phillips, Shi and Yu (2015, hereafter PSY), who establish that a bubble-detection methodology should satisfy two symmetric criteria: it must reject the null of no explosive behavior in series where a bubble is independently corroborated in the literature (*positive validation*), and it must fail to reject that null in series where no explosive equity dynamic is theoretically expected (*negative control*). For each case, the left panel overlays the BSADF statistic path (red, right axis) against its 95% Monte Carlo critical-value path (grey dashed), with detected episode windows shaded in red and literature reference windows shaded in blue. The right panel quantifies the timing match as the *union coverage*, the fraction of reference-window calendar days that fall within the union of all detected episodes.




Figure 47 - Validation of the developed methods against historical bubble episodes

The first row applies the test to the NASDAQ Composite (^IXIC , 1990–2003). GSADF statistics of 7.04 substantially exceeds the 95% critical value of 2.47, rejecting the null of no explosive dynamics. The detected episodes cover 96% of the PSY reference window (origination $\approx$ June 1995, collapse $\approx$ September 2000), replicating the central empirical result of PSY (2015) and confirming that the implementation correctly date-stamps the dot-com bubble on its canonical index.

The second row serves as the negative control: the S&P 500 (^GSPC , 2003–2010) over the Global Financial Crisis period. The GFC originated in the US housing and credit markets rather than in broad equity valuations; explosive price dynamics are therefore theoretically absent from the S&P 500 price index over this window. The GSADF statistic of 1.68 falls below the critical value of 2.41, correctly yielding a failure to reject.

The third row applies the procedure to Bitcoin (BTC-USD, 2016–2022). With a GSADF statistic of 11.2 against a 95% critical value of 2.41, the null of no explosive dynamics is decisively rejected. Rather than broad approximation windows, the navy reference shading now corresponds to eight precisely dated bubble episodes drawn from two independent studies. Li et al. (2021) identify five episodes using a GSADF-based procedure on daily BTC data: 8–18 June 2016, 30 December 2016–15 January 2017, 22 February–7 March 2017, 27 April–5 July 2017, and 15 July 2017–9 January 2018, the last of these capturing the full extent of the celebrated 2017 speculative surge. Oldani, Bruno and Signorelli (2025) extend the chronology with three subsequent episodes: 24 November–8 December 2018, 11 May–28 June 2019, and 19 December 2020–13 March 2021, the last coinciding with the first peak of the post-pandemic institutional wave. The combined weighted coverage across all eight reference windows is 95%, where the denominator is the total number of reference-window calendar days. It should be noted that five of the eight windows span fewer than 21 trading days, shorter than the minimum-duration filter applied to episode detection in this study, so the GSADF is not expected to flag those windows individually; they are nonetheless included in the reference set to preserve fidelity to the source literature. The positive detection of the main 2017 episode and the 2020–2021 phase, combined with the overall rejection of the null, confirms that the procedure correctly identifies the broad explosive dynamics in Bitcoin pricing over the sample period.




Figure 48 - Validation of the SV-ADF method in python against the historical dot com bubble documented by Sarkar

The SV-ADF validation exercise applies the recursive coefficient-based test to the NASDAQ Composite over the dot-com episode, a period for which the ground truth is well established in the literature. The reference configuration — a 1982–2012 sample ($T \approx 7\,560$ trading days), origination run length $M = 360$, post-bridge confirmation run $R = 360$, and a two-year bridge window of 720 calendar days — exactly replicates the parameters used by Sarkar and Wells (2026) in their Figure 9, which dates bubble origination to April 1995 and collapse to September 2000.




Figure 49 - SV-ADF parameter sensitivity - NASDAQ Composite dot-com episode under two different parameter scenarios

One alternative configuration uses a shorter 1990–2003 sample ($T \approx 3\,480$ trading days) with parameters scaled proportionally to series length ($M = 175$, $R = 175$, bridge = 245 calendar days). As Figure 49 illustrates, the two configurations produce markedly different results despite being applied to the same underlying price history. Under the alternative parameterisation the detected episode is substantially compressed: the shorter bridge window (eight months rather than two years) means that the first significant intra-bubble correction (the NASDAQ drawdown of approximately 25% in mid-1996) drops the recursive SV-ADF statistic below the origination threshold within the bridge period, triggering an immediate bridge collapse and terminating the episode far earlier than the true market peak. The mechanism is straightforward: the bridge period is designed to tolerate transient dips that occur inside a genuine explosive phase; when the bridge is too short relative to the typical duration of intra-bubble corrections, it becomes a liability rather than a safeguard. The minimum-run parameter M compounds the effect because a smaller value allows origination to be declared on a shorter burst of explosive behaviour, after which the brief bridge fires and collapses the episode. The Sarkar parameters, calibrated on a thirty-year sample that encompasses the full expansion and the post-crash recovery, set M and the bridge at a length that is commensurate with the timescale of the dot-com cycle. This exercise underscores a known limitation of the SV-ADF framework: its episode-detection performance is sensitive to the joint choice of sample window, M , R and bridge, and these parameters should be calibrated to the expected duration of the underlying economic cycle rather than derived mechanically from sample length alone.

5.2.6 Robustness checks

Any empirical study that relies on a single parameterisation of its test procedures is vulnerable to the criticism that its findings are an artefact of those choices. This section stress-tests every major design decision in the analysis by varying one dimension at a time while holding all others fixed. Seven distinct checks are reported, grouped into two families. Checks A1 through A5 concern the Generalised Supremum ADF (GSADF) methodology and its interaction with the data transformation and test parameters. Checks B1 and B2 concern the Sarkar-Wells Stochastic Volatility ADF (SV-ADF) and its own parameterization. The objective is not to find the parameter setting that maximizes or minimizes rejection rates, but to establish which conclusions are stable, meaning they hold in the same direction across the range of plausible configurations, and which are sensitive to researcher discretion.

5.2.6.1 Scope

The parametric robustness analysis in this section evaluates the sensitivity of results pertaining to H1 (explosive dynamics in AI equities), H2 (capital-cycle ordering of upstream versus downstream segments), H3 (index dilution), and H4 (post-ChatGPT acceleration). Two research hypotheses are addressed outside this framework. H5 (historical validation of the SV-ADF methodology) is examined separately, where the parameter sensitivity analysis on the NASDAQ Composite dot-com episode constitutes the robustness argument. H6 (commodity-market propagation) is examined separately; because the strong-form GSADF test fails to discriminate between AI-relevant and placebo commodity groups, H6 evidence rests on the conditional timing test measuring the share of episode-days occurring after Q4 2023. No systematic parameter sensitivity analysis has been conducted for that metric, and H6 conclusions should accordingly be read as indicative rather than as a robustness-verified finding.

5.2.6.2 Baseline

The baseline against which all variants are compared is the main-analysis configuration: the full sample period from 1 January 2020 to 1 May 2026, GSADF applied to price levels with the Phillips-Shi-Yu minimum-window formula, episode detection with a 60-trading-day minimum duration and a 42-day merge gap, and a 95% Monte Carlo critical value. Under this baseline, 70.3% of equity tickers reject the null of no explosive dynamics, with an upstream

segment rate of 85.7% and a downstream rate of 62.5%. No index rejects the null. The mean number of detected GSADF episodes per equity ticker is 1.51.

The SV-ADF baseline applies the method across two non-overlapping windows that straddle the ChatGPT inflection point. Window W1, covering the post-ChatGPT period from 1 November 2021 to 1 May 2026, yields an equity episode rate of 37.8%, with 57.1% of upstream tickers and 12.5% of downstream tickers recording a detected bubble episode; the upstream-downstream differential is accordingly preserved. Window W2, covering the pre-ChatGPT period from 1 January 2020 to 31 October 2022, yields a markedly lower overall equity episode rate of 5.7%, with upstream and downstream rates of 7.1% and 12.5% respectively - the only configuration in the robustness suite where the downstream rate nominally exceeds the upstream rate, driven by a single detected episode on each side of a much smaller denominator; read together with W1, this indicates the upstream-over-downstream gradient is itself a post-ChatGPT phenomenon. In both windows, episode detection uses a minimum origination exceedance of $M = 42$ consecutive trading days above the $\log(\tau)/10$ threshold, a minimum post-bridge collapse period of $R = 21$ consecutive trading days below the $\log(\tau)/2$ threshold, and a bridge window of 90 calendar days. The Panel B heatmap benchmarks all SV-ADF robustness variants against the W1 configuration with default parameters, which constitutes the primary SV-ADF reference for the post-ChatGPT analysis.

The seven checks vary across four distinct dimensions: sample window (A1), episode-dating parameters (A2, B2), statistical thresholds (A3), and model specification (A4, A5, B1). The most sensitive finding is A5, the log-price transformation, which causes the upstream GSADF rejection rate to collapse from 85.7% to 14.3% and reverses the upstream-downstream ordering (14.3% vs 25.0%), indicating that the levels-based results partly reflect convex price-level growth rather than a pure departure from a log-linear benchmark. All other checks preserve the core qualitative conclusions: the upstream-over-downstream ordering (H2) is maintained in 22 of 25 configurations; the index non-rejection (H3) holds in 18 of 19 applicable configurations; the SV-ADF results are directionally stable but not perfectly invariant across the six B1 window variants (equity rate ranges 5.7%-43.2%) and the five B2 threshold configurations (37.8%-54.1%), with the upstream rate exceeding the downstream rate in every post-ChatGPT (W1) variant. The complete per-configuration metrics are consolidated in `rob_summary.csv` and visualised in the stability heatmap (Figure 51).

Table 9 - Robustness Check inventory

Check	Short name	Repository files (rob_prefix)	Baseline → variants	What it tests
GSADF checks				
A1	Sub-period split	A1_pre_ChatGPT	Full sample 2020–2026 →	Whether explosive dynamics are temporally concentrated after the November 2022 ChatGPT launch; the direct operational test of H1
		A1_post_ChatGPT	pre-ChatGPT (Jan 2020 – Oct 2022) vs. post-ChatGPT (Nov 2022 – May 2026); new MC critical values computed for each sub-period length	
A2	Episode filter grid	A2_baseline (md 60, mg 42)	md: 21 / 42 / 60 / 60 / 90 days	Whether the upstream–downstream episode ordering and overall rejection rates survive more lenient or more restrictive noise-filter settings; GSADF paths loaded from cache, not re-computed
		A2_md21_mg21 A2_md42_mg21	mg: 21 / 21 / 42 / 60 / 42 days (5 configurations)	
		A2_md60_mg60 A2_md90_mg42		
A3	Significance level	A3_alpha90	MC critical value quantile: 90th	Whether rejection decisions hold across conventional significance levels; uses pre-stored quantiles from the MC cache, no re-simulation required
		A3_alpha95_baseline	/ 95th / 99th percentile	
		A3_alpha99		
A4	Minimum-window variants	A4_mw080	PSY formula scale factor: 0.80×	Whether the initial recursion window length drives rejection outcomes; tests the sensitivity to the Phillips-Shi-Yu (2015) formula choice
		A4_mw100	/ 1.00× / 1.20× (mw ≈ 77 / 96 / 115 for T = 1,842); new GSADF runs and MC simulations per variant	
		A4_mw120		
A5	Log-price transformation	A5_logprices	Price series: levels → natural logarithm; no new MC simulation required	Whether explosive dynamics reflect genuine departure from a log-linear random walk or are partly attributable to the convex curvature inherent in rapid exponential growth at price levels; the most important single sensitivity check
SV-ADF checks				
B1	Window boundary shift	B1_W1_early	W1 start: Aug 2021 / Nov 2021 / Feb 2022	Whether SV-ADF episode detection is sensitive to the exact placement of the pre/post-ChatGPT split date
		B1_W1_baseline		
		B1_W1_late	W2 end: Jul 2022 / Oct 2022 / Jan 2023 (6 combinations, W1 and W2 varied independently); SV-ADF re-estimated per variant	
		B1_W2_early		
		B1_W2_baseline		
		B1_W2_late		

B2	Episode threshold grid	B2_up21_dn10	up: 21 / 21 / 42 / 42 / 63 days dn: 10 / 21 / 21 / 42 / 21 days (5 configurations); coefficient paths loaded from cache, no re-estimation required	Whether SV-ADF episode counts are sensitive to the sustained-run thresholds governing origination confirmation and collapse confirmation
		B2_up21_dn21		
		B2_baseline (up 42 , dn 21)		
		B2_up42_dn42		
		B2_up63_dn21		

Summary across all 25 configurations: *rob_summary.csv*. Visual synthesis: *rob_stability_heatmap.png*.

Error! Reference source not found. below collects the key metrics for every variant across all seven checks. It is the primary reference for the discussion that follows.

Baseline values in the header row. Upstream = segments 1-5; Downstream = segments 11-13. Episode rate in A2 is based on dated episodes rather than formal GSADF rejections; the two measures differ because the global supremum can exceed the CV without producing an episode that passes the duration filter. SV-ADF metrics refer to the W1 (post-ChatGPT) window unless otherwise stated.

Table 10 - Robustness Results summary

Check	Variant		Equity rate	Upstream rate	Downstream rate	Index rate
Baseline	—		70.3%	85.7%	62.5%	0%
A1	Pre-ChatGPT	sub-sample	20.0%	28.6%	12.5%	0%
A1	Post-ChatGPT	sub-sample	56.8%	71.4%	50.0%	0%
A2	Loosest	filter (md=21, mg=21)	91.9% (ep.)	100%	75.0%	100%
A2	Baseline	filter (md=60, mg=42)	75.7% (ep.)	85.7%	75.0%	0%
A2	Strictest	filter (md=90, mg=42)	59.5% (ep.)	78.6%	37.5%	0%
A3	$\alpha = 90\%$		81.1%	92.9%	75.0%	33%
A3	$\alpha = 99\%$		59.5%	85.7%	37.5%	0%
A4	mw = 0.80× / 1.00× / 1.20×		64.9%	85.7%	50.0%	0%
A5	Log-price series		27.0%	14.3%	25.0%	0%
B1	All window variants	6 W1/W2	5.7%–43.2%	7.1%–57.1%	12.5%–25.0%	—

<i>B2</i>	All 5 threshold configurations	37.8%–54.1%	57.1%–71.4%	12.5%–25.0%	—
-----------	--------------------------------	-------------	-------------	-------------	---

5.2.6.3 A1 — Sub-period split

Running the full GSADF separately on the pre-ChatGPT sub-sample (January 2020 - October 2022) and on the post-ChatGPT sub-sample (November 2022 - May 2026) provides the most direct test of H1. The pre-ChatGPT rejection rate of 20.0% is roughly one-third of the post-ChatGPT rate of 56.8%. The upstream-downstream ordering (H2) and the index non-rejection (H3) hold in both sub-periods. The full-sample baseline rate of 70.3% sits between both sub-period rates because the supremum taken over the longer window naturally picks up dynamics from either phase; the sub-period split isolates what is happening within each regime and confirms that explosive dynamics are close to three times more prevalent after the ChatGPT launch.

5.2.6.4 A2 — Episode filter grid

The baseline episode-detection filter (minimum duration 60 days, merge gap 42 days) is a researcher choice. Testing five configurations from most lenient (21/21) to most restrictive (90/42) shows that the episode-based equity rate ranges from 59.5% to 91.9%. The critical observation is the loosest filter: at a 21-day minimum, every index series also produces detected episodes (index rate 100%), revealing that raw BSADF crossings below 21 days are predominantly Monte Carlo noise rather than genuine explosive dynamics. This confirms the necessity of the minimum-duration filter as a noise-control device, and it is the one configuration in the entire suite where H3 fails outright. At every other configuration, no index rejects and the upstream-downstream direction is maintained, so the core qualitative conclusions are stable across all meaningful filter choices.

5.2.6.5 A3 — Significance level

Comparing the 95% baseline with the 90th and 99th Monte Carlo quantiles requires no new simulations as the quantiles are pre-stored from the main cache. At alpha = 90%, the equity rate rises to 81.1% and one index crosses the rejection threshold (index rate 33.3%); this crossing disappears at alpha = 95% and is absent at alpha = 99%. At alpha = 99%, the equity rate falls to 59.5% and the downstream rate drops to 37.5%, but upstream holds at 85.7%. The widening gap between upstream and downstream as significance tightens is itself informative:

upstream explosive dynamics are strong enough to survive the strictest conventional threshold while downstream dynamics are more borderline, consistent with a cycle that is most intense at the supply-constrained infrastructure layer.

5.2.6.6 A4 — Minimum-window variants

Perturbing the Phillips-Shi-Yu initial recursion window by $\pm 20\%$ ($mw \sim 77, 96$, and 115 observations for $T = 1,842$) produces identical outcomes in all three cases: 64.9% equity rate, 85.7% upstream, 50.0% downstream, zero index rejections. The exact replication across three different window sizes indicates that the rejecting tickers exceed the critical value by a wide margin and the non-rejecting tickers fall short by a wide margin; no firm is on the boundary.

5.2.6.7 A5 — Log-price transformation

Applying the GSADF to log-transformed price series rather than price levels remains the most consequential specification choice in the analysis. The overall equity rejection rate falls from 70.3% to 27.0% , and the upstream rejection rate collapses from 85.7% to 14.3% - with the upstream-downstream ordering reversing outright (14.3% vs 25.0%) under this specification. This is the single red finding in the robustness exercise and requires honest interpretation. It does not mean the main results are wrong; it means they are partly dependent on the choice between testing in levels versus log levels. The log-price specification compresses the upward convex curvature that characterizes rapid exponential price growth. If the theoretical model is that investors form expectations in log-price space, the 27.0% rate is the more theory-consistent benchmark; if investors respond to price levels directly, the levels-based results are appropriate. This thesis reports both, bases the main narrative on levels (consistent with Basele, Phillips and Shi (2025) and other recent GSADF applications to technology equities), and explicitly flags the sensitivity.

5.2.6.8 B1 — SV-ADF window boundary shift

Shifting the W1/W2 split date by three months in either direction - making W1 start in August 2021 or February 2022, and W2 end in July 2022 or January 2023 - produces small, consistent movements rather than the complete invariance seen in the original panel. The three W1 variants give episode rates of 43.2% (early), 37.8% (baseline), and 43.2% (late); the three W2 variants are flat at 5.7% across all boundary choices. The upstream rate exceeds the downstream rate in every W1 variant; in every W2 variant, the ordering nominally reverses

(upstream 7.1% vs downstream 12.5%), for the same single-episode/small-denominator reason discussed under the baseline. The $W1 > W2$ ordering that underlies H1 and H4, and the upstream $>$ downstream direction within W1 that underlies H2, are both unaffected by the exact boundary choice.

5.2.6.9 B2 — SV-ADF episode threshold grid


Varying the minimum origination run (21, 42, 63 days) and the minimum collapse run (10, 21, 42 days) across five configurations gives episode rates of 54.1% for the two most permissive origination settings (21 days) and 37.8% for the three configurations using the 42-day or 63-day origination run. The upstream rate exceeds the downstream rate in every configuration (71.4% vs 25.0%, or 57.1% vs 12.5%). No configuration reverses the upstream-downstream ordering, and the split falls cleanly along the origination-run-length dimension rather than the collapse-run-length dimension, which has no visible effect at all (up42/dn21, up42/dn42, and up63/dn21 all produce the identical 37.8%/57.1%/12.5% result).

5.2.6.10 Synthesis

Heatmap

Figure 51 - Robustness Check Heatmap provides a quantitative summary of how each robustness configuration deviates from the baseline across six key metrics, organised as two stacked panels. Panel A covers the fourteen GSADF configurations (checks A1 through A5), benchmarked against the $\alpha = 95\%$ baseline. Panel B covers the eleven SV-ADF configurations (checks B1 and B2), benchmarked against the default episode-threshold setting (B2, up=42 days, dn=21 days).

Figure 50 - Robustness Check Heatmap



Each row of the heatmap corresponds to one configuration; each column to one metric. The six columns are: (i) equity rejection rate [H1], representing the share of AI equity tickers for which the null of no explosive behaviour is rejected; (ii) upstream segment rejection rate [H2↑]; (iii) downstream segment rejection rate [H2↓]; (iv) index rejection rate [H3 control], measuring detection in broad market benchmarks; (v) H2 holds, a binary cell indicating whether the upstream rate exceeds the downstream rate in that configuration; and (vi) H3 holds, a binary

cell indicating whether the equity rate exceeds the index rate. For the four rate columns, cells are colour-coded by absolute percentage-point deviation from the panel reference: deviations of eight percentage points or fewer are shaded green (stable); deviations between eight and eighteen percentage points are shaded orange (moderate sensitivity in magnitude, direction preserved); deviations exceeding eighteen percentage points or a directional reversal are shaded red (qualitatively different conclusion). For the two binary columns, green signals that the directional ordering matches the reference and red signals a reversal. Baseline variant configurations, those that replicate the main-analysis parameter choices exactly, are highlighted in teal within the panel to serve as internal reference anchors; the primary panel reference is rendered in dark blue at the bottom of each panel.

Reading the heatmap: a predominantly green row indicates that the configuration is a minor perturbation of the baseline with no qualitative impact on conclusions. A row with orange cells indicates magnitude sensitivity without directional change. Red cells in the rate columns denote a quantitatively large departure; red cells in the binary columns denote a direction reversal, which is the more consequential type of departure from the perspective of hypothesis support. The upstream and downstream columns should be read together: as long as the upstream value exceeds the downstream value within the same row, H2 is supported for that configuration regardless of absolute level.

Panel A (GSADF checks) reveals a structured pattern. The A4 rows are uniformly stable across all columns, confirming that the minimum-window choice is not a material driver of any reported metric. The A3 rows display a clean gradient: relaxing the threshold to $\alpha = 90\%$ shifts the equity and index columns to moderate sensitivity while H2 and H3 remain intact, whereas tightening to $\alpha = 99\%$ shifts the downstream rate to a large deviation while preserving the directional ordering throughout. The A2 rows show a similar pattern except at the most permissive setting (21-day minimum), which produces a large deviation in both the equity and index columns and is the one configuration in the entire panel where H3 fails - a minimum-duration filter short enough to admit noise detections does so indiscriminately across all asset classes, erasing the equity-index contrast that underpins H3. The A1 rows follow the expected structural pattern: the pre-ChatGPT sub-sample shows large deviations in the equity, upstream, and downstream rates by construction, while the post-ChatGPT sub-sample shows only moderate deviation and remains directionally consistent with the full-sample baseline. The A5 row (log-price transformation) is the most visually prominent departure in Panel A: the equity rate falls to 27.0%, the upstream rate falls to 14.3%, and - uniquely among all 25

configurations - the upstream-downstream ordering itself reverses (14.3% vs 25.0%), so H2 fails here.

Panel B (SV-ADF checks) presents a more stable picture than Panel A, though less perfectly invariant than in the original panel. The three W1 boundary variants are close to the reference across all rate columns and all confirm H2. The three W2 boundary variants show large deviations in the equity and upstream columns (rates of 5.7% and 7.1% against a reference of 37.8% and 57.1%), and are the only configurations in Panel B where H2 fails - for the same small-denominator reason noted throughout this section. The H3 column is not applicable for any B1/B2 configuration, since the SV-ADF estimation windows do not include index benchmark tickers. The five B2 configurations range from stable (three configurations identical to the reference) to moderate (the two 21-day-origination configurations, at 54.1%/71.4%/25.0%), with H2 holding in every case.

Hypothesis scorecard

Figure 51 - Robustness Scorecard translates the quantitative heatmap into a direct hypothesis-level assessment. Each of the 25 configurations is evaluated against four criteria: H1 (equity rejection rate exceeds 45% for GSADF checks or 35% for SV-ADF checks), H2 (upstream rate exceeds downstream rate), H3 (equity rate exceeds index rate, where applicable), and H4 (post-ChatGPT rate exceeds pre-ChatGPT rate, assessed via the A1 split only).

Robustness scorecard — hypothesis-level assessment across all configurations
H1: AI equity explosivity · H2: capital-cycle pattern (upstream > downstream) · H3: index dilution · H4: post-ChatGPT acceleration

Configuration	H1 Explosivity	H2 Capital cycle	H3 Index dilution	H4 ChatGPT break (A1 only)	Verdict
A1a Sub-period: pre-ChatGPT (2020-Oct 2022)	✗	✓	✓	—	✓✓
A1b Sub-period: post-ChatGPT (Nov 2022-2025)	✓	✓	✓	✓	✓✓
A2a Episode filter: md=21d, gap=21d	✓	✓	✗	—	○
A2b Episode filter: md=42d, gap=21d	✓	✓	✓	—	✓✓
A2c Episode filter: baseline (md=60d, gap=42d)	✓	✓	✓	—	✓✓
A2d Episode filter: md=90d, gap=42d	✓	✓	✓	—	✓✓
A2e Episode filter: md=60d, gap=60d	✓	✓	✓	—	✓✓
A3a Significance level: $\alpha = 90\%$	✓	✓	✓	—	✓✓
A3b Significance level: $\alpha = 95\%$ [baseline]	✓	✓	✓	—	✓✓
A3c Significance level: $\alpha = 99\%$	✓	✓	✓	—	✓✓
A4a Min window: $0.8 \times \text{PSY}$	✓	✓	✓	—	✓✓
A4b Min window: $1.0 \times \text{PSY}$ [baseline]	✓	✓	✓	—	✓✓
A4c Min window: $1.2 \times \text{PSY}$	✓	✓	✓	—	✓✓
A5 Log-price transformation	✗	✗	✓	—	✗
B1a Window W1 start: early (Aug 2021)	✓	✓	—	—	✓✓
B1b Window W1 start: baseline (Nov 2021) [baseline]	✓	✓	—	—	✓✓
B1c Window W1 start: late (Feb 2022)	✓	✓	—	—	✓✓
B1d Window W2 end: early (Jul 2022)	✗	✗	—	—	✗
B1e Window W2 end: baseline (Oct 2022) [baseline]	✗	✗	—	—	✗
B1f Window W2 end: late (Jan 2023)	✗	✗	—	—	✗
B2a SV-ADF thresholds: M=21, R=10	✓	✓	✓	—	✓✓
B2b SV-ADF thresholds: M=21, R=21	✓	✓	✓	—	✓✓
B2c SV-ADF thresholds: baseline (M=42, R=21) [baseline]	✓	✓	✓	—	✓✓
B2d SV-ADF thresholds: M=42, R=42	✓	✓	✓	—	✓✓
B2e SV-ADF thresholds: M=63, R=21	✓	✓	✓	—	✓✓

* grey ✗ = expected failure by design (pre-boom sub-period) · ✓✓ = fully robust · ✓ = majority hypotheses hold · ○ = partial · ✗ = failure
Summary counts: H1 ✓ in 20/25 · H2 ✓ in 21/25 · H3 ✓ in 18/19 configurations · H4 confirmed (post=0.57 > pre=0.20)

Figure 51 - Robustness Scorecard

The scorecard gives the following counts. H1 holds in 20 of 25 configurations. Of the five non-confirmations, four are expected by design - the pre-ChatGPT sub-sample (A1) and all three pre-ChatGPT SV-ADF window variants (B1_W2), since H1 is specifically a post-ChatGPT claim - and one is a genuine, disclosed sensitivity: A5 (log prices, equity rate 27.0%). H2 holds in 22 of 25 configurations; the three exceptions are the pre-ChatGPT SV-ADF window variants (B1_W2), where a single episode on each side of a much smaller downstream denominator produces a nominally higher downstream rate - a reversal confined entirely to the pre-ChatGPT period, plus the log-price specification (A5) as a separate, fourth-style failure driven by a different mechanism. H3 holds in 18 of the 19 applicable configurations (the six B1 window-boundary checks are excluded from this count, since no index ticker is tested in the SV-ADF windows); the one exception is the most permissive GSADF episode filter (A2, 21-day

minimum), which inflates index detections to 100% and is explicitly flagged as a noise-filter failure rather than a substantive one. H4 is confirmed by the A1 sub-period split: the post-ChatGPT equity rejection rate (56.8%) substantially exceeds the pre-ChatGPT rate (20.0%).

Overall assessment

The core directional findings - explosive dynamics concentrated in the post-ChatGPT period, more prevalent upstream than downstream, and largely absent from broad index benchmarks - survive across 22 of 25 configurations on H2 and 18 of 19 applicable configurations on H3, making these the most robustly supported results of the analysis. H1 is confirmed in 20 of 25 configurations, with four of the five non-confirmations fully explained by structural design choices (the checks deliberately probing the pre-ChatGPT period) rather than methodological fragility. The principal source of sensitivity remains the price-transformation choice (A5): applying a log specification substantially reduces detected explosivity and is the only configuration in the entire suite that reverses the upstream-downstream ordering outright. This is a deliberate econometric choice, disclosed rather than concealed: the main analysis follows the levels convention established by Phillips et al. (2015) and adopted by Sarkar and Wells (2026). A secondary, smaller source of sensitivity is the pre-ChatGPT SV-ADF window (B1_W2), where the upstream-downstream ordering nominally reverses due to a small-sample effect rather than a substantive one, and which is confined to the pre-period the ordering is not expected to hold in regardless. All remaining parameter dimensions (episode-filter grid, PSY minimum window, and SV-ADF episode thresholds) leave the directional findings undisturbed.

6. Discussion: Insights and limits of the test

6.1 Interpreting and contextualizing findings

The empirical results of Chapter 5 establish three findings that together form the analytical core of this thesis. First, most AI-exposed equities exhibit explosive price dynamics detectable at the 95% confidence level by the GSADF procedure of Phillips, Shi and Yu (2015), with a rejection rate of 70.3% among the main equity panel and a rate of 91.9% when the less demanding criterion of at least one dated BSADF episode is applied. Second, these dynamics are temporally concentrated in the post-ChatGPT period: the pre-ChatGPT subsample rejection rate of 20.0% rises to 56.8% in the post-ChatGPT window under the GSADF, and the SV-ADF W1 (post-ChatGPT) episode rate of 37.8% is materially higher than the W2 rate of 0.0%. Third, the episode timing propagates downstream along the value chain in a sequence that mirrors the structural capital-deployment sequence documented in Chapter 3: chip design (Q1-Q2 2023), foundry and memory (Q1 2024), cloud and enterprise (Q3-Q4 2024), application software (Q4 2024 onwards).

Interpreting these findings requires that the distinction between statistical explosiveness and theoretical bubble formation, introduced in §2.2.2, be applied with care. Rejection of the GSADF null establishes that the price process exhibits a locally non-random-walk autoregressive structure consistent with the mild explosiveness that Phillips, Wu and Yu (2011) associate with bubble dynamics; it does not, in isolation, establish that prices have deviated from fundamental value in the sense of any of the theoretical frameworks reviewed in Chapter 2. Two non-bubble explanations for the observed explosiveness deserve explicit consideration before a bubble interpretation is advanced.

The first is rational fundamental revision. As Pástor and Veronesi (2009) demonstrate in their general-equilibrium model of technological revolutions, the arrival of a genuinely transformative technology triggers a rational shift in growth expectations that can, under uncertainty about the technology's ultimate productivity impact, generate explosive price dynamics that are observationally indistinguishable from a bubble in any time-series-based test. The AI cycle, as documented in Chapter 3, does involve a genuine technological discontinuity (large-scale language models whose capabilities accelerated sharply after November 2022) and several of the explosive episodes documented here are at least partially consistent with a rational revision of long-run cash flow expectations. Nvidia's revenue

trajectory (from USD 5B quarterly data-centre revenue in Q1 2023 to approximately USD 50B by Q4 2024) constitutes a real fundamental shift, and NVDA's explosive episode coincides with this shift.

The second is the log-price sensitivity identified in robustness check A5: when the GSADF is applied to log-transformed price series rather than price levels, the upstream rejection rate falls from 85.7% to 14.3%, and the upstream-downstream ordering itself reverses (14.3% versus 25.0%). This finding indicates that a portion of the explosiveness detected in the levels-based analysis reflects the convex curvature inherent in any sufficiently rapid exponential growth trajectory rather than a departure from the log-linear random walk that standard valuation theory implies. A strictly log-linear investor who revised Nvidia's expected growth rate upward in January 2023 would, under continuous compounding, generate a price series that the levels-based GSADF detects as explosive even though no overvaluation relative to the revised fundamental was ever present. The log-price result is not decisive: as argued in §5.2.6, the choice between levels and log-levels depends on assumptions about the bubble formation process, and the recent literature applying GSADF to high-growth technology equities largely uses levels, but it disciplines the confidence that can be placed in the levels results: the upstream rejections, in particular, should be interpreted as explosive dynamics whose source includes but may not be limited to speculative excess.

With these qualifications acknowledged, the conjunction of three features of the empirical results shifts the interpretation toward speculative rather than purely fundamental dynamics. The first feature is the upstream-downstream temporal sequence. A story based purely on changes in fundamentals would predict that all segments are repriced at the same time, since the technology's impact on productivity is common knowledge. Instead, the data show a staggered pattern, whose timing follows the physical capacity-deployment cycle documented in Chapter 3. The second feature is the divergence between index-level non-rejection and firm-level majority rejection. Under efficient markets, the explosive dynamics present at the firm level should aggregate to the index level unless they are driven by idiosyncratic speculative episodes whose origination dates are distributed uniformly across calendar time; the fact that they do not aggregate may be a reassuring sign that the market exuberance within certain segments is compensated by opposite trends in others. The third feature is the magnitude and duration of the largest individual episodes: NVDA's 704-day episode, PLTR's 505-day episode, and the sustained Oracle episodes around the Stargate announcement are the

statistically sustained, directional deviation from the random-walk benchmark that persists across regimes of both rising and declining price momentum.

Taken together, and mapped against Chapter 3, the findings support the following conclusion: the AI capital cycle exhibits econometrically identifiable explosive price dynamics across a majority of its equity-market representations, these dynamics are concentrated in the post-ChatGPT period and propagate along the value chain in a pattern consistent with a supply-side capital-deployment narrative, and they are attributable, in proportions that the test cannot separately identify, to a combination of genuine fundamental revision, convex price-level effects, and speculative excess in the sense understood by the explosive-root literature. The balance of these three components differs by segment: the upstream chip design episodes, which coincide most directly with documented revenue acceleration, contain the largest fundamental component; the application-layer episodes, which originate later and in the context of low demonstrated profitability, contain the largest speculative component.

6.2 Where the tests fall short

The methodological framework developed in this thesis has three classes of limitation that condition the strength of the interpretive conclusions just drawn. The first and most significant is the confinement of all explosive-root tests to publicly traded, continuously priced assets. As §2.2.2 established, the GSADF, SADF, and SV-ADF procedures require a sufficiently long time series of regularly observed prices; they produce no output for assets that are valued only at discrete, irregular intervals. This restriction excludes precisely the market segment in which a large share of the AI capital cycle is concentrated: the equity valuations of privately held foundation-model developers. At the time of writing, the two largest pure-play AI model developers (OpenAI and Anthropic) are privately held, with valuations available only from secondary-transaction pricing and periodic disclosure to investors. The inability to apply explosive-root tests to these valuations is a major gap in the analysis's exhaustiveness. The foundation-model developers occupy the layer at which the circular financing arrangements documented in Chapter 3 are most tightly concentrated, the layer whose revenue is most endogenous to the capital cycle's own mechanics, and the layer at which any future demand shortfall would be most immediately felt. This gap pertains to the influx of capital as well into the broader AI-startup and venture capital markets. A bubble diagnosis that is confined to the publicly listed portions of the AI

equity universe is necessarily an incomplete diagnosis of the cycle, because the agents most centrally exposed to the speculative dynamics in question are not subject to the analytical tools applied here.

The second class of limitation concerns the interpretation of episode dates in the context of heteroskedastic markets. The robustness analysis of §5.2.6 establishes that the SV-ADF, explicitly designed to be robust to time-varying volatility, detects substantially fewer episodes than the GSADF and agrees with it on only 68% of verdicts (25 of the 37 equity tickers). While the preferred interpretation of this disagreement is an improvement in the diagnostic precision rather than error on the SV-ADF's part, the gap raises a legitimate question about the GSADF episodes that are not confirmed by the SV-ADF: are they genuine explosive departures from the random walk, or are they volatility clusters masquerading as explosiveness? For the long, sustained episodes (e.g. NVDA's 704 days, PLTR's 505 days) the question is less pressing because a 700-day continuous deviation is hardly a transient heteroskedastic shock. For the shorter GSADF episodes in names like META, ANET, and CSCO, which are between 90 and 180 days and not confirmed by the SV-ADF, greater interpretive caution is warranted. The two-method triangulation strategy, which the analysis adopts explicitly, provides a partial resolution: findings that are confirmed by both methods can be stated with greater confidence than findings that rest on the GSADF alone.

The third limitation is the joint-hypothesis problem that applies to all indirect bubble tests, first identified by Gürkaynak (2008) in the context of econometric bubble tests generally and directly relevant here. Rejection of the GSADF null establishes that the price process is locally explosive; it does not distinguish between explosiveness driven by a bubble and explosiveness driven by a model-consistent revision in discount rates or growth expectations. The log-price robustness check (A5) shows that the level of explosiveness detected is sensitive to the specification of the price-level process, which means that the test outcome implicitly depends on assumptions about the data-generating process that cannot be fully tested within the framework itself. This is not a deficiency unique to this analysis; it is a structural feature of all indirect bubble-detection methods, including those of Shiller (1981) and Diba and Grossman (1988), and has been explicitly acknowledged in the empirical literature since the Gürkaynak (2008) survey. The practical consequence for this thesis is that the econometric results are best understood if accompanied by the evidence of Chapter 3: the circular financing arrangements, the financing-gap calculation of §3.2.1, and the leverage and

risk-transfer documentation of §3.3.3. It is the convergence of the econometric and the structural evidence, rather than either alone, that motivates the bubble characterization of §6.1.

6.3 Integrating circular financing and debt into the analysis of results

The mapping account of Chapter 3 and the empirical results of Chapter 5 are analytically separable but interpretively interdependent. The structural account documents a set of commercial and financial arrangements, reciprocal equity investments, multi-year compute commitments, GPU-collateralised private credit, off-balance-sheet SPV debt, whose collective effect is to render a portion of the reported revenue at each node of the AI value chain partially endogenous to the capital deployed by an adjacent node. The empirical results document explosive price dynamics that originate in the chip design layer and propagate downstream along the capital-deployment cycle. The question of whether these two bodies of evidence point to the same underlying phenomenon requires an integration that neither can perform on its own.

The strongest point of contact between the two bodies of evidence is temporal. The chip design episodes (NVDA from March 2023, AVGO from May 2023) originate in the months immediately following the ChatGPT public release and the subsequent acceleration of hyperscaler capital expenditure commitments documented in §3.2.1. The hyperscaler cloud episodes (ORCL from June 2023 and more prominently from September 2024, GOOGL from September 2025) originate later, approximately contemporaneous with the major reciprocal investment and compute-commitment announcements documented in §3.3.2. The Oracle episodes in particular show a tight temporal alignment with the Stargate consortium announcement (January 2025) and Oracle's USD 300B OpenAI compute commitment (September 2025): the GSADF episode originating in September 2024 predates the public Stargate announcement by approximately four months, suggesting that the market might have began pricing the arrangement before its public disclosure, while the June 2025 episode follows the formal contract disclosure closely. This alignment is consistent with hypothesis H2's claim that explosive episodes coincide with the intensification of structural drivers, though the four-month anticipation window requires care in causal attribution: the market may have been responding to broader AI infrastructure signals rather than to any specific arrangement.

The power-and-grid segment adds another independent point of temporal contact. Vistra (VST) and Constellation Energy (CEG), the two names added to the universe to represent electricity generation tied to data-center demand, both originate their SV-ADF W1 episodes in January-February 2024, within days of the NVDA and AVGO chip-design episodes rather than following them with a lag. This near-simultaneity is informative: it indicates that the electricity-generation narrative was priced in at essentially the same moment as the chip-design re-rating, consistent with the account in §3.2.2 that power and grid constraints emerged as a binding, publicly documented input from early 2024 onward, rather than as a delayed downstream consequence of the compute buildout.

The debt-financing dimension adds another layer to the structural account. The hyperscaler corporate bond issuances documented in §3.3.3 (approximately USD 121B in 2025, a fourfold increase over the 2020-2024 average) mark a significant shift in financing the AI capital cycle from internally generated cash flow to external debt. The timing of this shift is broadly contemporaneous with the peak explosive-episode density in the hyperscale cloud segment (late 2024), suggesting that the equity market's detection of explosive dynamics and the credit market's accommodation of expanded debt issuance are two facets of the same event: a market-wide revision of AI revenue and growth expectations that makes both equity re-rating and incremental debt issuance feasible. The off-balance-sheet SPV structures documented in §3.3.3 are particularly relevant in this context: by removing debt from the consolidated balance sheets of the principal hyperscalers, these structures reduce the observable leverage ratios of the most visible AI participants, potentially attenuating the market's perception of financial fragility and extending the speculative phase. The mechanism is an exact analogue of the vendor-financing arrangements documented by Sidak (2003) for the telecommunications precedent: just as Lucent's extension of customer loans financed a portion of its own revenue, the hyperscalers' SPV-mediated debt financing of data-center buildout supports an infrastructure capacity that generates the reported backlog commitments on which AI equity valuations rest.

The practical implication for the structural test of hypothesis H2 is that the explosive dynamics identified in the data are not merely a financial phenomenon; they are the price-market expression of a capital allocation process whose mechanics are documented in the structural chapter. A bubble characterisation that rests only on the statistical detection of

explosive price dynamics would, as Gürkaynak (2008) cautions, be vulnerable to the alternative explanation that the dynamics reflect rational fundamental revision. A bubble characterisation that rests only on the structural documentation of circular financing would be vulnerable to the objection that circular financing is not inherently pathological (many industries involve bilateral supply relationships and equity cross-holdings) and that the question of whether such arrangements constitute a destabilising feedback loop depends on whether the external revenue base is sufficient to sustain the cycle. The conjunction of the two bodies of evidence — explosive equity dynamics with a timing pattern consistent with the structural mechanics, in a sector whose financing structure is demonstrably self-referential and whose revenue base is demonstrably insufficient to close the capital expenditure gap at the observed scale — is more informative than either alone. It does not establish that a collapse is certain or imminent; it establishes that the conditions for a bubble characterisation in the sense of this thesis are met.

6.4 Implications for bubble diagnostics in systemic innovation sectors

Two lessons generalize beyond the AI case, to any sector undergoing a comparable general-purpose technological transition. First, index-level surveillance is an unreliable proxy for firm-level speculative behaviour whenever explosive dynamics concentrate in a subset of high-momentum constituents: the diversification-dilution mechanism documented in §6.1 will recur wherever a technology re-rates some but not all of its exposed firms, so any monitoring framework should be built at the firm level and aggregated upward rather than read off an index. Second, where price series plausibly combine genuine fundamental growth with speculative excess — the exact combination the Pástor and Veronesi (2009) model produces endogenously — no single test statistic can decompose the two; agreement across methodologically distinct procedures is what licenses a defensible candidate list. Applied here, the intersection of the GSADF and SV-ADF results yields fourteen tickers whose explosive dynamics are confirmed by both methods: LRCX, NVDA, AVGO, MU, TSM, ASX, WDC, STX, CIEN, CEG, VST, GOOGL, ORCL, and PLTR — the most methodologically conservative bubble candidate list the dataset supports.

7. Conclusions and recommendations

The analysis pursued in this thesis contributed to the diagnosis whether AI-exposed asset prices exhibit the statistical signature of explosive, speculative dynamics, and whether any such dynamics could be related to the structural and financing characteristics of the sector. This concluding chapter draws together the answers, states the verdict the evidence supports, and sets out what follows from it for practice and for future research.

7.1 Principal findings and the verdict they support

The evidence answers the central research question affirmatively. The GSADF rejects the random-walk null for 70.3% of the equity panel at the 95% level (91.9% under the weaker at-least-one-episode criterion); the rejection rate rises from 20.0% before the ChatGPT launch to 56.8% after it; and episode timing propagates downstream along the value chain in the sequence documented in Chapter 3. That headline result is not read at face value. As §6.1 and §6.2 establish, the test cannot separate speculative excess from a rational revision of growth expectations; the upstream rejection rate falls from 85.7% to 14.3% and the upstream-downstream ordering reverses under a log-price specification; and the GSADF and SV-ADF agree on only 68% of verdicts. The fourteen tickers confirmed by both methods (§6.4) are accordingly the most defensible bubble candidates the data support.

The verdict is therefore conditional rather than categorical. Explosive price dynamics alone are vulnerable to the fundamental-revision objection; the structural evidence of circular financing alone is vulnerable to the objection that bilateral supply and equity relationships are not inherently pathological. It is their conjunction, set out in §6.3 — explosive dynamics whose timing tracks the structural mechanics, in a sector whose financing is demonstrably self-referential and whose external revenue base is demonstrably insufficient to close the capital-expenditure gap — that supports the following characterization: the AI capital cycle meets the conditions for a bubble in the explosive-root sense of this thesis. That is a statement about the present structure of prices and financing, not a prediction that a collapse is certain or imminent.

7.2 Recommendations and directions for future research

Three recommendations follow for analysts, investors, and supervisors assessing speculative risk in comparable technological transitions. Conduct surveillance at the firm level: all three benchmark indices failed to reject the explosive-root null even though the majority of their AI-exposed constituents did, so an index-only view would have missed the signal entirely (§6.4). Require agreement across methodologically distinct tests before treating a bubble claim as defensible, rather than relying on any single statistic (§6.2, §6.4). And read price-based diagnostics alongside structural and financing evidence: the econometric result is a necessary but not a sufficient condition, and the sufficient condition is supplied by the circular-financing and leverage evidence of Chapter 3 (§6.3).

The principal direction for future research follows from the limitation identified in §6.2: explosive-root tests require a continuous price series, which excludes the privately held foundation-model developers most centrally exposed to the cycle's speculative dynamics. Extending diagnostics to irregularly sampled private valuations — adapting, for example, the cross-sectional and option-implied approaches reviewed in Chapter 2 — is the principal outstanding gap in the bubble-detection literature as it applies to this cycle, and the natural extension of the present work. Refining the volatility-robust procedure further, and integrating the leverage and off-balance-sheet financing channels into a formal joint test of structural and price evidence, are complementary avenues.

The broader contribution of this thesis is to have moved the question of an AI bubble from the register of opinion to that of testable diagnosis. It does not pronounce on whether the AI expansion is desirable, nor on whether the capital being deployed is well spent; it offers a reproducible procedure for asking where, when, and to what degree the prices of an innovation cycle have begun to behave explosively, and a framework for reading that evidence honestly, neither dismissing the genuine technological transformation underway nor mistaking its narrative for its valuation. In a period of genuine uncertainty, that is the more useful contribution: not a verdict to be trusted in authority, but a method that others can apply, contest, and improve.

8. Appendices

8.1 Tables and Figures

8.1.1 Index

Figure 1 - Peak U.S. equity market concentration across speculative episodes. Leading group at its peak, as a share of total U.S. market capitalisation.	8
Figure 2 – Productivity J-curve by Brynjolfsson, Rock, and Syverson	19
Figure 3 - Perez Technology surge cycle (2002). Source : ResearchGate	20
Figure 4 - The AI value chain..	40
Figure 5 - Nvidia Vs AMD, data center segment revenue	44
Figure 6 - Foundry segment – Pure play Market Share	45
Figure 7 - Cloud infrastructure, Market Share	46
Figure 8 - CoreWeave: scale, profitability, customer concentration.....	47
Figure 9 - OpenAI and Anthropic reported ARR trajectory	48
Figure 10 - Microsoft AI business and contracted backlog	50
Figure 11 - Aggregate annual capital expenditure of hyperscalers in Billions of USD dollars	51
Figure 12 - Data center electricity demand and US grid interconnection queue	56
Figure 13 - OpenAI and Anthropic projected financial trajectories.....	59
Figure 14 - Inference cost reduction over time (as \$ per M tokens)	59
Figure 15 - Vendor Financing during the dot-com boom	63
Figure 16 - "A guide to the Circular Deals Underpinning the AI Boom” by Bloombers News	64
Figure 17 - Circular Financing by layer of the AI value-chain, author's synthesis.....	66
Figure 18 - Hyperscaler corporate bond issuance and forecast.....	70
Figure 19 - GSADF results, Index/control	92
Figure 20 - BSADF datestamping - Semiconductor Equipment.....	95
Figure 21 - BSADF datestamping - Chip Design	96
Figure 22 - BSADF datestamping - Memory	97
Figure 23 - BSADF datestamping - Foundry and Advanced Packaging	97
Figure 24 - BSADF datestamping - Storage	98
Figure 25 - BSADF datestamping - DC Networking & Optical	99
Figure 26 - BSADF datestamping - Data-Centre REITs.....	100
Figure 27 - BSADF datestamping - Power and Grid	100
Figure 28 - BSADF datestamping - Hyperscale Cloud.....	101
Figure 29 - BSADF datestamping - Neoclouds	102
Figure 30 - BSADF datestamping - 9 Enterprise AI / Hybrid Cloud.....	102
Figure 31 - BSADF datestamping - Application & AI Platform Software.....	103
Figure 32 - BSADF datestamping - Edge AI & Devices	104
Figure 36 - Post ChatGPT BSADF paths – explosivity propagation evidence.....	105
Figure 37 - Top 8 tickers by SV-ADF explosivity	107
Figure 38 - SV-ADF date stamping - 2 Chip Design.....	108
Figure 39 - SV-ADF datestamping - Foundry and memory	109
Figure 40 - SV-ADF datestamping - 5 Storage.....	110
Figure 41 - SV-ADF datestamping - 6 Data Center Networking & Optical.....	111
Figure 42 - SV-ADF datestamping - 8 Hyperscale Cloud	112
Figure 43 - SV-ADF datestamping - Power & Grid (planned)	113

Figure 44 - Commodity Market evidence for AI-demand propagation	115
Figure 45 - GSADF results, commodity sector, energy	116
Figure 46 - GSADF results, commodities, Industrial metals	117
Figure 47 - GSADF results, commodities, placebo (Agriculture and Livestock).....	118
Figure 48 . GSADF results, commodities, metals and broad commodity indices	119
Figure 49 . SV-ADF results, commodities indices.....	120
Figure 50 - Validation of the developed methods against historical bubble episodes	122
Figure 51 - Validation of the SV-ADF method in python against the historical dot com bubble documented by Sarkar	124
Figure 52 - SV-ADF parameter sensitivity - NASDAQ Composite dot-com episode under two different parameter scenarios	125
Figure 53 - Robustness Check Heatmap	133
Figure 54 - Robustness Scorecard.....	136

8.1.2 Tables

8.1.2.1 AI industry structure

Table 11 - Market structure and summary concentration by layer

<i>Layer</i>	<i>Leading firms</i>	<i>Indicative share / position</i>	<i>Source</i>
<i>EUV lithography</i>	ASML	Effective monopoly (~100% of EUV; sole supplier)	ASML 6-K Q4 2025
<i>Semi equipment (broader)</i>	ASML, AMAT, LAM, TEL, KLA	Top 5 hold the dominant share	10-K filings
<i>AI accelerators</i>	Nvidia, AMD, Intel, custom silicon	Nvidia DC FY26 ~\$195B; AMD DC FY25 \$16.6B; Intel marginal	NVDA, AMD 10-K / 8-K
<i>HBM</i>	SK Hynix, Samsung, Micron	Q2 2025: 62 / 17 / 21 (Counterpoint); Q1 2025: 46-49 / 42-45 / 4-6 (TrendForce)	Counterpoint; TrendForce
<i>Foundry (pure-play)</i>	TSMC, Samsung Foundry, SMIC, UMC, GF	Q3 2025: TSMC 72%; Samsung ~8%; SMIC ~6%; UMC ~5%; GF ~5%	Counterpoint; TrendForce Q4 2025: TSMC 70.4%
<i>Advanced packaging (CoWoS)</i>	TSMC dominant; ASE; Amkor	TSMC dominant; precise share not publicly disclosed	TSMC earnings commentary
<i>Hyperscale cloud</i>	AWS, Azure, GCP, Oracle	Q4 2025: AWS 28% / Azure 21% / GCP 14% / Big Three 63%	Synergy Research Group
<i>Specialised AI cloud</i>	CoreWeave (largest); Nebius; Lambda; Crusoe	CoreWeave FY25 revenue \$5.13B; 67% from single customer	CoreWeave 10-K

<i>Foundation models</i>	OpenAI; Anthropic; Google DeepMind; Meta; Mistral; xAI; DeepSeek	OpenAI ARR ~\$24–25B; Anthropic ARR ~\$47B (May 2026)	The Information; Financial Times
<i>Applications & enterprise</i>	Microsoft, Salesforce, ServiceNow, Adobe, GitHub	Microsoft AI run-rate \$37B (Q3 FY26 disclosure, up 123% YoY)	MSFT 8-K, Apr 2026

Sources: company filings (10-K, 10-Q, 8-K, 6-K, 20-F); Counterpoint Research; TrendForce; Synergy Research Group; The Information and Financial Times for private-company revenue (flagged). Where commercial estimators diverge, both figures are reported.

8.1.2.2 Stationarity Tests (ADF, KPSS, PP)

Table 12 - Full Stationarity Test Results

Sample: 2019-01-02 to 2026-04-30 (end-date for all); start varies for post-2019 listings.

Tests: ADF = Augmented Dickey-Fuller (AIC lag selection, constant); PP = Phillips-Perron (Bartlett-kernel, short bandwidth, $L = \lfloor 4(T/100)^{0.25} \rfloor$); KPSS = Kwiatkowski–Phillips–Schmidt–Shin (level, automatic lags).

Critical values (5%): ADF ≈ -2.863 (sample-adjusted); PP = -2.860 (fixed asymptotic); KPSS = 0.463.

Decision: R = reject H_0 at 5%; FTR = fail to reject. ADF/PP H_0 : unit root ($R \rightarrow I(0)$ signal). KPSS H_0 : level stationarity ($R \rightarrow I(1)$ signal).

p-values: ADF: MacKinnon (1994) asymptotic; PP: not available (fixed asymptotic CVs); KPSS: interpolated by statsmodels (≤ 0.01 when stat exceeds table bounds; ≥ 0.10 when below bounds). **I(d):** majority vote across three tests; * = 2-of-3 majority.

Panel A — AI Hardware Supply Chain (Segments 1–5)

<i>Ticker</i>	<i>Segment</i>	<i>T</i>	<i>ADF</i>	<i>ADF p</i>	<i>ADF</i>	<i>PP</i>	<i>PP</i>	<i>KPSS</i>	<i>KPSS p</i>	<i>KPSS</i>	<i>I(d)</i>
			<i>stat</i>		<i>dec.</i>	<i>stat</i>	<i>dec.</i>	<i>stat</i>		<i>dec.</i>	
<i>ASML</i>	Sem. Equip.	1,842	+0.001	0.959	FTR	−0.290	FTR	4.808	≤ 0.01	R	I(1)
<i>AMAT</i>	Sem. Equip.	1,842	+2.041	>0.999	FTR	+1.081	FTR	4.969	≤ 0.01	R	I(1)
<i>LRCX</i>	Sem. Equip.	1,842	+3.108	>0.999	FTR	+2.406	FTR	4.161	≤ 0.01	R	I(1)
<i>NVDA</i>	Chip Design	1,842	+1.068	0.995	FTR	+0.768	FTR	5.426	≤ 0.01	R	I(1)
<i>AMD</i>	Chip Design	1,842	+0.544	0.986	FTR	+1.034	FTR	4.774	≤ 0.01	R	I(1)
<i>INTC</i>	Chip Design	1,842	−0.409	0.909	FTR	+0.248	FTR	2.918	≤ 0.01	R	I(1)
<i>AVGO</i>	Chip Design	1,842	+1.978	>0.999	FTR	+1.611	FTR	5.127	≤ 0.01	R	I(1)

MRVL	Chip Design	1,842	+0.353	0.980	FTR	+0.271	FTR	4.639	≤0.01	R	I(1)
MU	Memory (HBM)	1,842	+4.357	>0.999	FTR	+4.438	FTR	2.846	≤0.01	R	I(1)
TSM	Foundry	1,842	+2.642	0.999	FTR	+2.288	FTR	4.685	≤0.01	R	I(1)
WDC	Storage (ctrl)	1,842	+6.653	>0.999	FTR	+9.189	FTR	1.994	≤0.01	R	I(1)
STX	Storage (ctrl)	1,842	+7.073	>0.999	FTR	+8.569	FTR	2.988	≤0.01	R	I(1)

Panel B — Data-Centre Infrastructure (Segments 6–9)

Ticker	Segment	T	ADF	ADF p	ADF	PP	PP	KPSS	KPSS p	KPSS	I(d)
			<i>stat</i>		<i>dec.</i>	<i>stat</i>	<i>dec.</i>	<i>stat</i>		<i>dec.</i>	
ANET	DC Network	1,842	+1.126	0.995	FTR	+0.854	FTR	5.636	≤0.01	R	I(1)
CSCO	DC Network	1,842	+1.305	0.997	FTR	+0.728	FTR	4.157	≤0.01	R	I(1)
NOK	DC Network	1,842	+3.539	>0.999	FTR	+3.469	FTR	1.116	≤0.01	R	I(1)
CIEN	DC Network	1,842	+8.298	>0.999	FTR	+9.567	FTR	2.137	≤0.01	R	I(1)
EQIX	DC REIT	1,842	−1.402	0.581	FTR	−1.845	FTR	4.698	≤0.01	R	I(1)
DLR	DC REIT	1,842	−0.747	0.834	FTR	−1.198	FTR	3.765	≤0.01	R	I(1)
MSFT	Hyperscale	1,842	−1.473	0.547	FTR	−1.576	FTR	6.023	≤0.01	R	I(1)
AMZN	Hyperscale	1,842	−0.958	0.768	FTR	−0.966	FTR	3.832	≤0.01	R	I(1)
GOOGL	Hyperscale	1,842	+2.347	0.999	FTR	+2.722	FTR	4.829	≤0.01	R	I(1)
META	Hyperscale	1,842	−0.820	0.813	FTR	−0.929	FTR	4.684	≤0.01	R	I(1)
ORCL	Hyperscale	1,842	−1.432	0.567	FTR	−1.498	FTR	5.182	≤0.01	R	I(1)
IBM	Enterprise AI	1,842	−0.770	0.828	FTR	−0.858	FTR	5.329	≤0.01	R	I(1)

Panel C — AI Software & Edge (Segments 10–11)

Ticker	Segment	T	ADF	ADF p	ADF	PP	PP	KPSS	KPSS p	KPSS	I(d)
			<i>stat</i>		<i>dec.</i>	<i>stat</i>	<i>dec.</i>	<i>stat</i>		<i>dec.</i>	
PLTR	AI Software	1,402	+0.073	0.964	FTR	−0.294	FTR	3.657	≤0.01	R	I(1)
SNOW	AI Software	1,412	−2.190	0.210	FTR	−2.525	FTR	1.966	≤0.01	R	I(1)
DDOG	AI Software	1,662	−2.684	0.077	FTR	−2.770	FTR	2.105	≤0.01	R	I(1)

<i>AI^a</i>	AI Software	1,353	-5.462	<0.001	R	-2.720	FTR	2.200	≤0.01	R	I(1)*
<i>QCOM</i>	Edge AI	1,842	-1.950	0.309	FTR	-2.181	FTR	4.317	≤0.01	R	I(1)
<i>AAPL</i>	Edge AI	1,842	-0.921	0.781	FTR	-0.957	FTR	6.116	≤0.01	R	I(1)
<i>TSLA</i>	Edge AI	1,842	-1.671	0.446	FTR	-1.687	FTR	4.148	≤0.01	R	I(1)

Panel D — Peripheral / Emerging AI Infrastructure

<i>Ticker</i>	<i>Segment</i>	<i>T</i>	<i>ADF</i> <i>stat</i>	<i>ADF p</i>	<i>ADF</i> <i>dec.</i>	<i>PP</i> <i>stat</i>	<i>PP</i> <i>dec.</i>	<i>KPSS</i> <i>stat</i>	<i>KPSS p</i>	<i>KPSS</i> <i>dec.</i>	<i>I(d)</i>
<i>CRWV^b</i>	Neocloud	274	-2.279	0.179	FTR	-2.271	FTR	0.259	≥0.10	FTR	I(1)*
<i>NBIS</i>	Neocloud	382	-0.744	0.835	FTR	-0.709	FTR	2.827	≤0.01	R	I(1)
<i>CEG</i>	Power & Grid	1,074	-1.031	0.742	FTR	-1.200	FTR	5.119	≤0.01	R	I(1)
<i>VST</i>	Power & Grid	1,590	-0.325	0.922	FTR	-0.750	FTR	4.659	≤0.01	R	I(1)
<i>ASX</i>	Adv. Packaging	1,590	+4.191	>0.999	FTR	+5.084	FTR	4.067	≤0.01	R	I(1)

Panel E — Commodity Controls & Benchmarks

<i>Ticker</i>	<i>Segment</i>	<i>T</i>	<i>ADF</i> <i>stat</i>	<i>ADF p</i>	<i>ADF</i> <i>dec.</i>	<i>PP</i> <i>stat</i>	<i>PP</i> <i>dec.</i>	<i>KPSS</i> <i>stat</i>	<i>KPSS p</i>	<i>KPSS</i> <i>dec.</i>	<i>I(d)</i>
<i>USO</i>	Energy	1,842	-0.692	0.849	FTR	-0.014	FTR	0.880	≤0.01	R	I(1)
<i>UNG</i>	Energy	1,842	-1.887	0.338	FTR	-1.745	FTR	3.700	≤0.01	R	I(1)
<i>BNO</i>	Energy	1,842	+0.977	0.994	FTR	+0.847	FTR	4.662	≤0.01	R	I(1)
<i>JJC</i>	Ind. Metals	1,842	-1.805	0.378	FTR	-2.129	FTR	2.847	≤0.01	R	I(1)
<i>PPLT</i>	Ind. Metals	1,842	-0.742	0.835	FTR	-1.137	FTR	2.018	≤0.01	R	I(1)
<i>PALL</i>	Ind. Metals	1,842	-1.457	0.555	FTR	-1.955	FTR	3.333	≤0.01	R	I(1)
<i>GLD</i>	Prec. Metals	1,842	+1.919	0.999	FTR	+1.215	FTR	4.481	≤0.01	R	I(1)
<i>SLV</i>	Prec. Metals	1,842	+0.401	0.982	FTR	-0.665	FTR	3.157	≤0.01	R	I(1)
<i>DBA</i>	Agriculture	1,842	+0.355	0.980	FTR	+0.280	FTR	6.152	≤0.01	R	I(1)
<i>WEAT</i>	Agriculture	1,842	-1.755	0.403	FTR	-1.902	FTR	1.605	≤0.01	R	I(1)
<i>CORN</i>	Agriculture	1,842	-1.363	0.600	FTR	-1.441	FTR	1.930	≤0.01	R	I(1)
<i>SOYB</i>	Agriculture	1,842	-1.246	0.653	FTR	-1.505	FTR	3.457	≤0.01	R	I(1)
<i>COW</i>	Livestock	1,842	-2.772	0.062	FTR	-2.596	FTR	0.687	0.015	R	I(1)

DBC	Commodity idx.	1,842	+0.508	0.985	FTR	+0.313	FTR	5.024	≤ 0.01	R	I(1)
GSG	Commodity idx.	1,842	+0.677	0.989	FTR	+0.589	FTR	4.768	≤ 0.01	R	I(1)
PDBC	Commodity idx.	1,842	+0.384	0.981	FTR	+0.205	FTR	4.952	≤ 0.01	R	I(1)

Panel F — Market Indices (Control)

<i>Ticker</i>	<i>Segment</i>	<i>T</i>	<i>ADF</i>	<i>ADF p</i>	<i>ADF</i>	<i>PP</i>	<i>PP</i>	<i>KPSS</i>	<i>KPSS p</i>	<i>KPSS</i>	<i>I(d)</i>
			<i>stat</i>		<i>dec.</i>	<i>stat</i>	<i>dec.</i>	<i>stat</i>		<i>dec.</i>	
^IXIC	Index	1,842	+0.031	0.961	FTR	−0.142	FTR	5.299	≤ 0.01	R	I(1)
^SOX	Index	1,842	+2.367	>0.999	FTR	+2.215	FTR	5.406	≤ 0.01	R	I(1)
^GSPC	Index	1,842	+0.045	0.962	FTR	−0.155	FTR	5.817	≤ 0.01	R	I(1)

Notes:

^a **AI (C.AI)**: ADF rejects the unit-root null ($\tau = -5.46$, $p < 0.001$), driven by a structural break in the series following the 2022–23 collapse from IPO highs; the PP statistic (-2.72) narrowly fails to reject, and KPSS strongly rejects level stationarity ($\eta = 2.20 \gg \text{cv } 0.463$). Majority vote: 2/3 tests indicate I(1). Treated as I(1) in all subsequent tests.

^b **CRWV**: Listed March 2025; $T = 274$ observations. The KPSS statistic ($\eta = 0.259$) falls below the 5% critical value, formally failing to reject level stationarity — but this reflects low power at $T = 274$, not genuine stationarity. ADF and PP both fail to reject the unit root. Majority vote: 2/3 indicate I(1). Treated as I(1).

COW: KPSS p -value is interpolated exactly at 0.015 (not bounded), indicating rejection is marginal relative to the rest of the sample, though the statistic (0.687) clearly exceeds the 5% cv of 0.463.

Universe totals: $N = 55$ securities — $53 \times \text{I}(1)$ unanimous; $2 \times \text{I}(1)^*$ (majority, AI and CRWV); $0 \times \text{I}(0)$ of any kind.

8.1.2.3 GSADF results

Table 13 - GSADF test results

<i>Segment</i>	<i>Ticker</i>	<i>Company</i>	<i>T</i>	<i>GSADF</i>	<i>CV</i>	<i>p-</i>	<i>Decision</i>	<i>#</i>	<i>Episode periods</i>
				<i>stat</i>	<i>(95%)</i>	<i>value</i>		<i>Ep.</i>	
1. Semiconductor Equipment	ASML	EUV lithography monopoly	1842	1.996	2.390	> 0.10	Fail to reject	1	2021-01 → 2021-11
1. Semiconductor Equipment	AMAT	Applied Materials	1842	2.218	2.390	< 0.10	Fail to reject	2	2020-11 → 2021-06; 2026-01 → 2026-04
1. Semiconductor Equipment	LRCX	Lam Research	1842	4.523	2.390	< 0.01	Reject	3	2019-10 → 2020-02; 2020-12 → 2021-04; 2025-09 → 2026-04

2. Chip Design (AI accelerators)	NVDA	Nvidia	1842	6.340	2.390	< 0.01	Reject	4	2020-07 → 2020-11; 2021-06 → 2022-01; 2023-03 → 2025-02; 2025-06 → 2026-04
2. Chip Design (AI accelerators)	AMD	Advanced Micro Devices	1842	2.902	2.390	< 0.01	Reject	1	2025-07 → 2025-11
2. Chip Design (AI accelerators)	INTC	Intel	1842	3.894	2.390	< 0.01	Reject	0	—
2. Chip Design (AI accelerators)	AVGO	Broadcom	1842	4.002	2.390	< 0.01	Reject	3	2023-05 → 2023-08; 2023-11 → 2025-02; 2025-05 → 2026-04
2. Chip Design (AI accelerators)	MRVL	Marvell	1842	3.603	2.390	< 0.01	Reject	0	—
3. Memory (HBM- relevant)	MU	Micron	1842	7.863	2.390	< 0.01	Reject	3	2020-11 → 2021-03; 2024-03 → 2024-07; 2025-09 → 2026-04
4. Foundry (leading-edge)	TSM	TSMC	1842	2.923	2.390	< 0.01	Reject	5	2019-10 → 2020-01; 2020-07 → 2021-04; 2024-01 → 2026-07; 2024-10 → 2025-01; 2025-09 → 2026-04
5. Storage (non- HBM control)	WDC	Western Digital	1842	8.494	2.390	< 0.01	Reject	2	2024-01 → 2024-07; 2025-07 → 2026-04
5. Storage (non- HBM control)	STX	Seagate	1842	7.944	2.390	< 0.01	Reject	2	2020-12 → 2021-06; 2025-06 → 2026-04
6. DC Networking & Optical	ANET	Arista Networks	1842	3.035	2.390	< 0.01	Reject	2	2023-12 → 2025-02; 2025-08 → 2025-11
6. DC Networking & Optical	CSCO	Cisco	1842	1.373	2.390	> 0.10	Fail to reject	1	2024-10 → 2025-02
6. DC Networking & Optical	NOK	Nokia	1842	5.141	2.390	< 0.01	Reject	0	—
6. DC Networking & Optical	CIEN	Ciena	1842	7.937	2.390	< 0.01	Reject	1	2025-09 → 2026-04
7. DC REITs	EQIX	Equinix	1842	1.545	2.390	> 0.10	Fail to reject	0	—
7. DC REITs	DLR	Digital Realty	1842	1.114	2.390	> 0.10	Fail to reject	0	—
8. Hyperscale Cloud	MSFT	Microsoft	1842	2.568	2.390	< 0.05	Reject	0	—
8. Hyperscale Cloud	AMZN	Amazon	1842	2.318	2.390	< 0.10	Fail to reject	1	2020-06 → 2020-10
8. Hyperscale Cloud	GOOGL	Alphabet	1842	2.746	2.390	< 0.05	Reject	2	2021-04 → 2021-12; 2025-09 → 2026-04
8. Hyperscale Cloud	META	Meta	1842	2.140	2.390	< 0.10	Fail to reject	3	2023-05 → 2023-08; 2024-01 → 2024-07; 2024-09 → 2025-02
8. Hyperscale Cloud	ORCL	Oracle	1842	3.494	2.390	< 0.01	Reject	4	2021-03 → 2021-12; 2023-06 → 2023-09; 2024-09 → 2025-01; 2025-06 → 2025-10
9. Enterprise AI / Hybrid Cloud	IBM	IBM	1842	2.691	2.438	< 0.05	Reject	2	2023-11 → 2024-04; 2024-09 → 2026-02
10. Application & AI Platform Software	PLTR	Palantir	1402	6.718	2.365	< 0.01	Reject	1	2024-09 → 2026-01
10. Application & AI Platform Software	SNOW	Snowflake	1412	0.973	2.359	> 0.10	Fail to reject	0	—
10. Application & AI Platform Software	DDOG	Datadog	1662	1.944	2.407	> 0.10	Fail to reject	1	2020-05 → 2020-08
10. Application & AI Platform Software	AI	C3.ai	1353	4.333	2.364	< 0.01	Reject	0	—
11. Edge AI & Devices	QCOM	Qualcomm	1842	2.140	2.390	< 0.10	Fail to reject	1	2020-11 → 2021-02
11. Edge AI & Devices	AAPL	Apple	1842	2.432	2.390	< 0.05	Reject	1	2019-10 → 2020-02
11. Edge AI & Devices	TSLA	Tesla	1842	7.833	2.390	< 0.01	Reject	2	2019-10 → 2020-03; 2020-06 → 2021-04

8.1.2.4 SV-ADF results

<i>Segment</i>	<i>Ticker</i>	<i>Company</i>	<i>W1 Detected</i>	<i>W1 Episode</i>	<i>W1 Dur. (d)</i>	<i>W1 Collapse</i>	<i>W2 Detected</i>
1. Semiconductor Equipment	ASML	ASML — EUV lithography monopoly	No	—	—	—	No
1. Semiconductor Equipment	AMAT	Applied Materials — wafer fab equipment	No	—	—	—	No
1. Semiconductor Equipment	LRCX	Lam Research — etch/deposition	Yes	2025-11 → 2026-04	157	end_of_sample	No
2. Chip Design	NVDA	Nvidia — GPU/AI accelerators, CUDA	Yes	2024-01 → 2024-07	189	post_bridge	No
2. Chip Design	AMD	Advanced Micro Devices — Instinct/EPYC	No	—	—	—	No
2. Chip Design	INTC	Intel — CPUs, Gaudi accelerators	No	—	—	—	No
2. Chip Design	AVGO	Broadcom — networking + custom AI ASICs	Yes	2024-01 → 2024-04	91	post_bridge	No
2. Chip Design	MRVL	Marvell — custom silicon, AI networking	No	—	—	—	No
3. Memory (HBM)	MU	Micron — DRAM/HBM	Yes	2025-10 → 2026-04	211	end_of_sample	No
4. Foundry	TSM	TSMC — foundry + CoWoS packaging	Yes	2025-09 → 2025-12	90	post_bridge	No
5. Storage (control)	WDC	Western Digital — HDD/NAND (control)	Yes	2025-08 → 2026-04	248	end_of_sample	No
5. Storage (control)	STX	Seagate — HDD (control)	Yes	2025-06 → 2026-04	308	end_of_sample	No
6. DC Networking & Optical	ANET	Arista Networks — DC switching	No	—	—	—	No
6. DC Networking & Optical	CSCO	Cisco — networking incumbent	No	—	—	—	No
6. DC Networking & Optical	NOK	Nokia — telecom networks	No	—	—	—	No
6. DC Networking & Optical	CIEN	Ciena — optical networking	Yes	2025-09 → 2026-04	234	end_of_sample	No
7. DC REITs	EQIX	Equinix — colocation	No	—	—	—	No
7. DC REITs	DLR	Digital Realty — DC REIT	No	—	—	—	No
8. Hyperscale Cloud	MSFT	Microsoft — Azure, OpenAI partner	No	—	—	—	No
8. Hyperscale Cloud	AMZN	Amazon — AWS, Anthropic partner	No	—	—	—	No
8. Hyperscale Cloud	GOOGL	Alphabet — Google Cloud, Gemini	Yes	2025-10 → 2026-02	116	post_bridge	No
8. Hyperscale Cloud	META	Meta — Llama, AI advertising	No	—	—	—	No
8. Hyperscale Cloud	ORCL	Oracle — OCI, AI infrastructure	Yes	2024-09 → 2024-12	90	post_bridge	No
9. Enterprise AI	IBM	IBM — enterprise AI, hybrid cloud	No	—	—	—	No
10. App & AI Platform SW	PLTR	Palantir — enterprise/government AI	Yes	2024-10 → 2025-02	145	post_bridge	No
10. App & AI Platform SW	SNOW	Snowflake — data cloud	No	—	—	—	No
10. App & AI Platform SW	DDOG	Datadog — observability with AI	No	—	—	—	No
10. App & AI Platform SW	AI	C3.ai — enterprise AI platform	No	—	—	—	No
11. Edge AI & Devices	QCOM	Qualcomm — mobile/edge AI silicon	No	—	—	—	No
11. Edge AI & Devices	AAPL	Apple — devices, on-device AI	No	—	—	—	No
11. Edge AI & Devices	TESLA	Tesla — autonomy/robotics (sentinel)	No	—	—	—	No

8.2 Technical Implementation and Repository Structure

8.2.1 Overview

The empirical analysis in this thesis was implemented entirely in Python (version 3.12), using a custom codebase structured as a reproducible research project. All raw data, intermediate computations, and final outputs are separated into distinct directories so that any stage of the pipeline can be re-run independently without affecting other stages. The codebase is designed around three principles: *separation of concerns* (methods, pipelines, data I/O, and visualization live in separate modules), *aggressive caching* (computationally expensive Monte Carlo simulations are written to disk and re-used transparently on subsequent runs, to make subsequent runs more rapid), and *numbered script execution* (the full analysis can be reproduced by running the numbered scripts in sequence, orchestrated end-to-end by `00_run_pipeline.py`). The repository is publicly available on GitHub at: [The-AI-Capital-Cycle-A-Multi-method-Explosive-Root-analysis-on-Equities-across-the-AI-Value-Chain](https://github.com/0x00000000/The-AI-Capital-Cycle-A-Multi-method-Explosive-Root-analysis-on-Equities-across-the-AI-Value-Chain)

8.2.1 Design Choices

- **Caching strategy:** The Monte Carlo simulations needed to calculate the critical values of the GSADF tests are computationally the most expensive component (approximately 5–10 minutes per unique (T, mw) pair at $nrep = 999$, depending on machine capacity). The cache transparently eliminates this cost on repeated runs. Adding a new ticker only triggers a new MC simulation if its sample size T differs from all previously analysed tickers, which in practice is rare once the initial run is complete.
- **Separation of outputs:** The robustness outputs are written to `outputs/robustness/` rather than `outputs/`. This is intentional: it prevents a robustness re-run from overwriting the main thesis tables and figures, and makes it easy to identify on GitHub which files are primary results and which are sensitivity checks.
- **Notebook versus scripts:** The numbered scripts (00 through 12, plus a few lettered and standalone variants) are designed to be run from the command line and handle all computationally intensive work. The three notebooks are reserved for visualisation and interactive exploration: they read from the cached results in `data/results/` and write to `outputs/figures/`. This separation means the notebook can be re-run quickly (in under a minute) without triggering any re-computation, and the scripts can be re-run on a server without requiring a GUI or Jupyter kernel.
- **Determinism and reproducibility:** All random processes in the analysis use `seed = 42` (configurable in `settings.py`). Combined with the Monte Carlo cache, this guarantees that every table and figure in the thesis is fully reproducible from the raw price data.

8.2.2 Environment and Dependencies

The project uses a standard scientific Python stack. All dependencies are pinned in `requirements.txt` at the repository root.

Core libraries

<i>Library</i>	<i>Role</i>
<i>numpy, scipy</i>	Numerical computation, random-walk simulation
<i>pandas</i>	Data manipulation, time-series alignment
<i>statsmodels</i>	ADF, PP, and KPSS test implementations
<i>numba</i>	JIT acceleration of the GSADF inner loop
<i>yfinance</i>	Automated price download from Yahoo Finance
<i>pyarrow</i>	Parquet serialisation of BSADF path arrays
<i>matplotlib</i>	All figures
<i>jupyter, jupyterlab</i>	Notebook environment for visualisation cells

Installation

```
git clone <repository-url>
cd ai-bubble-detection
python -m venv .venv
source .venv/bin/activate # Windows: .venv\Scripts\activate
pip install -r requirements.txt
```

8.2.3 Repository Structure

```
ai-bubble-detection/
├── config/
│   └── settings.py      ← universe, date windows, episode filter params
├── src/
│   ├── methods/        ← one file per test family
│   │   ├── adf_engine.py
│   │   ├── stationarity.py  (ADF, PP, KPSS)
│   │   ├── sadf.py         (PWY 2011)
│   │   ├── gsadf.py        (PSY 2015, Numba-accelerated)
│   │   ├── svadf.py        (Sarkar & Wells 2026, vectorized)
│   │   └── episodes.py
│   ├── monte_carlo.py   ← Monte Carlo + incremental CSV cache
│   ├── io/              ← structured load / save helpers
│   │   ├── data.py
│   │   ├── results.py    (per-method JSON + Parquet)
│   │   └── window.py
│   ├── pipelines/       ← one orchestrator per method
│   │   ├── stationarity.py
│   │   ├── sadf.py
│   │   ├── gsadf.py
│   │   └── svadf.py
│   ├── plotting/        ← reads results from disk only
│   │   ├── diagnostic.py
│   │   ├── comparison.py
│   │   └── tables.py
│   ├── api.py           ← pure functions for app use
└── scripts/             ← numbered = run order; 08b/09b are companion figure scripts for 08/09
```

- | └─ 00_run_pipeline.py ← end-to-end runner (calls 01–05 in order)
- | └─ 01_download_data.py ← fetch prices via yfinance
- | └─ 02_run_montecarlo.py ← pre-compute GSADF/SADF critical values
- | └─ 03_run_analysis.py ← run all or selected methods (--methods flag)
- | └─ 04_run_validation.py ← historical validation (dot-com, GFC)
- | └─ 05_generate_figures.py ← batch-generate all per-ticker + panel PNGs
- | └─ 06_stationarity_table.py ← ADF/PP/KPSS results table (CSV + LaTeX)
- | └─ 07_stationarity_figure.py ← ADF/PP/KPSS heatmap + forest-plot figure
- | └─ 08_robustness.py ← A1–A5/B1–B2 robustness checks + stability heatmap
- | └─ 08b_robustness_figures.py ← thesis-quality versions of 08's figures
- | └─ 09_svadf_validation.py ← SV-ADF standalone historical validation
- | └─ 09b_svadf_validation_comparison.py ← Sarkar's params vs proportional scaling
- | └─ 10_svadf_sector_panels.py
- | └─ 11_volatility_analysis.py
- | └─ 12_gsadf_equity_table.py ← GSADF date-stamping table (CSV + LaTeX)
- | └─ 13_svadf_equity_table.py ← SV-ADF date-stamping table (CSV + LaTeX)
- | └─ 14_commodity_propagation_figure.py ← Figure 6 (H6 commodity evidence)
- └─ notebooks/
 - | └─ 01_visualization.ipynb ← interactive per-ticker viewer (no file saving)
 - | └─ 02_svadf_explorer.ipynb ← interactive single-ticker SV-ADF demo
 - | └─ 03_summary_graphs.ipynb ← publication figures → outputs/figures/thesis/
- └─ app/ ← Hugging Face Space scaffold
 - | └─ README.md
 - | └─ gradio_app.py.example
- └─ data/
 - | └─ prices/prices.csv ← committed (fixed window 2019-2026)
 - | └─ mc_cache/ ← committed (saves hours of Monte Carlo)
 - | └─ results/
 - | └─ stationarity/ ← <ticker>.json
 - | └─ sadf/ ← <ticker>.json
 - | └─ gsadf/ ← <ticker>_summary.json + _paths.parquet
 - | └─ svadf/ ← <ticker>_<start>_<end>_*.json.parquet
- └─ outputs/ ← gitignored; regenerated by scripts
 - | └─ tables/ ← one subfolder per producing method
 - | └─ stationarity/ ← script 06
 - | └─ gsadf/ ← script 12
 - | └─ svadf/ ← script 13
 - | └─ summary/ ← notebook 01 (comparison/episodes/segment tables)
 - | └─ figures/
 - | └─ thesis/ ← publication-ready figures (fig1–fig8 etc.)
 - | └─ gsadf/ ← per-ticker GSADF diagnostic PNGs
 - | └─ sadf/
 - | └─ svadf/
 - | └─ comparison/ ← GSADF-vs-SV-ADF overlays
 - | └─ gsadf_panels/ ← segment-level GSADF grid panels
 - | └─ svadf_sector/
 - | └─ svadf_volatility/
 - | └─ validation_svadf/ ← scripts 09/09b

8.2.4 Configuration: config/settings.py

All user-facing parameters are centralised in a single file. Editing this file is the only change required to alter the scope or parameterisation of the analysis. The most important settings are:

<i>Parameter</i>	<i>Default</i>	<i>Effect</i>
<i>FIXED_START / FIXED_END</i>	2020-01-01 / 2026-05-01	Price download window and main sample
<i>UNIVERSE</i>	37 AI-chain equities	Tickers included in the main panel
<i>COMMODITIES</i>	17 commodity ETFs	Commodity panel for H6
<i>INDICES</i>	^IXIC, ^SOX, ^GSPC	Index benchmarks for H3
<i>GSADF_MC_REPS</i>	999	Monte Carlo replications for GSADF critical values
<i>EPISODE_MIN_DURATION</i>	60 trading days	Minimum episode length (noise filter)
<i>EPISODE_MERGE_GAP</i>	42 trading days	Gap within which two episodes are merged
<i>SVADF_DEFAULT_START/END</i>	2021-11-01 / 2026-05-01	SV-ADF W1 (post-ChatGPT) window
<i>SVADF_PRE_GPT_START/END</i>	2020-01-01 / 2022-10-31	SV-ADF W2 (pre-ChatGPT) window
<i>VALIDATION_CASES</i>	NASDAQ 1990–2003	Historical validation ticker and window

8.2.5 Execution Pipeline

The full analysis is reproduced by running the numbered scripts in the order shown below. Each script is self-contained and can also be run in isolation if only one stage needs to be repeated.

Step 1 — Download prices

```
python scripts/01_download_data.py
```

→ writes data/prices/prices.csv

→ requires internet connection; takes ~2 minutes

Step 2 — Pre-compute Monte Carlo critical values [optional — done automatically]

```
python scripts/02_run_montecarlo.py
```

→ writes data/mc_cache/global_critical_values.csv

and data/mc_cache/bsadf_paths/T{T}_mw{mw}_nrep{n}_seed{s}.csv

→ ~15-30 minutes depending on universe size

→ subsequent runs are instant (cache hits)

Step 3 — Run main analysis

```
python scripts/03_run_analysis.py --methods all
```

→ runs ADF/PP/KPSS, SADF, GSADF, SV-ADF on all tickers

→ writes data/results/{stationarity,sadf,gsadf,svadf}/

→ MC is called internally and cached automatically
→ ~20-40 minutes on first run, much faster on re-runs

Step 4 — Historical validation [optional]

python scripts/04_run_validation.py
→ downloads NASDAQ, BTC, Case-Shiller data
→ runs GSADF on validation cases

Step 5 — Build output tables

python scripts/06_stationarity_table.py # full-sample stationarity table
python scripts/06_stationarity_table.py \ # robustness: custom date window
--start 2022-11-01 --end 2026-05-01
→ writes outputs/tables/stationarity/

Step 6 — Build stationarity figure

python scripts/07_stationarity_figure.py
→ writes outputs/figures/thesis/fig_stationarity.png

Step 7 — Run robustness checks

python scripts/08_robustness.py # all checks (~30-35 minutes)
python scripts/08_robustness.py --skip-mc # skip MC-heavy checks (~10 min)
python scripts/08_robustness.py --checks A2 A3 B2 # fast checks only (~2 min)
→ writes outputs/robustness/tables/ and outputs/robustness/figures/
python scripts/08b_robustness_figures.py # thesis-quality stability heatmap + scorecard
→ writes outputs/robustness/figures/rob_stability_heatmap_v2.png, rob_scorecard.png

Step 8 — Generate figures

Open notebooks/03_summary_graphs.ipynb in JupyterLab
Run all cells in order
→ writes outputs/figures/thesis/fig1 through fig8
(or, non-interactively: python scripts/05_generate_figures.py)

Step 9 — Build date-stamping tables and method-specific figures

python scripts/12_gsadf_equity_table.py # GSADF date-stamping table (CSV + LaTeX)
python scripts/13_svadf_equity_table.py # SV-ADF date-stamping table (CSV + LaTeX)
python scripts/09_svadf_validation.py # SV-ADF historical validation
python scripts/09b_svadf_validation_comparison.py # SV-ADF parameter-config comparison
python scripts/10_svadf_sector_panels.py # SV-ADF segment panels
python scripts/11_volatility_analysis.py # volatility diagnostics
python scripts/14_commodity_propagation_figure.py # Figure 6 (H6 commodity evidence, live-computed)

Alternatively, the entire pipeline can be run with a single command:

8.2.6 Data Layer: data/

data/prices/prices.csv A wide-format CSV with trading dates as rows and ticker symbols as columns, containing daily adjusted close prices for all 55 tickers over the full sample window. The file is generated once by 01_download_data.py via Yahoo Finance and read from cache on all subsequent runs. It is the single source of truth for all price data in the analysis.

data/mc_cache/ The Monte Carlo cache stores pre-computed critical values for the GSADF and BSADF tests. Each Monte Carlo run is uniquely identified by four parameters: sample size T , minimum window mw , number of replications $nrep$, and the random seed. These parameters form the cache key. When a new ticker has a different T from any previously analysed ticker, a new simulation is run and written to disk; on all subsequent calls with the same key, the cached file is read directly. The index file `global_critical_values.csv` records which $(T, mw, nrep, seed)$ combinations have been computed and what GSADF quantiles they produced. Individual BSADF path quantile arrays are stored as `T{T}_mw{mw}_nrep{n}_seed{s}.csv` files, one per unique parameter combination.

data/results/ Per-ticker results from each statistical method, written as JSON summaries and binary Parquet files for array-valued path data.

- `stationarity/<ticker>.json` — ADF, PP, and KPSS test statistics, p-values, critical values, and reject decisions at the 5% level, plus the sample size T and date range.
- `sadf/<ticker>.json` — SADF test statistic and binary signal.
- `gsadf/<ticker>_summary.json` — GSADF statistic, 95% critical value, reject decision, and a list of detected episode dictionaries (start date, end date, duration in days). `<ticker>_paths.parquet` stores the full BSADF statistic path and CV path as a time-indexed DataFrame, used for plotting and robustness re-processing.
- `svadf/<ticker>_<window_id>_summary.json` — SV-ADF result for a specific analysis window, including the episode dictionary if one was detected (origination date, collapse date, collapse type). `<ticker>_<window_id>_paths.parquet` stores the expanding-window coefficient path and the origination and collapse threshold paths.

8.2.7 Outputs: outputs/

outputs/tables/

File	Contents
<code>comparison_table.csv</code>	Master per-ticker table: all test statistics, decisions, and episode counts for all methods
<code>episodes_table.csv</code>	One row per detected episode, with ticker, method, start, end, and duration
<code>segment_summary.csv</code>	Episode rates and reject rates aggregated by value-chain segment
<code>stationarity_table.csv</code>	ADF, PP, KPSS statistics and decisions with implied order of integration (full sample)
<code>stationarity_table_<start>_<end>.csv</code>	Same table computed for a custom date window (generated by <code>--start/--end</code> flags)

8.3 Explosivity tests implemented in Python

8.3.1 Phillips Perron test method

```
def run_pp(series: pd.Series, lags="short") -> dict | None:
    s = np.asarray(series.dropna(), dtype=float)
    n = len(s)
    if n < 30: return None
    L = (int(np.floor(4 * (n / 100) ** 0.25)) if lags == "short" else
         int(np.floor(12 * (n / 100) ** 0.25)) if lags == "long" else int(lags))
    L = max(L, 1)
    y = s[1:]; ylag = s[:-1]
    X = add_constant(ylag)
    ols = OLS(y, X).fit()
    rho, se_rho = ols.params[1], ols.bse[1]
    resid = ols.resid; T = len(resid)
    gamma0 = float(np.dot(resid, resid) / T)
    s2 = gamma0
    for k in range(1, L + 1):
        gk = float(np.dot(resid[k:], resid[:-k]) / T)
        s2 += 2 * (1 - k / (L + 1)) * gk
    s2 = max(s2, 1e-12)
    t_stat = (rho - 1) / se_rho
    Z_t = (math.sqrt(gamma0 / s2) * t_stat
           - 0.5 * (s2 - gamma0) * (T * se_rho
                                   / (math.sqrt(s2) * math.sqrt(np.var(ylag, ddof=0) * T))))
    cv = {"1%": -3.43, "5%": -2.86, "10%": -2.57}
    return {"stat": float(Z_t), "lag": L,
            "cv_1pct": cv["1%"], "cv_5pct": cv["5%"], "cv_10pct": cv["10%"],
            "reject_5pct": bool(Z_t < cv["5%"])}
```

8.3.2 SADF test method

```
def sadf(y, min_window: int | None = None,
         max_lags: int = 0, lag_method: str = "fixed") -> dict:
    y = np.asarray(y, dtype=float)
    T = len(y)
    if min_window is None:
        min_window = minimum_window(T)
    path = np.full(T - min_window + 1, np.nan)
    for i, end in enumerate(range(min_window, T + 1)):
        res = adf_regression(y[:end], max_lags=max_lags, lag_method=lag_method)
        if res is not None:
            path[i] = res["t_stat"]
    return {"statistic": float(np.nanmax(path)),
            "path": path,
            "argmax_index": int(np.nanargmax(path)) + min_window - 1,
            "min_window": min_window}
```

8.3.3 GSADF

```
def run_one(ticker: str, series: pd.Series,
            mc_reps: int | None = None,
            force: bool = False,
            force_mc: bool = False,
            min_duration: int | None = None,
            merge_gap: int | None = None,
            verbose: bool = True) -> dict | None:
```

```
"""
```

Runs GSADF + BSADF episode dating, saves both the summary JSON and a parquet with the full BSADF and CV paths.

If results already exist but the episode-filter settings differ from what was used at last run, only the episode-detection step is re-applied on the stored BSADF/CV paths — the expensive GSADF computation and Monte Carlo are not repeated.

```
"""
```

Resolve filter params first so we can compare against cached values.

```
min_duration = min_duration or settings.EPISODE_MIN_DURATION
```

```
merge_gap = merge_gap or settings.EPISODE_MERGE_GAP
```

```
if not force:
```

```
    existing = load_gsadf(ticker)
```

```
    if existing is not None:
```

```
        stored = existing["summary"].get("episode_filters", {})
```

```
        if (stored.get("min_duration_days") == min_duration and
```

```
            stored.get("merge_gap_days") == merge_gap):
```

```
            return None # fully cached, filters match — skip
```

```
if existing["paths"] is not None:
```

```
    paths = existing["paths"]
```

```
    summary = existing["summary"]
```

```
    episodes = detect_bsadf_episodes(
```

```
        paths["bsadf_stat"].to_numpy(),
```

```
        paths["cv_path"].to_numpy(),
```

```
        paths.index,
```

```
        min_duration=min_duration,
```

```
        merge_gap=merge_gap,
```

```
    )
```

```
    summary["episodes"] = episodes
```

```
    summary["reject"] = bool(summary["statistic"] > summary["cv_value"])
```

```
    summary["episode_filters"] = {
```

```
        "min_duration_days": int(min_duration),
```

```
        "merge_gap_days": int(merge_gap),
```

```
    }
```

```
    save_gsadf(ticker, summary, paths)
```

```

        return summary

mc_reps = mc_reps or settings.GSADF_MC_REPS

s = series.dropna().astype(float)
if len(s) < 100:
    return None
y, T = s.to_numpy(), len(s)
mw = minimum_window(T)

g = gsadf(y, min_window=mw)
cv = get_critical_values(
    T=T, min_window=mw, nrep=mc_reps, seed=settings.SEED,
    quantiles=(settings.GSADF_QUANTILE,),
    force_recompute=force_mc, verbose=verbose,
)
cv_path = cv["bsadf_path"][settings.GSADF_QUANTILE]
cv_val = cv["gsadf"][settings.GSADF_QUANTILE]

bsadf_dates = s.index[mw - 1:]
episodes = detect_bsadf_episodes(
    g["path"], cv_path, bsadf_dates,
    min_duration=min_duration, merge_gap=merge_gap,
)

summary = {
    "ticker": ticker,
    "T": int(T),
    "start_date": s.index[0].isoformat()[:10],
    "end_date": s.index[-1].isoformat()[:10],
    "min_window": int(mw),
    "statistic": float(g["statistic"]),
    "cv_value": float(cv_val),
    "cv_quantile": float(settings.GSADF_QUANTILE),
    "reject": bool(g["statistic"] > cv_val),
    "mc_reps": int(mc_reps),
    "mc_seed": int(settings.SEED),
    "episode_filters": {
        "min_duration_days": int(min_duration),
        "merge_gap_days": int(merge_gap),
    },
    "episodes": episodes,
}

paths_df = pd.DataFrame({
    "bsadf_stat": g["path"],
    "cv_path": cv_path,
}, index=bsadf_dates)
paths_df.index.name = "date"
save_gsadf(ticker, summary, paths_df)

```

8.3.4 Monte Carlo simulation and caching mechanism

```

def simulate_critical_values(T: int, min_window: int, nrep: int, seed: int,
                             quantiles: Iterable[float] = (0.90, 0.95, 0.99),
                             verbose: bool = True) -> dict:
    """Simulate `nrep` random walks and return empirical quantiles."""
    rng = np.random.default_rng(seed)
    npath = T - min_window + 1
    gs_stats = np.empty(nrep)
    bs_paths = np.empty((nrep, npath))
    if verbose:
        print(f" MC: {nrep} reps, T={T}, mw={min_window} ...", end=" ", flush=True)
    t0 = time.time()
    for r in range(nrep):
        rw = np.cumsum(rng.standard_normal(T))
        out = gsadf(rw, min_window=min_window)
        gs_stats[r] = out["statistic"]
        bs_paths[r, :] = out["path"]
    if verbose:
        print(f" {time.time() - t0:.1f}s")
    quantiles = list(quantiles)
    return {
        "gsadf": {q: float(np.quantile(gs_stats, q)) for q in quantiles},
        "bsadf_path": {q: np.quantile(bs_paths, q, axis=0) for q in quantiles},
    }

def _save_cache(T, mw, nrep, seed, cv, quantiles):
    settings.MC_PATHS_DIR.mkdir(parents=True, exist_ok=True)
    df = pd.DataFrame({f"q{int(q*1000):03d}": cv["bsadf_path"][q] for q in quantiles})
    df.index.name = "path_index"
    df.to_csv(_path_file(T, mw, nrep, seed))

    row = {"T": T, "min_window": mw, "nrep": nrep, "seed": seed}
    for q in quantiles:
        row[f"gsadf q{int(q*1000):03d}"] = cv["gsadf"][q]
    idx_path = settings.MC_GLOBAL_INDEX_FILE
    if idx_path.exists():
        existing = pd.read_csv(idx_path)
        mask = ~((existing["T"] == T) & (existing["min_window"] == mw)
                  & (existing["nrep"] == nrep) & (existing["seed"] == seed))
        out = pd.concat([existing.loc[mask], pd.DataFrame([row])], ignore_index=True)
    else:
        out = pd.DataFrame([row])
    out.sort_values(["T", "min_window", "nrep", "seed"]).to_csv(idx_path, index=False)

```

```

def _load_cache(T, mw, nrep, seed, quantiles) -> dict | None:
    p = _path_file(T, mw, nrep, seed)
    idx_path = settings.MC_GLOBAL_INDEX_FILE
    if not p.exists() or not idx_path.exists():
        return None
    idx = pd.read_csv(idx_path)
    row = idx[(idx["T"] == T) & (idx["min_window"] == mw)
              & (idx["nrep"] == nrep) & (idx["seed"] == seed)]
    if row.empty: return None
    row = row.iloc[0]
    df = pd.read_csv(p, index_col="path_index")
    out_gs, out_paths = {}, {}
    for q in quantiles:
        gcol = f"gsadf_q{int(q*1000):03d}"
        pcol = f"q{int(q*1000):03d}"
        if gcol not in row.index or pcol not in df.columns:
            return None
        out_gs[q] = float(row[gcol]); out_paths[q] = df[pcol].to_numpy()
    return {"gsadf": out_gs, "bsadf_path": out_paths}

def get_critical_values(T: int, min_window: int,
                       nrep: int, seed: int = 42,
                       quantiles: Iterable[float] = (0.95,),
                       force_recompute: bool = False,
                       verbose: bool = True) -> dict:
    """
    Cache-or-compute. Cache key: (T, mw, nrep, seed).
    First call simulates and caches; subsequent calls just load.
    """
    quantiles = sorted(set(float(q) for q in quantiles))
    if not force_recompute:
        cached = load_cache(T, min_window, nrep, seed, quantiles)
        if cached is not None:
            if verbose:
                print(f" MC: cache hit T={T}, mw={min_window}, nrep={nrep}")
            return cached
    sim_qs = sorted(set(quantiles) | {0.90, 0.95, 0.99})
    cv = simulate_critical_values(T, min_window, nrep, seed, sim_qs, verbose=verbose)
    _save_cache(T, min_window, nrep, seed, cv, sim_qs)
    return {"gsadf": {q: cv["gsadf"][q] for q in quantiles},
            "bsadf_path": {q: cv["bsadf_path"][q] for q in quantiles}}

def is_cached(T: int, min_window: int, nrep: int, seed: int) -> bool:
    """Check whether MC results exist on disk for this key."""
    return _load_cache(T, min_window, nrep, seed, [0.95]) is not None

```

8.3.4.1 SV-ADF method

```
def svadf(y, min_fraction: float = 0.10) -> dict | None:
    """
    Vectorized SV-ADF coefficient statistic (demeaned OLS, with intercept).

    Parameters
    -----
    y          : 1-D array of prices
    min_fraction : minimum window as fraction of total pairs (default 0.10)

    Returns
    -----
    dict with:
        coef_stat : length-T array (nan before min window), indexed by last obs
        orig_thr  : length-T array of  $\log(\tau)/10$ 
        screen_thr : length-T array of  $\log(\tau)$ 
        coll_thr  : length-T array of  $\log(\tau)/2$ 
        min_window : int (minimum  $\tau$  actually used)
        dates_idx  : 1-D int array of valid indices into the above arrays
    """
    y = np.asarray(y, dtype=np.float64)
    T = len(y)
    pairs = T - 1 # number of ( $X_t, X_{t-1}$ ) pairs

    if pairs < 2:
        return None

    tau_min = max(2, int(min_fraction * pairs))
    if tau_min >= pairs:
        return None

    # Pairs:  $X_t = y[1:]$ ,  $X_{t-1} = y[:-1]$ 
    xt = y[1:] #  $X_t$ 
    x1 = y[:-1] #  $X_{t-1}$ 

    # Cumulative sums (length = pairs)
    cxt = np.cumsum(xt) #  $\sum X_t, t = 1 \dots \tau$ 
    cx1 = np.cumsum(x1) #  $\sum X_{t-1}, t = 1 \dots \tau$ 
    cxt1 = np.cumsum(xt * x1) #  $\sum X_t \cdot X_{t-1}$ 
    cx12 = np.cumsum(x1 ** 2) #  $\sum X_{t-1}^2$ 

    #  $\tau$  grid:  $\tau_{\min} \dots \text{pairs}$  (inclusive)
    tau_arr = np.arange(tau_min, pairs + 1, dtype=np.float64)
    idx = (tau_arr - 1).astype(int) # 0-indexed into cumsum arrays

    # Demeaned OLS numerator / denominator
```

```

# numer =  $\Sigma(X_{t-1} \cdot X_t) - \tau \cdot \text{mean}(X_{t-1}) \cdot \text{mean}(X_t)$ 
# denom =  $\Sigma(X_{t-1}^2) - \tau \cdot \text{mean}(X_{t-1})^2$ 
numer = cxtl[idx] - cxl[idx] * cxt[idx] / tau_arr
denom = cxl2[idx] - cxl[idx] ** 2 / tau_arr

valid = denom > 1e-14
delta_hat = np.where(valid, numer / denom, np.nan)

coef_vals = tau_arr * (delta_hat - 1)

# Date indexing: for window of size  $\tau$ , the last observation is prices[ $\tau$ ],
# matching Sarkar's grid_df$Date <- date_vector[tau + 1] (1-indexed R).
date_idx = tau_arr.astype(int) # ==  $\tau$ 

coef_path = np.full(T, np.nan)
orig_path = np.full(T, np.nan)
screen_path = np.full(T, np.nan)
coll_path = np.full(T, np.nan)

log_tau = np.log(tau_arr)
coef_path[date_idx] = coef_vals
orig_path[date_idx] = log_tau / 10.0
screen_path[date_idx] = log_tau
coll_path[date_idx] = log_tau / 2.0

return {
    "coef_stat": coef_path,
    "orig_thr": orig_path,
    "screen_thr": screen_path,
    "coll_thr": coll_path,
    "min_window": int(tau_min),
    "dates_idx": date_idx,
}

```

8.4 Autoregression

A statistical model is said autoregressive if the current value of a variable is modeled as a function of its own past values.

For a price series, a simple autoregressive model of order 1, written AR(1), says:

$$p_t = \mu + \rho \cdot p_{t-1} + \varepsilon_t$$

Today's price p_t equals a constant μ , plus some fraction ρ of yesterday's price, plus a random shock ε_t . The parameter ρ is called the autoregressive root or persistence parameter, and it controls how strongly today's price depends on yesterday's.

Three cases matter for bubble detection:

- If $\rho < 1$: the series is **stationary**: prices tend to revert toward their long-run mean. Shocks die out over time.
- If $\rho = 1$: the series is a **random walk**: today's price equals yesterday's price plus a random shock, with no tendency to revert. This is the standard null hypothesis for asset prices under efficient markets.
- If $\rho > 1$: the series is **explosive**: each period's price is a magnified version of the previous period's, so shocks accumulate and the series diverges. This is what bubble behavior looks like statistically.

The entire SADF/GSADF testing framework is built around asking: is the estimated ρ in this price series significantly greater than 1 in some subperiod?

8.5 Unit Root testing

Unit root testing

A unit root test is a statistical procedure designed to determine whether a time series contains a unit root, that is, whether its autoregressive root satisfies $\rho = 1$. In practice, the most widely used test is the Augmented Dickey–Fuller (ADF) test, which estimates a regression of the form:

$$\Delta p_t = \alpha + \gamma p_{t-1} + \sum_{i=1}^k \phi_i \Delta p_{t-i} + \varepsilon_t$$

and tests the null hypothesis $\gamma = 0$. This null corresponds to a unit root ($\rho = 1$), meaning the series follows a random walk and is non-stationary.

The alternative hypothesis $\gamma < 0$ corresponds to stationarity ($\rho < 1$), implying mean reversion. Standard unit root tests are therefore left-tailed tests: they are designed to detect whether a series is less persistent than a random walk. In the context of bubble detection, this creates a fundamental limitation. Rational bubble theory predicts explosiveness ($\rho > 1$), not mean reversion. However, conventional unit root tests do not test against this explosive alternative; they test only whether the series is stationary. As a result, failure to reject the unit root null does not imply absence of a bubble, it simply indicates that the data are consistent with a random walk. This limitation motivated the development of right-tailed unit root tests, such as the SADF and GSADF procedures, which explicitly test whether the autoregressive root exceeds unity over some subsample.

8.6 Cointegration

A time series y_t is said to be integrated of order d , written $y_t \sim I(d)$, if it must be differenced d times before becoming stationary. The most common case in financial data is $I(1)$: a series that is non-stationary in levels but stationary in first differences. The canonical example is the random walk:

$$y_t = y_{t-1} + \varepsilon_t,$$

$$\varepsilon_t \sim \text{i.i.d.}(0, \sigma^2)$$

This series has no tendency to revert to a fixed mean, it accumulates shocks indefinitely. Taking the first difference, however, yields $\Delta y_t = \varepsilon_t$, which is $I(0)$. Most financial price series, including stock prices and dividends, exhibit this $I(1)$ behavior.

Two $I(1)$ series p_t and d_t are said to be *cointegrated* if there exists a scalar $\beta \neq 0$ such that the linear combination

$$u_t = p_t - \beta d_t$$

is $I(0)$. Algebraically, this means the two series share a common stochastic trend: each individually contains a random walk component, but that component is proportional across the two series and cancels in the combination u_t , leaving only stationary noise. Economically, cointegration corresponds to a long-run equilibrium relationship — prices and dividends may diverge in the short run, but they cannot drift apart without bound.

The relevance for bubble detection follows directly. In the absence of a bubble, the present-value model implies that stock prices should be cointegrated with dividends with cointegrating vector $\beta = 1/r$, where r is the discount rate. The spread $p_t - \beta d_t$ should be $I(0)$. A rational bubble component B_t , however, satisfies $E_t[B_{t+1}] = (1+r)B_t$: it is explosive ($B_t = (1+r)^t B_0 + \text{noise}$), because the growth rate exceeds unity. Substituting into the price equation:

$$p_t = \beta d_t + B_t$$

the spread $p_t - \beta d_t = B_t$ is no longer $I(0)$, and cointegration between prices and dividends breaks down. Diba and Grossman (1988) formalized this implication: the existence of a rational bubble is a necessary condition for the failure of cointegration between real stock prices and real dividends, so testing for cointegration constitutes an indirect test for the absence of bubbles.

As discussed in §2.2.2, Evans (1991) showed that this inference is not reliable for periodically collapsing bubbles. When a bubble inflates, collapses, and reinflates repeatedly within a long sample, the explosiveness is present only during subperiods. Averaging over the full sample, the price-dividend relationship may appear cointegrated even when genuine bubble episodes occurred within it — the standard cointegration test lacks the power to detect intermittent explosive behavior. This limitation motivates the recursive right-tailed unit root procedures of Phillips, Wu and Yu (2011) and Phillips, Shi and Yu (2015), which test for explosiveness within rolling subsamples rather than across the full estimation window.

9. Bibliography

- Abreu, D., & Brunnermeier, M. K. (2003). Bubbles and crashes. *Econometrica*, 71(1), 173–204. <https://doi.org/10.1111/1468-0262.00393>
- Adrian, T., & Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, 106(7), 1705–1741.
- Aldasoro, I., Doerr, S., & Rees, D. (2026). *Financing the AI boom: from cash flows to debt* (BIS Bulletin No. 120). Bank for International Settlements.
- Anderson, K., & Brooks, C. (2014). Speculative bubbles and the cross-sectional variation in stock returns. *International Review of Financial Analysis*, 35, 20–31.
- Anderson, K., Brooks, C., & Katsaris, A. (2010). Speculative bubbles in the S&P 500: Was the tech bubble confined to the tech sector? *Journal of Empirical Finance*, 17(3), 345–361.
- Basele, R. B., Phillips, P. C. B., & Shi, S. (2025). Explosive dynamics in the Nasdaq and Magnificent Seven. *Journal of Time Series Analysis*.
- Blanchard, O. J., & Watson, M. W. (1982). Bubbles, rational expectations and financial markets. In P. Wachtel (Ed.), *Crises in the economic and financial structure* (pp. 295–316). D.C. Heath and Company.
- Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time series analysis: Forecasting and control* (5th ed.). Wiley.
- Bresnahan, T. F. (2010). General purpose technologies. In B. H. Hall & N. Rosenberg (Eds.), *Handbook of the economics of innovation* (pp. 761–791). Elsevier.
- Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies: Engines of growth? *Journal of Econometrics*, 65(1), 83–108.
- Brockwell, P. J., & Davis, R. A. (2016). *Introduction to time series and forecasting* (3rd ed.). Springer.
- Brunnermeier, M. K., & Nagel, S. (2004). Hedge funds and the technology bubble. *Journal of Finance*, 59(5), 2013–2040. <https://doi.org/10.1111/j.1540-6261.2004.00690.x>
- Brunnermeier, M. K., & Oehmke, M. (2013). Bubbles, financial crises, and systemic risk. In G. M. Constantinides, M. Harris, & R. M. Stulz (Eds.), *Handbook of the economics of finance* (Vol. 2B, pp. 1221–1288). Elsevier.
- Brunnermeier, M. K., Rother, S., & Schnabel, I. (2020). Asset price bubbles and systemic risk. *Review of Financial Studies*, 33(9), 4272–4317.
- Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1), 333–372.
- Cahn, D. (2024, June 20). AI's \$600B question. Sequoia Capital. <https://sequoiacap.com/article/ais-600b-question/>
- Campbell, G. (2012). Myopic rationality in a Mania. *Explorations in Economic History*, 49(1), 75–91.

- Campbell, G., & Turner, J. D. (2012). Dispelling the myth of the naive investor during the British Railway Mania, 1845–46. *Business History Review*, 86(1), 3–41.
- Cortes, G., Lindner, P., Malik, S., & Suntheim, F. (2025). Recycling risk: Synthetic risk transfers (IMF Working Paper No. WP/25/200). International Monetary Fund.
- Courvisanos, J. (2009). The timing of innovation: A stylised Schumpeterian model of the long wave. *Research Policy*, 38(7), 1117–1124.
- De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of Political Economy*, 98(4), 703–738.
- Dell’Acqua, F., McFowland, E., III, Mollick, E., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Krayner, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality (Harvard Business School Working Paper No. 24-013). Harvard Business School.
- Diba, B. T., & Grossman, H. I. (1987). On the inception of rational bubbles. *Quarterly Journal of Economics*, 102(3), 697–700.
- Diba, B. T., & Grossman, H. I. (1988). Explosive rational bubbles in stock prices? *American Economic Review*, 78(3), 520–530.
- Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366), 427–431.
- Evans, G. W. (1991). Pitfalls in testing for explosive bubbles in asset prices. *American Economic Review*, 81(4), 922–930.
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2), 383–417.
- Frehen, R. G. P., Goetzmann, W. N., & Rouwenhorst, K. G. (2013). New evidence on the first financial bubble. *Journal of Financial Economics*, 108(3), 585–607.
<https://doi.org/10.1016/j.jfineco.2012.12.002>
- Gürkaynak, R. S. (2008). Econometric tests of asset price bubbles: Taking stock. *Journal of Economic Surveys*, 22(1), 166–186.
- Gambacorta, L., & Shreeti, V. (2025). The AI supply chain (BIS Papers No. 154). Bank for International Settlements.
- Gârleanu, N., Panageas, S., & Yu, J. (2012). Technological growth and asset pricing. *Journal of Finance*, 67(4), 1265–1292.
- Geraskin, P., & Fantazzini, D. (2013). Everything you always wanted to know about log-periodic power laws for bubble modeling but were afraid to ask. *European Journal of Finance*, 19(5), 366–391.
- Goldfarb, B., & Kirsch, D. A. (2019). Bubbles and crashes: The boom and bust of technological innovation. Stanford University Press.
- González, C. I., & Morar Triandafil, C. (2023). Non-bank financial intermediation and systemic risk (ESRB Occasional Paper No. 23). European Systemic Risk Board.
- Hamilton, J. D. (1994). *Time series analysis*. Princeton University Press.

- Harsha, G. S., & Ismail, M. T. (2019). A classification of stock market prediction methods. *Statistical Journal of the IAOS*, 35(1), 1–11.
- Helpman, E., & Trajtenberg, M. (1998). A general purpose technology at work: The electric dynamo in the United States, 1880–1940. In E. Helpman (Ed.), *General purpose technologies and economic growth* (pp. 55–108). MIT Press.
- Homm, U., & Breitung, J. (2012). Testing for speculative bubbles in stock markets: A comparison of alternative methods. *Journal of Financial Econometrics*, 10(1), 198–231.
- Hoshi, T., & Kashyap, A. K. (2004). Japan's financial crisis and economic stagnation. *Journal of Economic Perspectives*, 18(1), 3–26.
- Jarrow, R. A., & Kwok, S. S. (2021). Inferring financial bubbles from option data. *Journal of Applied Econometrics*, 36(7), 1013–1046.
- Jarrow, R. A., & Lamichhane, S. (2021). Asset price bubbles, market liquidity, and systemic risk. *Mathematics and Financial Economics*, 15(1), 5–40.
- Johansen, A., Ledoit, O., & Sornette, D. (2000). Crashes as critical points. *International Journal of Theoretical and Applied Finance*, 3(2), 219–255.
- Johnson, T. C. (2007). Optimal learning and new technology puzzles. *Journal of Monetary Economics*, 54(8), 2486–2511.
- Kindleberger, C. P., & Aliber, R. Z. (2005). *Manias, panics, and crashes: A history of financial crises* (5th ed.). Wiley.
- Kregel, J. (2009). The global crisis and the implications for developing countries and the BRICs. In W. Drechsler, R. Kattel, & E. S. Reinert (Eds.), *Techno-economic paradigms: Essays in honour of Carlota Perez* (pp. 167–185). Anthem Press.
- Kwiatkowski, D., Phillips, P. C. B., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics*, 54(1–3), 159–178.
- Li, Y., Wang, Z., & Wang, H. (2021). Identifying price bubble periods in the Bitcoin market based on the GSADF model. *Quality & Quantity*, 55(5), 1829–1844.
- MacKinnon, J. G. (1994). Approximate asymptotic distribution functions for unit-root and cointegration tests. *Journal of Business & Economic Statistics*, 12(2), 167–176.
- Maghyereh, A. I., & Awartani, B. (2025). AI bubble dynamics: Liquidity, sentiment, and uncertainty. *Economics and Business Letters*.
- Martin, A., & Ventura, J. (2018). The macroeconomics of rational bubbles: A user's guide. *Annual Review of Economics*, 10(1), 505–539.
- McKinsey & Company. (2025). *The cost of compute: A \$7T race to scale data centers*.
- Minsky, H. P. (1986). *Stabilizing an unstable economy*. Yale University Press.
- Odlyzko, A. (2010). Collective hallucinations and inefficient markets: The British Railway Mania of the 1840s. *B.E. Journal of Economic Analysis & Policy*, 10(1), Article 6. <https://ssrn.com/abstract=1537338>

- Ofek, E., & Richardson, M. (2003). DotCom mania: The rise and fall of internet stock prices. *Journal of Finance*, 58(3), 1113–1137.
- Okina, K., Shirakawa, M., & Shiratsuka, S. (2001). The asset price bubble and monetary policy: Japan's experience in the late 1980s and the lessons. *Monetary and Economic Studies (Special Edition)*, 19(S-1), 395–450.
- Oldani, C., Bruno, G. S. F., & Signorelli, M. (2025). Collapsing bubbles in the prices of cryptocurrencies. *The Journal of Economic Asymmetries*, 31, e00420. <https://doi.org/10.1016/j.jeca.2025.e00420>
- Pástor, L., & Veronesi, P. (2006). Was there a Nasdaq bubble in the late 1990s? *Journal of Financial Economics*, 81(1), 61–100.
- Pástor, L., & Veronesi, P. (2009). Technological revolutions and stock prices. *American Economic Review*, 99(4), 1451–1483.
- Papadovasilaki, D., Guerrero, F., Wuthisatian, R., & Gulati, B. (2022). The roaring twenties and the dot-com bubble: Technological booms and speculative bubbles. *SN Business and Economics*, 2(8).
- Pavlidis, E. G., & Vasilopoulos, K. (2020). Speculative bubbles in segmented markets: Evidence from Chinese cross-listed stocks. *Journal of International Money and Finance*, 109, Article 102222.
- Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. arXiv:2302.06590.
- Perez, C. (2002). Technological revolutions and financial capital: The dynamics of bubbles and golden ages. Edward Elgar.
- Phillips, P. C. B., & Shi, S. (2020). Real time monitoring of asset markets: Bubbles and crises. In H. D. Vinod & C. R. Rao (Eds.), *Handbook of statistics* (Vol. 42, pp. 61–80). Elsevier.
- Phillips, P. C. B., Shi, S., & Yu, J. (2015). Testing for multiple bubbles: Historical episodes of exuberance and collapse in the S&P 500. *Journal of Applied Econometrics*, 30(7), 1065–1109.
- Phillips, P. C. B., Wu, Y., & Yu, J. (2011). Explosive behavior in the 1990s Nasdaq: When did exuberance escalate asset values? *International Economic Review*, 52(1), 201–226.
- Potrykus, M. (2024). Are we in an AI bubble? A dynamic time warping approach. *Finance Research Letters*.
- Rappoport, P., & White, E. N. (1993). Was there a bubble in the 1929 stock market? *Journal of Economic History*, 53(3), 549–574.
- Rappoport, P., & White, E. N. (1994). Was the crash of 1929 expected? *American Economic Review*, 84(1), 271–281.
- Roman, S., & Bertolotti, F. (2025). AI boom-bust dynamics on innovation networks. *Frontiers in Artificial Intelligence*, 8.
- Sarkar, A., & Wells, M. T. (2026). Is there an AI bubble? Robust date-stamping for periods of exuberance. arXiv preprint arXiv:2604.12062. <https://arxiv.org/abs/2604.12062>

- Schot, J., & Kanger, L. (2018). Deep transitions: Emergence, acceleration, stabilization and directionality. *Research Policy*, 47(6), 1045–1059.
- Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71(3), 421–436.
- Shiller, R. J. (2000, 2015). *Irrational exuberance* (3rd ed.). Princeton University Press.
- Shiller, R. J. (2017). Narrative economics. *American Economic Review*, 107(4), 967–1004. <https://doi.org/10.1257/aer.107.4.967>
- Sidak, J. G. (2003). The failure of good intentions: The WorldCom fraud and the collapse of American telecommunications after deregulation. *Yale Journal on Regulation*, 20(2), 207–267.
- Sornette, D. (2003). *Why stock markets crash: Critical events in complex financial systems*. Princeton University Press.
- Stiglitz, J. E. (1990). Symposium on bubbles. *Journal of Economic Perspectives*, 4(2), 13–18.
- Temin, P., & Voth, H.-J. (2004). Riding the South Sea bubble. *American Economic Review*, 94(5), 1654–1668. <https://doi.org/10.1257/0002828043052268>
- Teti, E., & Maroni, D. (2021). Is there a new tech bubble? Comparing the valuations of the dot-com era and today. *Technology Analysis & Strategic Management*.
- Tirole, J. (1985). Asset bubbles and overlapping generations. *Econometrica*, 53(6), 1499–1528.
- Wei, Y., Li, Z., & Wang, Y. (2022). Speculation or fundamentals? Regime-dependent drivers in the crude oil market. *Energy Economics*, 112, Article 106232.
- White, E. N. (1990). The stock market boom and crash of 1929 revisited. *Journal of Economic Perspectives*, 4(2), 67–83.
- Woodman, C. (2021, March). Defining, detecting and measuring asset price bubbles. Cambridge Judge Business School. <https://www.jbs.cam.ac.uk/2021/defining-detecting-measuring-asset-price-bubbles/>