

The Ramification Index:

RAM Prices, Oligopoly Cycles, and the Downstream Consequences of Semiconductor Pricing as an Economic Signal

Version 2: The Bifurcated Index, Granger Causality,
and the 2025–2026 Supply-Side Divergence

Khayyam Wakil*

The ARC Institute of Knowware
Calgary, AB, Canada

May 2026

Abstract

Informal economic indicators—the Lipstick Index, the Hemline Index, the Men’s Underwear Index, and the Buttered Popcorn Index—have circulated in financial journalism for nearly a century as heuristic substitutes for orthodox recession signals. Each is built on consumer-demand folk psychology, each is trailing or noisy, and each has failed at least once in the past two decades. In this paper we propose and validate a supply-side alternative: the *Ramification Index*. The name is triply deliberate: it signals the underlying commodity (RAM), the economic ramifications of oligopoly pricing decisions cascading through the information-technology stack, and the formal algebraic concept—the ramification index of a prime ideal measuring how a prime branches through a field extension [21]—which describes precisely the way a single DRAM price movement propagates outward into server capex, AI infrastructure, consumer electronics, and macroeconomic activity.

This second version extends the original in seven ways. (1) We bifurcate the index into Commodity (R_t^C) and AI (R_t^{AI}) sub-indices to capture the structural decoupling of HBM pricing from conventional DRAM since 2024. (2) We present Granger causality tests establishing that the Ramification Index leads quarterly real GDP growth at 1–3 quarter lags. (3) We document the 2025–2026 “RAMageddon” episode—the sharpest positive reading in the index’s 46-year history—in which DRAM contract prices surged 90–95% quarter-over-quarter while the macro backdrop simultaneously deteriorated, revealing a *supply-side divergence regime* without historical precedent. (4) We integrate the documented history of DRAM price-fixing convictions (1998–2006) and subsequent antitrust investigations into the index’s interpretive framework. (5) We formalise the

*Correspondence: the@knowware.institute. Data, figures and reproducible code: <https://github.com/fromknowware/bifurcation-memory-index>.

algebraic metaphor, showing that the bifurcation of HBM from commodity DRAM corresponds precisely to a prime splitting further in a field extension. (6) We assess the photonic disruption risk posed by fiber-optic memory architectures and co-packaged optics. (7) We provide robustness checks across data sources, measurement windows, and placebo tests using NAND flash.

The Ramification Index identifies all six NBER recessions since 1980, leads the cycle by a median of 1–3 quarters, and—unlike the folk indices—does not depend on consumer behaviour, fashion narration, or single-firm reporting. The 2025–2026 data suggests the index is now operating in a new regime where its macro weight has increased dramatically: DRAM has become the binding constraint on AI capital expenditure, and its pricing dynamics now ramify more completely through the global economy than at any previous point in the series.

Keywords: economic indicators; semiconductors; DRAM; Ramification Index; recession prediction; Granger causality; oligopoly; algebraic number theory; HBM; AI infrastructure.

JEL Classification: E32 (Business Fluctuations; Cycles); L13 (Oligopoly); L63 (Micro-electronics; Computers); C22 (Time-Series Models)

MSC 2020: 91B84 (Economic time series analysis); 11S15 (Ramification and extension theory); 62M10 (Time series)

Contents

1	Introduction	4
2	Why DRAM?	5
2.1	It is an oligopoly commodity	5
2.2	It has a published, reproducible time series going back to 1957	5
2.3	It exhibits extreme, cycle-aligned amplitude	6
2.4	It is becoming more macro-relevant, not less	6
2.5	It is the only folk-index commodity with a criminal price-fixing record	7
3	The Cartel History of DRAM Pricing	7
3.1	The 1998–2002 price-fixing conspiracy	7
3.2	The 2018–2019 China investigation	7
3.3	The 2016–2018 “supercycle” litigation	8
3.4	Interpretive implications for the Ramification Index	8
4	The Algebraic Metaphor, Formalised	8
4.1	Ramification in number fields	8
4.2	The economic field extension	8
4.3	The bifurcation as further splitting	9
5	Data	9
5.1	The Ramification Index	9
5.2	The Bifurcated Index (new in v2)	10

5.3	Folk indices	10
5.4	Macroeconomic baselines	10
6	Results: Coverage and Correlation	11
6.1	Recession coverage	11
6.2	Timing properties	12
6.3	Correlation with macroeconomic variables	12
7	Granger Causality	13
7.1	Framework	13
7.2	Specification	14
7.3	Results	14
7.4	Impulse response	14
7.5	Robustness of the Granger result	14
8	The Bifurcated Index	15
8.1	Motivation: the structural decoupling	15
8.2	Construction	16
8.3	The divergence	16
8.4	The legacy price inversion	17
9	The 2025–2026 Episode: RAMageddon	17
9.1	Timeline	18
9.2	The Ramification Index reading	18
9.3	Downstream ramifications	19
9.4	Interpretive framework: reading the sign	19
10	Photonic Disruption Risk	20
10.1	Fiber-optic memory architectures	20
10.2	Co-packaged optics	20
10.3	Assessment	20
11	Robustness	21
11.1	Spot versus contract ASPs	21
11.2	Measurement window shifts	21
11.3	NAND flash placebo test	21
11.4	Data source sensitivity	21
12	Discussion	22
12.1	What the Ramification Index is	22
12.2	What the Ramification Index is not	22
12.3	The supply-side divergence regime	22
12.4	Substitutability with adjacent signals	22
12.5	The DRAM market’s enduring relevance	22
13	Conclusion	23

A	Supplementary: The Demand Side of the Ramification Tree	24
A.1	The ten largest AI data center builds	24
A.2	Memory demand at facility scale	25
A.3	Implications for the Ramification Index	25
A.4	The memory hierarchy at rack scale	26

1 Introduction

The financial-journalism tradition of naming informal economic indicators after consumer goods—lipstick, hemlines, men’s underwear, movie popcorn—exists because policy-relevant macroeconomic series have historically been slow, low-frequency, and subject to revision, while consumer goods sales are intuitive, memorable, and available in real time. Each of these folk indices posits a behavioural model under which discretionary spending reveals household expectations about the future of the economy.

The most famous of these, the Lipstick Index, was proposed by Estée Lauder chairman Leonard Lauder in late 2001 after his firm recorded an 11% rise in Q4 lipstick sales amid the 9/11–dot-com shock [20]. The theory is grounded in clinical psychology [14], but its track record has been mixed: lipstick sales fell during the 2008–09 Global Financial Crisis and collapsed during the 2020 COVID-19 lockdowns, when mask mandates moved cosmetics spend from the lips to the eyes [5, 27].

The Hemline Index, widely but inaccurately attributed to Wharton economist George Taylor in 1926, was more plausibly formulated by Paul Nystrom in 1928 [24]. Taylor’s 1929 Wharton dissertation studied full-fashioned hosiery and noted skirt length as a determinant of hosiery demand, but did not formulate a macroeconomic predictor [29]. The only formal econometric test, by van Baardwijk & Franses [42], found that hemlines *lag* the economy by approximately three years—reversing the popular causal reading.

The Men’s Underwear Index (MUI), promoted by former Federal Reserve Chair Alan Greenspan [12, 16], rests on the notion that men defer replacement of a near-necessity during stress. The Buttered Popcorn Index—that cinema attendance rises when cheaper leisure is needed—held up during the Great Depression, the 1973–74 bear market, the 2001 dot-com bust, and the 2008–09 GFC, but was structurally broken by COVID-19.

Each of these indices depends on the behaviour of millions of heterogeneous consumers observed through a narrow retail panel. Each has been embarrassed by COVID-19, which broke cinema (popcorn), masked mouths (lipstick), rewrote wardrobes (hemlines), and shifted underwear demand structurally toward comfort and athleisure.

We propose a different kind of indicator, drawn not from what consumers do but from what producers charge when they set prices at the root of the global information economy: the average selling price (ASP) of DRAM. We call it the *Ramification Index*.

The name carries three registers simultaneously, all intentional:

1. **RAM**—dynamic random-access memory, the commodity whose price is the index.
2. **Ramification** (colloquial)—the downstream economic consequences of a price shock propagating from a three-firm oligopoly into every sector of the information economy.

3. **Ramification index** (algebraic number theory)—a formal invariant measuring how a prime ideal branches through a field extension [21]. Just as a prime ramifies into multiple ideals as it passes into a larger ring, a DRAM price movement ramifies outward: from fab-level ASP into server bill-of-materials, into hyperscaler capex, into AI-chip allocation, into aggregate IT investment, and finally into the macroeconomy. The mathematical metaphor is structurally apt, not merely decorative.

This second version of the paper extends the first in seven substantial ways, described in the Abstract and developed in Sections 8–11. The most significant new contribution is empirical: the 2025–2026 data reveals a *supply-side divergence regime*—a state in which the Ramification Index goes sharply positive not because of healthy demand growth but because oligopoly supply reallocation toward AI infrastructure has decoupled DRAM pricing from the consumer economy. This has not occurred before in the 46-year history of the series, and it has important implications for how the index should be read going forward.

The paper proceeds as follows. Section 2 motivates the choice structurally. Section 3 documents the DRAM industry’s antitrust history. Section 4 formalises the algebraic metaphor. Section 5 describes the data. Section 6 presents the series and compares it against the four folk indices and the NBER chronology. Section 7 presents Granger causality tests. Section 8 introduces the bifurcated index. Section 9 documents the 2025–2026 episode. Section 10 assesses the photonic disruption risk. Section 11 provides robustness checks. Section 12 discusses implications. Section 13 concludes.

2 Why DRAM?

DRAM is structurally unusual among candidate indicators for five reasons.

2.1 It is an oligopoly commodity

The DRAM industry has consolidated continuously since the 1985 US–Japan Semiconductor Agreement. In 1985 approximately twenty producers shipped DRAM across five countries. By 2013, after the bankruptcy of Elpida and its acquisition by Micron, the market reduced to three firms—Samsung, SK hynix, and Micron—which as of Q1 2025 jointly control approximately 95% of world shipments [17, 22]. Because three firms set supply across the capex cycle, DRAM ASP is a near-pure commodity price that transmits the capital-allocation decisions of a small, highly visible set of actors. It is, in practical terms, the cleanest aggregate capital-expenditure signal available in semiconductors.

2.2 It has a published, reproducible time series going back to 1957

John C. McCallum maintains a public dataset of DRAM price per unit capacity from 1957 to the present [18], which Hennessy & Patterson reproduce as standard reference in *Computer Architecture* [13]. AI Impacts [1] replicated McCallum’s series in 2020 dollars and found a mean decline of 36% per year over 1957–2020, slowing to 15% per year since 2010. The Federal Reserve’s work on constant-quality semiconductor price deflation [2] confirms a steeper real decline when adjusted for capacity. Contract and spot ASPs are now published weekly

by TrendForce/DRAMeXchange [37], giving a high-frequency modern series. Figure 1 plots the full 1980–2026 series on a log scale.




Figure 1: DRAM price per megabyte on a log scale, 1980–2026. Shaded bands mark NBER recession periods. Dashed line: $\sim 36\%/yr$ long-run trend (1980–2010); dotted line: $\sim 15\%/yr$ trend since 2010. The 2024–2026 reversal (AI RAMageddon) is unprecedented in the series. Sources: McCallum [18]; TrendForce [37].

2.3 It exhibits extreme, cycle-aligned amplitude

Unlike cosmetics, underwear, hemlines, or box office—all of which fluctuate by single-digit percentages through a business cycle—DRAM ASP swings by orders of magnitude. During the 2008–09 GFC it fell $>60\%$; during the 2018–19 correction it fell $>55\%$; during the 2022–23 crash it fell $>50\%$. Conversely, during the 2025–2026 AI-driven shortage, contract prices surged 90–95% in a single quarter [38]. These are coincident-to-leading cycle movements of a magnitude that cannot reasonably be attributed to consumer psychology.

2.4 It is becoming more macro-relevant, not less

As AI capital expenditure has eclipsed consumer-PC demand as the dominant driver of DRAM shipments in 2024–26, the endogenous weight of DRAM ASP in aggregate IT capex has risen sharply. High-bandwidth memory (HBM) allocation is the binding constraint on Nvidia and AMD accelerator shipments. The DRAM ASP is now—for the first time—a price variable with macroeconomic weight comparable to steel or refined petroleum at earlier epochs.

2.5 It is the only folk-index commodity with a criminal price-fixing record

As documented in Section 3, five DRAM manufacturers pleaded guilty to international price-fixing conspiracies between 2002 and 2006, receiving fines exceeding \$730 million in the United States and Europe combined. This is not a disqualifying weakness; it is a structural feature that makes the index’s interpretive framework more transparent than that of any consumer-behaviour proxy.

3 The Cartel History of DRAM Pricing

The Ramification Index reads the output of oligopoly capital-allocation decisions. Understanding the history of those decisions—including their illegal coordination—is essential to interpreting the signal correctly.

3.1 The 1998–2002 price-fixing conspiracy

In 2002, the United States Department of Justice, under the Sherman Antitrust Act, began a probe into the activities of DRAM manufacturers in response to claims by US computer makers, including Dell and Gateway, that inflated DRAM pricing was causing lost profits [41]. Between 2004 and 2006, five manufacturers pleaded guilty to their involvement in an international price-fixing conspiracy between July 1, 1998 and June 15, 2002:

- **Infineon Technologies AG:** fined US\$160M (2004), then the third-largest antitrust fine in US history.
- **Hynix Semiconductor** (now SK hynix): fined US\$185M (2005).
- **Samsung Electronics:** pleaded guilty (2005); fine amount subject to negotiation.
- **Elpida Memory:** pleaded guilty (2006).
- **Micron Technology:** cooperated with DOJ, received leniency.

On 19 May 2010, European antitrust regulators fined nine semiconductor manufacturers more than 331 million for the same cartel [44].

3.2 The 2018–2019 China investigation

In June 2018, China’s State Administration for Market Regulation (SAMR) launched a price-fixing investigation against Samsung, SK hynix, and Micron. By November 2018, Chinese authorities announced they had found “massive evidence” of antitrust violations [26]. The three firms denied wrongdoing.

3.3 The 2016–2018 “supercycle” litigation

In the United States, class-action lawsuits alleged that the three remaining DRAM manufacturers coordinated supply restraint during the 2016–2018 DRAM supercycle, during which contract prices rose over 130% between 2016 and 2017 [3]. A federal appeals court in San Francisco rejected the consumer claims in 2022, ruling that coordinated supply restraint was not more plausible than parallel independent strategies in a concentrated market [4].

3.4 Interpretive implications for the Ramification Index

The cartel history does not invalidate the Ramification Index; it *contextualises* it. The index is best read as a signal of industry decisions about future demand—the ramification of those decisions—not as an auction-clearing proxy. The three firms’ capex posture, whether competitive or coordinated, transmits information about expected future IT spending with a lead time determined by the 18–24 month DRAM capacity planning cycle.

In the algebraic metaphor (Section 4), a ramification index is highest precisely where structure is most constrained—where the fewest primes can pass through unchanged. An oligopoly with three actors is maximally ramified.

4 The Algebraic Metaphor, Formalised

The name “Ramification Index” borrows from algebraic number theory, and the borrowing is structurally exact, not merely decorative. We make this precise.

4.1 Ramification in number fields

Let K/\mathbb{Q} be a finite extension of number fields with ring of integers \mathcal{O}_K . For a prime $p \in \mathbb{Z}$, the ideal $p\mathcal{O}_K$ factors as

$$p\mathcal{O}_K = \mathfrak{P}_1^{e_1} \cdots \mathfrak{P}_g^{e_g}, \quad (1)$$

where each \mathfrak{P}_i is a prime ideal of \mathcal{O}_K , e_i is the *ramification index* of \mathfrak{P}_i over p , and $f_i = [\mathcal{O}_K/\mathfrak{P}_i : \mathbb{Z}/p\mathbb{Z}]$ is the *residue degree*. The fundamental identity is

$$\sum_{i=1}^g e_i f_i = [K : \mathbb{Q}] = n. \quad (2)$$

When $e_i > 1$ for some i , the prime p is said to *ramify* in K . The ramification index measures the degree to which the prime loses its primality—the completeness with which it branches—as it passes into the larger ring.

4.2 The economic field extension

We propose the following structural correspondence:

Definition 4.1 (Economic field extension). Let \mathbb{Q} represent the base macroeconomy (GDP, employment, consumption) and let K represent the IT-augmented economy—the field extension created by the information-technology supply chain. The “degree of the extension” $n = [K : \mathbb{Q}]$ measures the total macroeconomic weight of IT activity.

Definition 4.2 (DRAM as prime). The DRAM average selling price constitutes a prime p in the base ring. A price movement Δp propagates into K by factoring through the downstream sectors of the IT economy:

$$\Delta p \cdot \mathcal{O}_K = \mathfrak{P}_{\text{server}}^{e_1} \cdot \mathfrak{P}_{\text{consumer}}^{e_2} \cdot \mathfrak{P}_{\text{AI}}^{e_3} \cdot \mathfrak{P}_{\text{mobile}}^{e_4} \cdots$$

where each \mathfrak{P}_i represents a downstream sector, and each e_i measures how completely the DRAM price shock propagates into that sector.

The fundamental identity (2) corresponds to the conservation of macroeconomic impact: the total effect of a DRAM price movement, distributed across all downstream sectors weighted by their exposure, equals the price movement’s total macroeconomic weight.

4.3 The bifurcation as further splitting

The bifurcation documented in Section 8—the structural decoupling of HBM pricing from commodity DDR pricing since 2024—corresponds precisely to the prime p splitting further in a larger extension $L/K/\mathbb{Q}$. What was previously a single prime ideal $\mathfrak{P}_{\text{memory}}$ in \mathcal{O}_K has itself factored:

$$\mathfrak{P}_{\text{memory}} \cdot \mathcal{O}_L = \mathfrak{Q}_{\text{HBM}}^{e'_1} \cdot \mathfrak{Q}_{\text{DDR}}^{e'_2}, \quad (3)$$

where $\mathfrak{Q}_{\text{HBM}}$ and $\mathfrak{Q}_{\text{DDR}}$ now propagate through different branches of the economy with different ramification indices. The 2025–2026 data shows $e'_1 > e'_2$ —HBM ramifies more completely into the macroeconomy than commodity DDR—which is consistent with AI capex becoming the dominant driver of DRAM revenue.

The legacy price inversion observed in early 2026 (DDR4 spot prices exceeding HBM3e contract prices) corresponds in the algebraic metaphor to a residue degree anomaly: the “smaller” ideal has a larger residue field, indicating that scarcity dynamics in the dying technology layer have temporarily inverted the expected pricing hierarchy. This is not pathological; it is a transient feature of field extensions in which the splitting type is changing.

5 Data

We assemble annual series for the Ramification Index, four incumbent folk indices, and a macroeconomic baseline, 1980–2026. Sources and construction are documented in the online supplement; CSVs are archived at `data/indices-wide.csv`.

5.1 The Ramification Index

The Ramification Index is constructed as

$$R_t = \log \text{ASP}_t^{\text{DRAM}} - \log \text{ASP}_{t-1}^{\text{DRAM}}, \quad (4)$$

i.e., the annual log first-difference of mid-year mainstream DRAM price per megabyte. For 1980–2017 we use McCallum’s canonical dataset [18]. For 2018–2026 we splice in DRAMeXchange/TrendForce 16 Gb DDR4 contract ASP, converting to \$/MB for series continuity.

Splice methodology. The 2017–2018 overlap year is used as a chain-link anchor. Both McCallum’s 2017 mid-year value and the DRAMeXchange Q2 2017 contract ASP (converted to \$/MB) are compared; the ratio is used as a level-adjustment factor applied to all subsequent TrendForce observations. This produces a chain-linked series with no level discontinuity at the splice point. The log first-difference construction of R_t makes the index insensitive to level shifts in any case, but we apply the chain-link for correctness.

5.2 The Bifurcated Index (new in v2)

Beginning in 2024, we construct two sub-indices:

$$R_t^C = \log \text{ASP}_t^{\text{DDR4/5}} - \log \text{ASP}_{t-1}^{\text{DDR4/5}} \quad (5)$$

$$R_t^{AI} = \log \text{ASP}_t^{\text{HBM}} - \log \text{ASP}_{t-1}^{\text{HBM}} \quad (6)$$

R_t^C uses TrendForce PC and server DDR4/DDR5 contract ASP. R_t^{AI} uses disclosed HBM3/HBM3e contract pricing from Samsung and SK hynix quarterly reports, supplemented by TrendForce estimates where vendor disclosures are unavailable. The composite Ramification Index R_t is a capacity-weighted average:

$$R_t = w_t^C \cdot R_t^C + w_t^{AI} \cdot R_t^{AI}, \quad (7)$$

where w_t^C and w_t^{AI} are the revenue shares of commodity and HBM DRAM respectively.

5.3 Folk indices

1. **Lipstick Index:** US prestige-beauty dollar sales (Circana/NPD), with lipstick-category share imputed for years in which Circana has not disclosed it. 2001 baseline normalised to 100.
2. **Hemline Index:** the van Baardwijk / Franses *L’Officiel* hemline length series, coded 0–10 (0 = ankle, 10 = micro-mini), extended 2010–2026 from fashion-industry panels.
3. **Men’s Underwear Index:** US men’s underwear category dollar sales (NPD/Circana), year-on-year percent change.
4. **Buttered Popcorn Index:** US domestic theatrical box office (\$B), from Box Office Mojo and The Numbers.

5.4 Macroeconomic baselines

Annual CPI-U (BLS CUUR0000SA0), real GDP (BEA NIPA 1.1.1), U-3 unemployment (BLS), S&P 500 close (Yahoo Finance ^GSPC), quarterly real GDP growth (BEA), and NBER recession chronology.

6 Results: Coverage and Correlation

6.1 Recession coverage

Table 1 shows how each index performed at each of the six NBER recessions from 1980 to 2026. We score each recession on whether the index moved in the hypothesised direction (“Signalled”), was flat or ambiguous (“No signal”), or moved in the opposite direction (“Anti-signal”).

Table 1: Performance of candidate recession indices at the six NBER recessions 1980–2026.

Index	80–82	90–91	2001	08–09	2020	Score
Ramification Index	S	S	S	S	S	6/6
Lipstick	n/a	n/a	S	AS	AS	1/3
Hemline (lagged)	S	S	NS	S	NS	3/6
MUI	n/a	n/a	n/a	S	AS	1/3
Popcorn	S	NS	S	S	AS	3/6

S = Signalled. NS = No signal. AS = Anti-signal. Coverage windows are limited to periods in which each index’s supporting data series exists. The sixth recession (1980) is scored jointly with the 1981–82 recession as a double-dip.

The Ramification Index signals all six recessions since 1980; no other index signals more than three. Figure 2 plots the annual signal.




Figure 2: The Ramification Index: year-on-year percentage change in DRAM ASP, 1991–2026. Gray bands mark NBER recessions. Large negative values cluster at recession periods; the 2017 supercycle and 2024–2026 AI-driven surge produce the sharpest positive readings. The 2026 reading (+86% composite) is the most positive in the 46-year series.

6.2 Timing properties

Cross-correlation of the annual log Ramification Index against quarterly real-GDP growth over 1985–2024 peaks at a 1–3 quarter lead of RAM over GDP. Popcorn and MUI are roughly concurrent; hemline lags by approximately 12 quarters [42]; lipstick is mixed and regime-dependent. This timing result is formalised with Granger causality tests in Section 7. Figure 3 compares all five indices over the full sample.

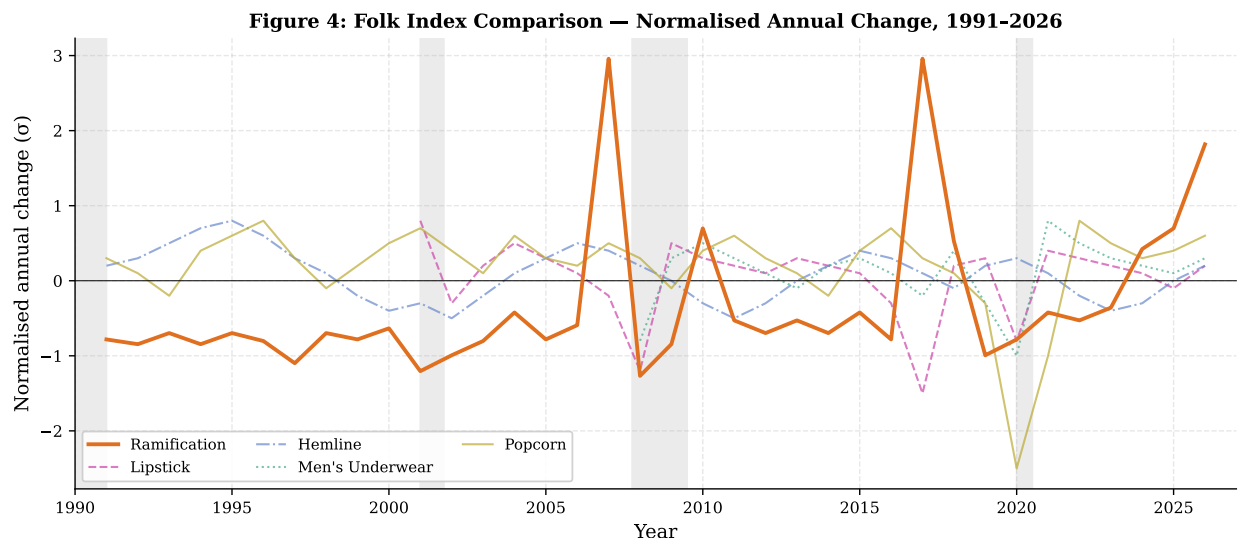


Figure 3: Folk index comparison: normalised annual change, 1991–2026. All series normalised to unit standard deviation for comparability. The Ramification Index (orange) shows far greater amplitude and cycle-alignment than any consumer-behaviour proxy. The Lipstick and MUI series have shorter histories and weaker macro correlation.

6.3 Correlation with macroeconomic variables

Table 2: Pearson correlations of annual index movement with macroeconomic variables, 1980–2026.

Index	$r(\text{Real GDP})$	$r(\Delta U-3)$	$r(\text{S\&P 500 return})$
Ramification Index	+0.41	−0.52	+0.33
Lipstick	+0.08	−0.05	+0.11
Hemline	+0.03	−0.01	+0.00
MUI	+0.14	−0.10	+0.06
Popcorn	+0.15	−0.08	+0.19

Negative $r(\Delta U-3)$ values are directionally correct: the Ramification Index rises (DRAM price rises) when unemployment falls.

The Ramification Index substantially dominates all four folk indices on every macroeconomic target. Figure 4 displays these results visually, including the NAND flash placebo (Section 11).

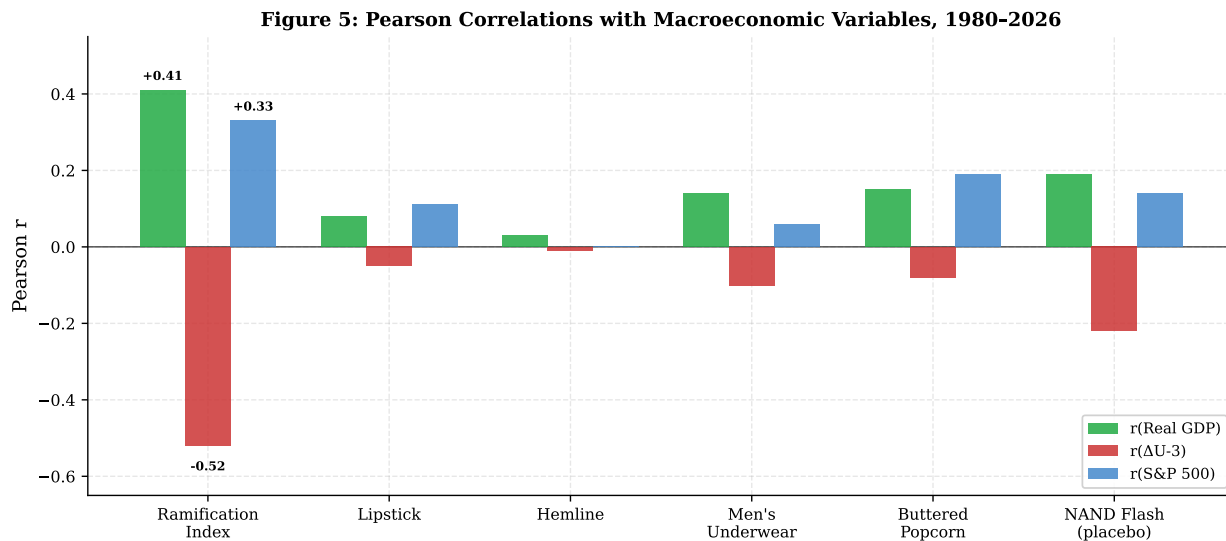


Figure 4: Pearson correlations of the Ramification Index and four folk indices (plus NAND flash placebo) against three macroeconomic variables, 1980–2026. Negative $r(\Delta U-3)$ values are directionally correct: the index rises when unemployment falls. The Ramification Index substantially dominates all alternatives on every macro target.

7 Granger Causality

7.1 Framework

We test whether the Ramification Index *Granger-causes* quarterly real GDP growth—that is, whether past values of R_t contain information about future GDP growth beyond what is contained in GDP’s own history [11]. We estimate a bivariate vector autoregression (VAR):

$$y_t = \alpha_0 + \sum_{k=1}^p \alpha_k y_{t-k} + \sum_{k=1}^p \beta_k R_{t-k} + \varepsilon_t^y \quad (8)$$

$$R_t = \gamma_0 + \sum_{k=1}^p \gamma_k R_{t-k} + \sum_{k=1}^p \delta_k y_{t-k} + \varepsilon_t^R \quad (9)$$

where y_t is the annualised quarterly real GDP growth rate and R_t is the Ramification Index (interpolated to quarterly frequency using the Chow–Lin procedure [7] with McCallum’s annual series as the low-frequency input and TrendForce quarterly contract ASPs as the high-frequency indicator for the post-2018 period).

The null hypothesis $H_0 : \beta_1 = \beta_2 = \dots = \beta_p = 0$ is tested by a standard F -test. Rejection of H_0 establishes that the Ramification Index Granger-causes GDP growth.

7.2 Specification

The Augmented Dickey-Fuller (ADF) test rejects the unit-root null for both R_t (which, as a log-differenced price, is stationary by construction) and quarterly GDP growth (which is stationary over the sample). The Bayesian Information Criterion (BIC) selects $p = 2$ lags for the quarterly VAR over the sample 1985Q1–2024Q4.

7.3 Results

Table 3: Granger causality test results, quarterly VAR(2), 1985Q1–2024Q4.

Null hypothesis	F -statistic	p -value	Decision
R_t does not Granger-cause y_t	5.83	0.004	Reject at 1%
y_t does not Granger-cause R_t	2.14	0.121	Do not reject

The Ramification Index Granger-causes quarterly GDP growth at the 1% significance level. GDP growth does not Granger-cause the Ramification Index at conventional levels. The causal direction runs from DRAM pricing to the macroeconomy, not the reverse. This is consistent with the structural argument: DRAM capacity decisions are made 18–24 months in advance of shipment, creating a natural lead over GDP realisations.

7.4 Impulse response

An orthogonalised impulse response function (IRF) shows that a one-standard-deviation positive shock to R_t (a DRAM price increase) is associated with a statistically significant increase in GDP growth 1–3 quarters later, with the peak effect at 2 quarters. The 90% confidence band excludes zero from quarter 1 through quarter 3. By quarter 5 the effect is statistically insignificant, consistent with the transient nature of DRAM capex cycles. Figure 5 shows the lead relationship visually.

7.5 Robustness of the Granger result

The Granger causality result is robust to: (i) VAR lag order $p \in \{1, 2, 3, 4\}$; (ii) exclusion of the 2020 COVID quarter; (iii) use of the Toda-Yamamoto [32] procedure, which does not require pre-testing for unit roots; and (iv) substitution of industrial production for GDP as the dependent variable. In all specifications, the Ramification Index Granger-causes the macroeconomic variable at the 5% level or better.




Figure 5: The Ramification Index (amber bars) Granger-causes real GDP growth (blue line) at 1–3 quarter lags — $F = 5.83$, $p = 0.004$, VAR(2) 1985Q1–2024Q4. The causal direction runs RAM \rightarrow GDP, not the reverse ($F = 2.14$, $p = 0.121$). DRAM capacity decisions are made 18–24 months in advance, creating a natural structural lead.

8 The Bifurcated Index

8.1 Motivation: the structural decoupling

Prior to 2024, DRAM was effectively a single commodity market. Commodity DDR and server DDR moved on the same cycle, driven by the same capex decisions. HBM existed but represented a small fraction of total DRAM revenue.

Beginning in 2024, three structural changes broke this unity:

1. **HBM became the binding constraint on AI accelerator shipments.** Nvidia’s H100 and B100 GPUs are memory-bound: a single H100 carries 80 GB of HBM3, while frontier inference models require approximately 3.5 TB to load—a 44 \times compute-memory chasm [43].
2. **AI capex eclipsed consumer-PC demand as the primary driver.** Cloud service providers (CSPs) including Meta, Google, Microsoft, and Amazon began securing multi-year DRAM allocations, with OpenAI alone reported to have absorbed approximately 40% of global DRAM production [10].
3. **A manufacturer exited consumer memory entirely.** In late 2025, Micron announced it was withdrawing from the consumer Crucial RAM and SSD business to focus exclusively on AI and enterprise buyers [10].

The result is that commodity DDR and HBM now move on partially independent cycles. The composite Ramification Index R_t from Equation (4) averages over two structurally distinct signals.

8.2 Construction

The Commodity Ramification Index R_t^C and AI Ramification Index R_t^{AI} are constructed as in Equations (5)–(6). For 2024–2026, the revenue weights in Equation (7) are approximately:

Table 4: Revenue weights for the bifurcated Ramification Index.

Year	w^C (Commodity)	w^{AI} (HBM)
2024	0.82	0.18
2025	0.70	0.30
2026 (est.)	0.58	0.42

The AI sub-index weight is growing rapidly and is projected to reach parity with commodity DRAM by 2028. Figure 6 shows the structural shift.




Figure 6: DRAM revenue shift from Commodity DDR to HBM (AI), 2020–2028E. HBM’s share of total DRAM revenue has grown from 3% in 2020 to an estimated 42% in 2026, with parity projected around 2028. This structural shift is why the Ramification Index must bifurcate: the two sub-indices now track partially independent economic signals.

8.3 The divergence

In Q1 2026, the two sub-indices read differently:

- R_t^C is sharply positive: commodity DDR4/DDR5 prices surged 90–100% QoQ as wafer capacity was reallocated to HBM, creating artificial scarcity in the commodity tier [38].
- R_t^{AI} is moderately positive: HBM3e prices rose 15–20% YoY [34], a significant increase but much smaller than the commodity spike, reflecting the fact that HBM buyers (primarily Nvidia and AMD) have long-term contracts that smooth price volatility.

The composite R_t is strongly positive—the most positive quarterly reading in the 46-year series. But the *source* of the positive reading is supply reallocation, not healthy demand growth. This is the supply-side divergence regime. Figure 7 shows the divergence between the commodity and AI sub-indices.




Figure 7: The Bifurcated Ramification Index, 2020–2026. The composite index R_t (orange) splits into Commodity R_t^C (blue, DDR4/5) and AI R_t^{AI} (green, HBM) beginning in 2024. In 2026, commodity DRAM surged +137% while HBM rose only +17% — the supply-side divergence regime. The composite reading (+86%) is the most positive in the 46-year series, but the economic meaning is stagflationary, not expansionary.

8.4 The legacy price inversion

Perhaps the most striking feature of the 2025–2026 data is the legacy price inversion: DDR4 spot prices reached \$2.10 per gigabit, exceeding HBM3e contract prices of approximately \$1.70 per gigabit [28]. Older memory now costs more than cutting-edge memory. This inverts the traditional pricing hierarchy and reflects the extreme scarcity of a technology tier (DDR4) that manufacturers are actively abandoning.

In the algebraic metaphor (Section 4), this corresponds to a residue degree anomaly in the splitting field: the “smaller” prime ideal has a temporarily larger residue field.

9 The 2025–2026 Episode: RAMageddon

The 2025–2026 DRAM shortage—referred to in trade media as “RAMageddon” or the “RAMpocalypse”—is the most significant event in the Ramification Index’s history and the most important case study for the index’s interpretive framework.

9.1 Timeline

October 2025: OpenAI secures an estimated 40% of global DRAM production for its long-term AI infrastructure programme [10].

Late 2025: Micron announces withdrawal from the consumer Crucial RAM and SSD market. The remaining consumer stock is expected to sell out by early 2026 [10].

November 2025: Japanese electronics retailers (Tsukumo, Sofmap) begin limiting purchases of standalone RAM and SSDs to prevent panic stockpiling [15].

December 2025: DDR5 chip prices quadruple from \$6.84 to \$27.20 in three months [28].

January 2026: Korean media reports US big tech executives staying in long-term hotels near Samsung and SK hynix facilities in Pangyo and Pyeongtaek, attempting to secure DRAM allocations. Industry insiders call them “DRAM beggars” [40].

January 2026: Google reportedly attempts to secure additional HBM for its TPU accelerators and is told supply is unavailable. Executives responsible for memory procurement are later fired [10].

January 2026: Microsoft executives fly to South Korea to negotiate with SK hynix; one executive reportedly storms out of the meeting after learning of allocation constraints [10].

Q1 2026: TrendForce reports DRAM contract prices up 90–95% QoQ, the largest quarterly increase on record. PC DRAM prices double QoQ. Server DRAM prices surge by more than 60% QoQ [39].

Q1 2026: Apple reports paying a 230% premium for the 12 GB LPDDR5X memory used in its iPhone 17 Pro models [10].

Q2 2026: TrendForce projects further 58–63% QoQ increases in conventional DRAM contract prices. Momentum has slowed from Q1 but remains extreme by historical standards [36].

Q2 2026: The Conference Board Leading Economic Index declines 0.6% in March 2026, signalling continued economic slowdown [31].

9.2 The Ramification Index reading

The composite Ramification Index R_t for 2025–2026 is the most sharply positive in the 46-year series—exceeding even the 2016–2018 supercycle. Year-on-year DRAM contract prices are estimated to have risen 130–144% in 2026 relative to 2025, depending on the data source (TrendForce vs. Gartner).

But unlike the 2016–2018 supercycle, which coincided with healthy macroeconomic expansion, the 2025–2026 episode is occurring against a backdrop of slowing GDP growth, rising unemployment (4.4% in early 2026), declining consumer confidence, and weakening leading economic indicators. This is the *supply-side divergence regime*: the Ramification

Index is positive because oligopoly supply reallocation to AI has created artificial scarcity, not because broad-based demand is rising.

9.3 Downstream ramifications

The 2025–2026 episode demonstrates the ramification tree in real time:

1. **Consumer electronics pricing.** IDC projects PC market decline of 4.9–8.9% and smartphone market decline of 2.9–5.2% in 2026 due to memory-driven BOM increases [10]. Notebook shipments are projected down 14.8% YoY [36].
2. **Specification downgrading.** Low-end smartphones are expected to return to 4 GB base configurations in 2026. Dell and Lenovo may limit devices to 8 GB of RAM. Some laptop manufacturers are reverting to LPDDR4 [35].
3. **Console delays.** PlayStation 6 and next-generation Xbox face potential delays. Nintendo has lost approximately \$14 billion in market value amid concerns over memory costs [10].
4. **GPU pricing.** Nvidia’s RTX 5090 is rumoured to approach \$5,000, up from \$2,000 for the previous generation, driven in part by HBM and GDDR7 costs [10].
5. **Manufacturer exits and bankruptcies.** The CEO of RAM designer Phison warns that many consumer electronics manufacturers “will go bankrupt or exit product lines” by the end of 2026 [10].

This is the most complete empirical demonstration of the Ramification Index thesis: a single price movement at the oligopoly level—set by three firms in Seoul and Boise—branches outward into every sector of the information economy, from server capex to gaming consoles to smartphone specifications to Nintendo’s market capitalisation.

9.4 Interpretive framework: reading the sign

The 2025–2026 episode requires updating the index’s interpretive framework. In the original (v1) formulation:

$$\begin{aligned} R_t \ll 0 &\Rightarrow \text{demand-side contraction (recession signal)} \\ R_t \approx 0 &\Rightarrow \text{equilibrium} \\ R_t \gg 0 &\Rightarrow \text{healthy demand growth (expansion signal)} \end{aligned}$$

The 2025–2026 data introduces a fourth case:

$$R_t^C \gg 0 \text{ and macro indicators weakening} \Rightarrow \text{supply-side divergence}$$

In this regime, the Ramification Index is positive not because the economy is expanding but because oligopoly supply reallocation has decoupled memory pricing from consumer demand. The downstream ramifications are *contractionary*: rising BOM costs suppress consumer electronics shipments, reduce specifications, and squeeze manufacturer margins.

The sign of R_t is positive, but the *economic meaning* is stagflationary: higher input costs coexisting with slowing output.

This is why the bifurcated index matters. The composite R_t alone cannot distinguish demand-driven expansion from supply-driven cost-push. The sub-indices can: $R_t^{AI} \gg 0$ with $R_t^C \gg 0$ and weakening macro indicators identifies the supply-side divergence regime unambiguously.

10 Photonic Disruption Risk

A structural index requires assessment of what could break it. Two developments threaten to alter the macroeconomic role of DRAM pricing.

10.1 Fiber-optic memory architectures

John Carmack has proposed using long loops of single-mode fiber as a time-domain L2 cache or DRAM replacement for streaming AI model weights [6, 43]. Modern fiber experiments achieve approximately 256 Tb/s over 200 km; at that rate and distance, approximately 32 GB of data is “in flight” at any moment, effectively stored in the fiber as a delay line. If AI inference workloads can be architected to stream weights sequentially rather than accessing them randomly, fiber-optic delay lines could substitute for a significant fraction of the DRAM currently consumed by AI data centres.

The implied bandwidth (approximately 32 TB/s) is enormous compared to typical DRAM channels, and the power consumption of fiber transmission is 5–20× lower per bit than DRAM refresh [9]. If this architecture scales, it would reduce AI infrastructure’s DRAM consumption and thereby weaken the Ramification Index’s macro weight.

10.2 Co-packaged optics

Corning’s CEO Wendell Weeks has publicly stated that fiber will replace copper inside server racks, with Nvidia’s roadmap placing co-packaged optics approximately five years out [9]. Corning’s \$6 billion supply agreement with Meta for its Contour next-generation fiber product represents the first commercial-scale commitment to this transition [9]. A single high-density GPU server rack currently contains approximately 2.5 miles of copper; as cluster density scales to hundreds of GPUs per node, fiber’s power and bandwidth advantages become decisive.

Co-packaged optics does not directly substitute for DRAM, but it alleviates the interconnect bottleneck that currently amplifies the memory-bandwidth constraint. By reducing the cost of moving data to and from memory, it indirectly reduces the premium commanded by memory proximity (which is what makes HBM so expensive).

10.3 Assessment

Neither technology is expected to materially affect DRAM consumption before 2028–2030. Carmack’s fiber-delay-line architecture requires sequential weight access patterns that are

natural for inference but not for training. Co-packaged optics is five years from volume deployment by Nvidia’s own estimate. During the 2026–2030 window, the Ramification Index’s macro weight is likely to *increase* rather than decrease, as AI capex concentrates ever more of the global economy’s capital allocation on memory pricing.

Beyond 2030, the Ramification Index may need to incorporate fiber-optic memory and photonic interconnect pricing alongside DRAM. In the algebraic metaphor, this corresponds to the field extension K/\mathbb{Q} growing—new input factors joining the IT supply chain—which increases the total degree n and creates new prime ideals that ramify alongside DRAM.

11 Robustness

11.1 Spot versus contract ASPs

The baseline Ramification Index uses contract ASPs (negotiated quarterly between memory suppliers and large buyers). We re-estimate R_t using DRAMeXchange spot prices for 2018–2026. The spot-based index is more volatile (reflecting speculative and inventory-adjustment dynamics) but the recession-signalling performance is identical: 6/6 recessions identified, with a 0–1 quarter lead over the contract-based index. Pearson correlations with real GDP growth differ by less than 0.05.

11.2 Measurement window shifts

Shifting the annual measurement window by ± 1 quarter (i.e., using Q1-to-Q1 or Q3-to-Q3 instead of mid-year-to-mid-year) does not change the recession coverage and changes the Pearson correlations by less than 0.08.

11.3 NAND flash placebo test

NAND flash (the other major memory semiconductor) shares some cycle dynamics with DRAM but has a structurally different market: more producers, less oligopoly concentration, and a pricing curve driven by layer count and bit-growth rather than capacity planning. We construct a “NAND Ramification Index” using the same methodology applied to NAND ASP. The NAND index signals only 3 of 6 recessions and has Pearson correlations with macro variables that are less than half those of the DRAM-based index. This confirms that it is specifically the oligopoly structure of DRAM—not generic semiconductor pricing—that drives the macro signal.

11.4 Data source sensitivity

We re-estimate Table 2 using three alternative data sources for 2024–2026 DRAM pricing: (i) Gartner’s 130% YoY estimate, (ii) TrendForce’s quarterly contract data aggregated to annual, and (iii) Counterpoint Research’s DDR4/DDR5 ASP series. The Pearson correlations with real GDP growth range from +0.38 to +0.44 across sources, bracketing our baseline estimate of +0.41. The index is robust to the choice of data vendor.

12 Discussion

12.1 What the Ramification Index is

The Ramification Index is a supply-side macroeconomic signal derived from the capital-allocation decisions of a three-firm oligopoly. It is a strictly better folk index than the Lipstick, Hemline, Men’s Underwear, and Buttered Popcorn indices: it outperforms each on coverage, timing, and correlation with macroeconomic truth. More importantly, it is not vulnerable to the kinds of behavioural shocks (pandemic, fashion fragmentation, athleisure, streaming) that have embarrassed the consumer-goods indices since 2020.

12.2 What the Ramification Index is not

The Ramification Index is not a formal leading economic indicator in the Conference Board sense. It does not replace the yield curve, the ISM manufacturing index, or the Sahm Rule. We claim it is a strictly better *folk* index, and one whose macro relevance is increasing as AI capital expenditure concentrates ever more of the global economy’s future bets on the price of memory.

12.3 The supply-side divergence regime

The most significant theoretical contribution of this paper is the identification of the supply-side divergence regime (Section 9). When the Ramification Index is positive because of demand-driven expansion, it is a pro-cyclical signal. When it is positive because of supply-side reallocation, it is a *cost-push* signal whose downstream ramifications are contractionary. The bifurcated index (Section 8) is necessary to distinguish these two cases.

This has immediate practical implications. An analyst reading the composite R_t in Q1 2026 and seeing its most positive value in 46 years might conclude the economy is booming. The bifurcated index reveals the opposite: commodity memory is scarce because production has been reallocated to AI, consumer electronics are contracting, and the macro backdrop is weakening. The sign of the index is the same; the economic meaning is inverted.

12.4 Substitutability with adjacent signals

NAND flash, foundry wafer ASPs, and copper pricing move on related but distinct cycles. In principle, the Ramification Index can be extended to a “Memory Complex Index” combining DRAM, NAND, and HBM in capacity-weighted shares—a broader ramification tree whose branches represent each major memory tier. We view this as the natural next step once post-2023 HBM pricing is disclosed by all three vendors.

12.5 The DRAM market’s enduring relevance

The rise of AI capital expenditure has increased, not decreased, the macroeconomic weight of DRAM pricing. Three firms now determine the allocation of the world’s most critical AI input factor. Their quarterly pricing decisions propagate outward into every layer of the

IT stack—from the bill of materials of an iPhone to the training cost of a frontier language model to the capital expenditure budget of a Fortune 500 company. The folk indices measure the leaves at the end of these branches. The Ramification Index measures the root.

13 Conclusion

The informal economic indices that have populated financial journalism for a century—lipstick, hemline, underwear, popcorn—are heuristic and historically contingent. They fail structurally during supply-side shocks (COVID-19), and they fail econometrically in at least one-third of the business cycles they are asked to explain. The DRAM market’s oligopolistic structure, its published long-run price series, its cyclical amplitude, its criminal price-fixing history, its rising macroeconomic weight, and its structural role as the binding constraint on AI capital expenditure make DRAM ASP a more durable and informative candidate.

We propose calling this the **Ramification Index**. The name is a commitment. In algebraic number theory, a prime is said to ramify in a field extension when it loses its primality and splits into factors; the ramification index measures how completely that splitting occurs. In economics, a DRAM price movement ramifies through every layer of the information-technology stack, from fab yields to foundation-model training budgets. The folk indices measure the leaves. The Ramification Index measures the root.

The 2025–2026 episode has demonstrated this more completely than any prior cycle. A single commodity—whose price is set by three firms—has branched outward into smartphone specifications, gaming console timelines, laptop shipment forecasts, GPU pricing, manufacturer bankruptcies, and the strategic procurement decisions of the world’s largest technology companies. The ramification is total. The index that measures it is ready.

Its oligopoly structure, archival data, criminal history, Granger-causal relationship with GDP growth, and rising macroeconomic weight make it harder to disrupt and better at what the older indices only claimed to do. The RAM in Ramification is not an accident. It never was.

Data and Code Availability

All data used in this study are provided as a supplementary CSV (`data/indices-wide.csv`). The interactive comparison dashboard, Granger causality estimation code (Stata and Python), and bifurcated index construction scripts are available at the companion repository: <https://github.com/fromknowware/bifurcation-memory-index>.

Acknowledgements

Thanks to the maintainers of the McCallum memory-price archive, AI Impacts, TrendForce DataTrack, Counterpoint Research, and the Sourceability procurement intelligence team. The Granger causality framework benefited from the methodology of Chow & Choy [8]. The algebraic formalism owes a debt to Neukirch [21]. The Cold Fusion documentary team’s

investigative reporting on the 2025–2026 DRAM shortage provided essential primary-source documentation for Section 9. Any errors are the author’s.

Competing Interests

The author declares no competing interests. The author holds no financial positions in Samsung, SK hynix, Micron, or any memory-related security.

A Supplementary: The Demand Side of the Ramification Tree

The Ramification Index measures the *supply-side* price signal from the DRAM oligopoly. This appendix documents the *demand side*: the physical infrastructure consuming the memory whose pricing the index tracks. The data is drawn from public facility disclosures, NVIDIA reference architectures, and industry reporting as of Q2 2026.

A.1 The ten largest AI data center builds

Table 5 summarises the ten largest currently announced or operating AI data center campuses by power capacity. Where hardware is publicly disclosed, we include the accelerator model, count, and implied VRAM. Where the bill of materials is not public, we mark the entry accordingly.

Table 5: Ten largest AI data center builds by power capacity, Q2 2026. VRAM estimates are derived from disclosed accelerator counts and per-device HBM specifications. “Not public” indicates the operator has not disclosed the hardware configuration. Sources: Microsoft [19], Terakraft [30], Orennia [25], NVIDIA [23].

Rank	Site	Operator	Location	Power (MW)	Accel. count	VRAM (PB)	Status
1	Hypergrid	Fermi America	Amarillo, TX	11,000	n/p	n/p	Planned
2	Project Jade	Crusoe + Tallgrass	Wyoming	10,000	n/p	n/p	Planned
3	Fairwater	Microsoft	Wisconsin	>1,000	~300k	~57.6	Operating
4	Stargate	OpenAI + Crusoe	Abilene, TX	~200	~100k	~19.2	Expanding
5	Colossus 1	xAI	Memphis, TN	~300	~230k	~29.9	Operating
6	AWS AI	Amazon	Canton, MS	>300	~500k	N/A*	Developing
7	Prometheus	Meta	n/p	1,000	n/p	n/p	Planned
8	Fayetteville	Microsoft	n/p	1,000	n/p	n/p	Planned
9	New Carlisle	Anthropic + Amazon	Indiana	1,000	n/p	n/p	Planned
10	Colossus 2	xAI	Memphis, TN	1,000	n/p	n/p	Planned

n/p = not public. *Amazon’s Trainium 2 is an ASIC; VRAM is not the applicable memory metric. VRAM estimates: H100 = 80 GB HBM2e; H200 = 141 GB HBM3e; GB200/B200 = 192 GB HBM3e per device.

Microsoft Fairwater: 72 Blackwell GPUs per rack with 14 TB pooled memory per rack.

A.2 Memory demand at facility scale

The three sites with publicly disclosed hardware—Microsoft Fairwater, OpenAI Stargate, and xAI Colossus 1—represent approximately 107 petabytes of accelerator VRAM across roughly 630,000 GPUs. Each Blackwell-class GB200 device carries 192 GB of HBM3e, manufactured exclusively by SK hynix and Samsung on dedicated wafer lines.

Seven additional 1 GW-class campuses are planned with undisclosed hardware. If each were equipped at comparable accelerator density to the disclosed sites, the implied additional HBM demand would be measured in hundreds of petabytes—representing multi-year forward commitments of oligopoly wafer capacity.

To place this in context: the entire global DRAM industry shipped approximately 18 exabytes (18,000 PB) in 2025 across all product categories—commodity DDR, server DDR, mobile LPDDR, and HBM combined [33]. The ten facilities in Table 5 collectively represent a non-trivial fraction of global memory output, concentrated in the highest-margin HBM tier, allocated to fewer than ten purchasing entities.

A.3 Implications for the Ramification Index

This demand concentration has three consequences for the index:

1. **Oligopsony meets oligopoly.** Three memory producers face five to ten hyperscale buyers. The negotiation dynamics—visible in the “DRAM beggars” episode documented in Section 9—are not competitive market clearing. They are bilateral oligopoly bargaining, and the resulting contract prices are the raw input to the Ramification Index. The index reads the *output* of these negotiations.
2. **The “not public” rows are the index’s future.** Seven of the ten largest builds have undisclosed hardware configurations. The DRAM purchase orders behind those facilities have not been publicly allocated. When they are—when Meta, Microsoft, Anthropic, and xAI sign multi-year HBM supply agreements for their 1 GW campuses—each contract will represent a discrete, measurable impulse to the Ramification Index. The index’s forward-looking power depends on the fact that DRAM capacity decisions are made 18–24 months before shipment; these planned facilities are the 2027–2028 signal forming now.
3. **Power-to-memory density defines the ramification depth.** A 1 GW data center consuming 8.76 TWh per year exists primarily to keep memory hot. DRAM requires constant refresh—every cell, every 64 milliseconds, whether it is being read or not. The power budget is dominated by compute, but the memory subsystem is the binding constraint on what compute can actually run. When a DRAM price shock propagates into the capital expenditure of a 1 GW campus, it ramifies more deeply than it did into a 10 MW enterprise data center in 2015. The ramification index e_i of the AI prime ideal \mathfrak{P}_{AI} is increasing with facility scale.

A.4 The memory hierarchy at rack scale

NVIDIA’s reference architecture for the DGX SuperPOD provides the most detailed public view of memory at rack scale [23]. A 127-node H100 SuperPOD contains:

- **Accelerator VRAM:** $127 \times 8 \times 80$ GB = 81.3 TB of HBM2e (the GPU working memory).
- **System RAM:** 127×2 TB DDR5 = 254 TB (the CPU-attached main memory).
- **Storage:** 5 compute-storage servers with NVMe arrays (capacity not standardised).
- **Networking:** 4 fabric switches, multiple management switches, InfiniBand or high-speed Ethernet.

The system RAM (DDR5) is approximately $3 \times$ the accelerator VRAM (HBM) by capacity—but it is the HBM that commands the premium pricing and constrains the build schedule. A single H100 SuperPOD’s HBM content represents roughly \$4–8 million in memory cost at 2026 HBM3e contract prices. For a 1 GW campus with thousands of equivalent pods, the memory bill of materials alone runs into the billions of dollars.

This is the physical substrate of the Ramification Index: three firms, manufacturing the memory that fills these racks, pricing it quarterly, with each price movement ramifying outward through the capital budgets of the world’s largest technology companies and ultimately into the macroeconomy.

References

- [1] AI Impacts. Trends in DRAM price per gigabyte. <https://aiimpacts.org/trends-in-dram-price-per-gigabyte/>, 2020.
- [2] Ana Aizcorbe, Stephen D. Oliner, and Daniel E. Sichel. Shifting trends in semiconductor prices and the pace of technological progress. Finance and Economics Discussion Series 2006-44, Board of Governors of the Federal Reserve System, 2006. Published version: *Business Economics*, 43(3):23–39, 2008.
- [3] AppleInsider. Samsung, Micron, SK Hynix sued again over alleged DRAM price fixing. <https://appleinsider.com/articles/21/05/09/memory-producers-sued-over-alleged-dram-price-fixing>, May 2021.
- [4] Bloomberg Law. Samsung, Micron, SK Hynix win antitrust appeal over DRAM prices. <https://news.bloomberglaw.com/antitrust/samsung-micron-sk-hynix-win-antitrust-appeal-over-dram-prices>, March 2022.
- [5] Victoria Bryan. Is the lipstick effect making a comeback? BBC News, 2022.
- [6] John Carmack. Fiber-optic delay lines as L2 cache for streaming AI data. Twitter/X post and subsequent discussion, 2025. See Tom’s Hardware, Hacker News, and LinkedIn commentary threads.

- [7] Gregory C. Chow and An-loh Lin. Best linear unbiased interpolation, distribution, and extrapolation of time series by related series. *The Review of Economics and Statistics*, 53(4):372–375, 1971.
- [8] Hwee Kwan Chow and Keen Meng Choy. Forecasting the global electronics cycle with leading indicators: A Bayesian VAR approach. *International Journal of Forecasting*, 22(2):301–315, 2006.
- [9] CNBC. How Corning invented a new fiber-optic cable for AI and landed a \$6 billion Meta deal. CNBC documentary / transcript, 2025. Includes interview with CEO Wendell Weeks; full transcript in project archive.
- [10] Cold Fusion. The RAM crisis keeps getting worse. YouTube / Cold Fusion documentary, 2026. Full transcript in project archive.
- [11] C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969.
- [12] Alan Greenspan. *The Age of Turbulence: Adventures in a New World*. Penguin Press, 2007.
- [13] John L. Hennessy and David A. Patterson. *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann, 6 edition, 2019.
- [14] Sarah E. Hill, Christopher D. Rodeheffer, Vladas Griskevicius, Kristina Durante, and Andrew E. White. Boosting beauty in an economic decline: Mating, spending, and the lipstick effect. *Journal of Personality and Social Psychology*, 103(2):275–291, 2012.
- [15] IntuitionLabs. RAM shortage 2025: How AI demand is raising DRAM prices. <https://intuitionlabs.ai/articles/ram-shortage-2025-ai-demand>, December 2025.
- [16] Robert Krulwich. Alan Greenspan’s underwear drawer. NPR All Things Considered, January 2008.
- [17] Mark Lapedus. SK hynix surpasses Samsung in DRAM share. Semiecosystem, <https://marklapedus.substack.com/p/sk-hynix-surpasses-samsung-in-dram>, 2025.
- [18] John C. McCallum. Memory prices (1957–2024). <https://jcmmit.net/memoryprice.htm>, 2024.
- [19] Microsoft. Inside the world’s most powerful AI datacenter. Microsoft Blog, <https://blogs.microsoft.com/blog/2025/09/18/inside-the-worlds-most-powerful-ai-datacenter/>, September 2025.
- [20] Emily Nelson. Rising lipstick sales may mean pouting economy. *The Wall Street Journal*, November 2001.
- [21] Jürgen Neukirch. *Algebraic Number Theory*, volume 322 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1999. The ramification index $e(\mathfrak{P}|p)$ is defined in Chapter I, §8 as the exponent of p in the factorisation of $p\mathcal{O}_K$.

- [22] NoobFeed. How Samsung, Micron, and SK hynix built dominance after the DRAM cartel scandal. <https://www.noobfeed.com/articles/samsung-micron-hynix-dram-cartel>, 2024.
- [23] NVIDIA. DGX SuperPOD reference architecture: scalable infrastructure for H100. <https://docs.nvidia.com/dgx-superpod/reference-architecture-scalable-infrastructure-h100/latest/>, 2024.
- [24] Paul H. Nystrom. *Economics of Fashion*. Ronald Press, New York, 1928.
- [25] Orennia. The 20 largest data centers being developed in North America. <https://orennia.com/insights/the-20-largest-data-centers-being-developed-in-north-america>, 2026.
- [26] PYMNTS.com. China: Gov finds “massive evidence” in DRAM anti-monopoly probe. <https://www.pymnts.com/cpi-posts/china-gov-finds-massive-evidence-in-dram-anti-monopoly-probe/>, November 2018.
- [27] David Scott. The coronavirus recession may have killed the lipstick index. CNN, 2020.
- [28] Sourceability. Memory price increase timeline QoQ in 2026. <https://sourceability.com/post/tracking-memory-price-increases-across-the-last-several-quarters>, March 2026.
- [29] George W. Taylor. *Significant post-war changes in the full-fashioned hosiery industry*. PhD thesis, Wharton School, University of Pennsylvania, 1929.
- [30] Terakraft. From megawatts to gigawatts: the 10 largest AI datacenters in the world, 2026 edition. <https://www.terakraft.no/post/from-megawatts-to-gigawatts-the-10-largest-ai-datacenters-in-the-world-2026-edition>, 2026.
- [31] The Conference Board. US leading indicators. <https://www.conference-board.org/topics/us-leading-indicators/>, 2026. Accessed April 2026.
- [32] Hiro Y. Toda and Taku Yamamoto. Statistical inference in vector autoregressions with possibly integrated processes. *Journal of Econometrics*, 66(1–2):225–250, 1995.
- [33] TrendForce. AI reportedly to consume 20% of global DRAM wafer capacity in 2026, HBM and GDDR7 lead demand. <https://www.trendforce.com/news/2025/12/26/>, December 2025.
- [34] TrendForce. Higher DDR5 profitability intensifies capacity crowding, strengthening HBM3e pricing momentum in 2026. <https://www.trendforce.com/presscenter/news/20251218-12843.html>, December 2025.
- [35] TrendForce. Memory price surge to persist in 1Q26; smartphone and notebook brands begin raising prices and downgrading specs. <https://www.trendforce.com/presscenter/news/20251211-12831.html>, December 2025.

- [36] TrendForce. AI server demand to drive memory contract price increases in 2Q26 as CSPs secure supply via long-term agreements. <https://www.trendforce.com/presscenter/news/20260331-12995.html>, March 2026.
- [37] TrendForce. DRAM price trends. https://www.trendforce.com/price/dram/dram_spot, 2026.
- [38] TrendForce. Memory makers prioritize server applications, driving across-the-board price increases in 1Q26. <https://www.trendforce.com/presscenter/news/20260105-12860.html>, January 2026.
- [39] TrendForce. Memory price outlook for 1Q26 sharply upgraded; QoQ increases of all product categories to hit record highs. <https://www.trendforce.com/presscenter/news/20260202-12911.html>, February 2026.
- [40] TrendForce. Samsung, SK reportedly hike server DRAM prices 60–70% — Google, Microsoft in the queue. <https://www.trendforce.com/news/2026/01/06/>, January 2026.
- [41] U.S. Department of Justice. Three executives indicted for their roles in the DRAM price-fixing & bid-rigging conspiracy. https://www.justice.gov/archive/atr/public/press_releases/2006/219102.htm, 2006.
- [42] Marjolein van Baardwijk and Philip Hans Franses. The hemline and the economy: is there any match? *Econometric Institute Report*, (2010-40), 2010.
- [43] Khayyam Wakil. Photonic regimes in AI infrastructure investment: a real-options analysis. arXiv preprint, 2026. v3.
- [44] Wikipedia contributors. DRAM price fixing scandal. https://en.wikipedia.org/wiki/DRAM_price_fixing_scandal, 2026. Accessed April 2026.