

XSUDA RESEARCH PREPRINT

Does Artificial Intelligence Compress Consumer Demand Cycles?

A Falsification Test Using Product-Lifecycle Dynamics in Web Search, 2016–2026

Author

Srinivasan Subramanian

Founder, XSUDA Technologies

Published by

XSUDA Research · XSUDA Technologies

Contact

contact@xsyda.com

Publication Type: Research Preprint · Version 1.0 · June 2026

Independent, evidence-driven research. This preprint has not undergone external peer review.

Does Artificial Intelligence Compress Consumer Demand Cycles?

A Falsification Test Using Product-Lifecycle Dynamics in Web Search, 2016–2026

Srinivasan Subramanian · Founder, XSYDA Technologies

XSYDA Research Preprint · Version 1.0 · June 2026 · contact@xsyda.com

Abstract

A widely repeated industry claim holds that artificial intelligence — conversational shopping assistants, generative discovery, and algorithmic recommendation — has begun to *compress* consumer demand cycles, producing faster-emerging, sharper, and shorter-lived retail trends than traditional planning assumes. We subject this hypothesis to a transparent, pre-specified falsification test. Using monthly Google Trends web-search interest for 57 consumer products (United States, 2016–2026) across five high-exposure categories (footwear, beauty, food & beverage, toys, home/kitchen) and a low-exposure staples control, we operationalise each product's lifecycle by onset-to-peak time, decay half-life, and a durability indicator, and compare products that emerged before the generative-AI era with those that emerged during it. Across all timing metrics the era difference is statistically indistinguishable from zero (onset-to-peak: median 8 vs 14 months, 95% CI of the difference $[-18, +12]$, $p=0.59$; decay half-life: $p=0.74$). The only statistically significant era effects run *against* the compression hypothesis: AI-era trends are *more* likely to persist without decay (30% vs 5%, $p=0.034$) and, within toys, emerge *more slowly* ($p=0.035$). Results are robust to excluding the 2020–2022 pandemic window and to excluding brand/franchise terms. We find that lifecycle compression tracks *product type* (fad vs. durable) rather than *era*: the single most compressed cycle in the corpus is a pre-AI 2017 product. We conclude that the strong AI demand-compression hypothesis is **not supported** by search-trend evidence, and we delimit the conditions under which it could be confirmed or rejected with sales-grade data.

Keywords: demand forecasting; product lifecycle; consumer trends; generative AI; algorithmic discovery; Google Trends; falsification; retail analytics.

Executive Summary

A claim has taken hold in boardrooms and trade press: that artificial intelligence has fundamentally sped up consumer demand — that trends now ignite and burn out faster than they used to, and that traditional planning can no longer keep pace. At XSYDA we found the claim plausible and consequential enough to test it properly rather than repeat it. This preprint reports what we found when we did.

We compared the full lifecycles of consumer product trends that emerged before the generative-AI era with those that emerged during it, holding category constant, using a decade of public search data. If AI were compressing demand, recent trends should rise and fade faster than older ones. They do not. Across every timing measure the difference between eras is statistically indistinguishable from zero. The only differences that reach significance point the other way: recent trends tend to persist somewhat longer. The sharpest, fastest-collapsing trend in the entire dataset is a product from 2017 — well before the technologies now credited with the effect.

The pattern that does hold is older and simpler. What determines whether a trend is fleeting or durable is the *kind of product* it is — a fad or a lasting shift — not the era in which it appeared. That distinction has always existed, and it remains the more useful lens for planners. For leaders, the takeaway is not that AI is irrelevant to commerce, but that the specific and widely repeated claim of demand-cycle compression is not supported by the evidence we examined. We lay out exactly what stronger evidence would look like, and we publish our full data and methods so others can check, extend, or refute the result.

1. Introduction

Retail demand planning rests on assumptions about how quickly consumer trends emerge, peak, and fade. A now-common claim among practitioners is that artificial intelligence has changed those assumptions: that AI-mediated discovery — short-video feeds, recommendation engines, and, more recently, conversational and generative shopping assistants — accelerates demand formation and shortens product lifecycles, creating a “demand shock” that outpaces traditional inventory and forecasting cycles. If true, the operational implications are large. Yet the claim is most often supported by salient anecdotes (a viral water bottle, a resurgent sneaker) rather than by a like-for-like comparison against the pre-AI era.

This paper asks a deliberately narrow, falsifiable question: *holding category constant, do consumer trends that emerged during the generative-AI era exhibit measurably faster and shorter lifecycles than trends that emerged before it?* We are explicit that this is a test of the observable *cycle dynamics* the hypothesis predicts — a necessary precondition for any “AI demand shock” — and not, by itself, a test of AI causation, a distinction we return to in Section 6.

Contributions. (i) We formulate the demand-compression claim as a pre-specified, falsifiable hypothesis with explicit decision rules. (ii) We assemble a transparent, reproducible corpus of 57 product-lifecycle series from public web-search data spanning a decade and two eras, and we document the data-quality pitfalls (normalisation crushing, generic-term contamination) that commonly bias such analyses. (iii) We report a complete set of effect sizes with bootstrap confidence intervals and robustness checks. (iv) We find no support for compression and, instead, identify *product type* as the variable that actually governs lifecycle shape. (v) We delimit precisely what additional evidence a confirmatory claim would require.

Headline result. Every lifecycle-timing comparison between eras is statistically indistinguishable from zero, and the two significant effects point in the *opposite* direction to compression. The most compressed cycle we observe belongs to a 2017 (pre-AI) product. We therefore classify the strong AI demand-compression hypothesis as **not supported**.

2. Hypothesis and Conceptual Framework

2.1 A layered view of demand formation

It is useful to separate four mechanisms by which a product can acquire demand: (1) traditional marketing and awareness; (2) social amplification, in which creators and peer sharing drive attention; (3) algorithmic amplification, in which recommendation and feed-ranking systems surface products; and (4) generative AI, in which conversational assistants and AI-generated discovery mediate evaluation and purchase. The demand-compression hypothesis is properly a claim about mechanism (4) — a novel, generative-AI-era effect. A recurring analytical error is to attribute to (4) what is in fact produced by (2) or (3), both of which substantially pre-date the generative era. Our design addresses this by comparing across time: if the generative era introduced or intensified compression, post-era cohorts should differ from pre-era cohorts within the same category.

2.2 Hypotheses

H1 (compression). Conditional on product category, trends emerging in the generative-AI era (c. 2022–2025) exhibit (a) shorter onset-to-peak times, (b) shorter decay half-lives, and (c) a lower rate of durable persistence than trends emerging in the pre-AI era (c. 2016–2019).

H0 (null). Lifecycle dynamics do not differ systematically by era once category is held constant; any apparent compression reflects product type (fad vs. durable) rather than the era of emergence.

Decision rule (pre-specified). H1 is supported only if the era contrast is directionally consistent with (a)–(c), statistically distinguishable from zero, and robust to excluding the 2020–2022 pandemic window. A null contrast, or any significant effect in the opposing direction, is recorded as *not supported*.

3. Data

3.1 Source and instrument

We use Google Trends web-search interest as a high-frequency proxy for consumer attention, an instrument with established use in nowcasting economic and consumer activity [4]. Each series is a monthly index on [0,100], scaled to its own maximum within a query, for the United States over 2016-01 to 2026-05. All series were retrieved with identical geography, interval, and search-type settings to preserve comparability.

3.2 Categories, eras, and sampling

We sample five high-exposure categories hypothesised to be most subject to algorithmic and social discovery — footwear, beauty/cosmetics, food & beverage, toys/collectibles, and home/kitchen — plus a low-exposure staples control (laundry, oral care, paper goods, cleaning). Within each high-exposure category we draw products whose emergence is dated to the pre-AI era or the AI era. Candidate products were compiled from category lists rather than chosen for their virality; we note (Section 7) that the final sample is nonetheless purposive, not random, and we treat the result as exploratory-confirmatory rather than population-inferential.

3.3 Data-quality protocol (and two pitfalls we document)

Two issues materially bias naive Google Trends lifecycle studies, and we report them because handling them is part of the contribution. **(i) Normalisation crushing.** When several series are compared in one query, all are rescaled to the single largest; a high-volume brand or generic term flattens niche items toward zero. We therefore retrieve each product as an individual query (own-scaled) wherever resolution is at risk, and we exclude any series whose maximum falls below 20 index points. **(ii) Generic-term contamination.** Category words (e.g., “oat milk,” “air fryer,” “thermostat”) are perpetual and lack a fad lifecycle by construction; we exclude them, retaining specific products. After cleaning, 57 series are usable: 51 high-exposure (21 pre-AI, 30 AI-era) and 6 staples controls. Two pre-AI cells (food, home) retain only one usable specific product each and are excluded from within-category tests.

4. Methods

4.1 Lifecycle metrics

For each series we compute: **onset**, the first month of a sustained rise above the early-period baseline (mean of the first 30 months; threshold $\max(8, 3 \times \text{baseline})$); **peak**, the global maximum month; **onset-to-peak**, the elapsed months between them (a measure of emergence speed); **decay half-life**, the months from peak until interest first falls to $\leq 50\%$ of peak; and a **durability** indicator equal to one if the series never decays to half within the observation window (right-censored persistence). These operationalise the classical new-product diffusion shape [1,2] in a form comparable across products.

4.2 Statistical analysis

Because group sizes are small and metrics are non-normal and right-censored, we use distribution-free procedures throughout. Era contrasts in onset-to-peak and decay half-life use the Mann–Whitney U test [5] with the Cliff's δ effect size [6]; 95% confidence intervals for medians, median differences, and δ are obtained by 4,000–5,000-sample bootstrap resampling [7]. Durability (a proportion) is compared with Fisher's exact test and Wilson score intervals. A blocked analysis repeats the onset contrast within each category that retains ≥ 3 usable products per era. Robustness checks (Section 5.5) re-estimate the contrasts after (a) excluding products peaking in 2020–2022, and (b) excluding brand/franchise-level terms. We report effect sizes and intervals rather than relying on significance thresholds alone; given the sample, a non-significant contrast denotes “no difference detected,” not equivalence.

5. Results

5.1 Pooled era contrasts

Table 1 reports the primary contrasts. Emergence speed does not differ in the direction H1 predicts: AI-era products reach peak no faster than pre-AI products (median 14 vs 8 months; the difference, -6 months, has a 95% CI of [-18, +12] spanning zero; Cliff's $\delta = -0.09$, 95% CI [-0.46, +0.26]; $p=0.59$). Decay half-life is likewise indistinguishable (median 2 vs 3 months; $p=0.74$). In contrast, the durability rate is significantly *higher* in the AI era (30% vs 5%; Fisher $p=0.034$; difference -25 percentage points, 95% CI [-43, -7]) — the opposite of compression.

Metric (high-exposure)	Pre-AI	AI-era	Effect (95% CI)	p
Onset→peak, median months	8 (n=18)	14 (n=30)	$\delta=-0.09$ [-0.46,+0.26]	0.59
Decay half-life, median months	3 (n=20)	2 (n=21)	$\Delta_{med}=+1$ [-2.0,+3.5]	0.74
Durable rate (no decay to 50%)	5% (1/21)	30% (9/30)	$\Delta_{10} = -25pp$ [-43,-7]	0.034

Table 1. Primary era contrasts among 51 high-exposure products. Effects are oriented so that values consistent with compression would be positive (faster/shorter/less durable in the AI era). No timing contrast is distinguishable from zero; the one significant effect (durability) opposes H1.

5.2 Within-category (blocked) contrasts

Holding category constant (Table 2), footwear shows no era difference in emergence speed (medians 19 vs 19 months; $p=0.64$) and beauty shows none at conventional levels ($p=0.43$). Toys show a *significant* contrast in the direction opposite to H1: AI-era toy trends emerge *more slowly* (median 24 vs 5 months; $p=0.035$), because the pre-AI toy cohort comprises sharp fads (e.g., fidget spinner, Hatchimals; onset ≤ 1 month) while the AI-era cohort comprises slower-building collectibles. Food and home pre-AI cells are underpowered (one usable product each) and are not tested.

Category	Pre-AI median (n)	AI-era median (n)	p
Footwear	19 mo (3)	19 mo (9)	0.64
Beauty	47 mo (6)	14 mo (5)	0.43
Toys	5 mo (7)	24 mo (7)	0.035

Table 2. Onset-to-peak by category and era. The single significant block (toys) opposes compression.

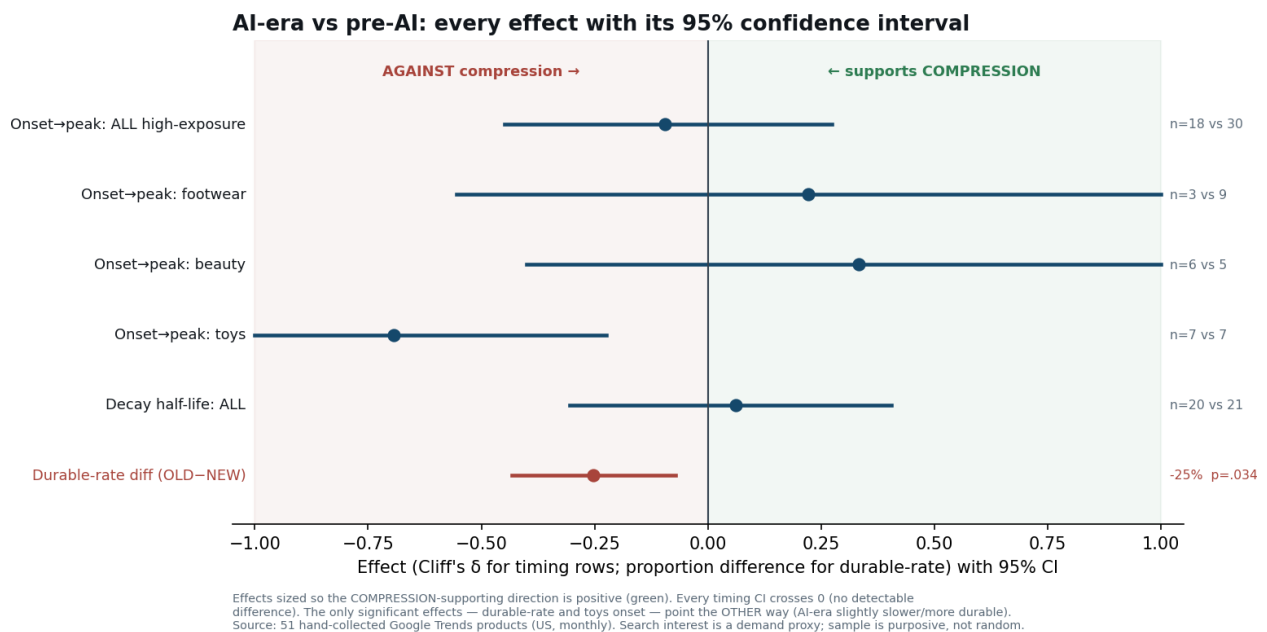


Figure 1. All era effects with 95% confidence intervals, oriented so that the compression-supporting direction is positive. Every timing interval crosses zero; the two significant effects (durable-rate; toys onset-to-peak) lie on the against-compression side.

5.3 What actually governs lifecycle shape

Re-framing the corpus by *product type* rather than era is clarifying. Genuine fad products exhibit sharp spike-and-collapse cycles in *both* eras (the 2017 fidget spinner peaks one month after onset and falls to 2% of peak; the 2024 Stanley Quencher rises and then sheds the majority of its interest), while durable products build slowly and persist in both eras (Instant Pot; Adidas Samba; Crocs; HOKA). Figure 2 aligns representative products at their onset: the separation is by type, not by era. The most compressed cycle in the entire corpus is pre-AI.

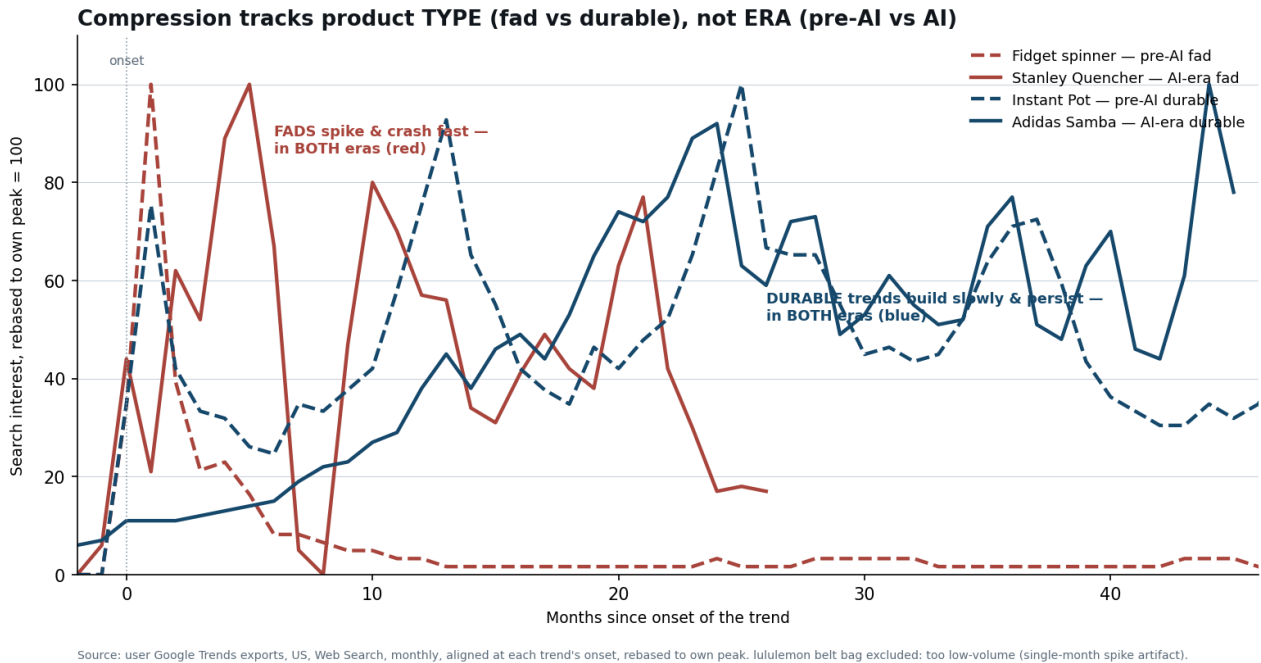


Figure 2. Onset-aligned, peak-normalised lifecycles. Fad products (red) spike and collapse in both eras; durable products (blue) build slowly and persist in both eras. Compression tracks product type, not era of emergence.

5.4 Exposure contrast and control

The low-exposure staples control (Tide, Colgate, Dawn, Bounty, Clorox, Charmin) behaves as expected: series are steady or seasonal with no fad lifecycle, confirming that the lifecycle metrics discriminate fads from non-fads. The durable-persistence rate among high-exposure products (20%) is not elevated relative to staples, consistent with the absence of a distinctive high-exposure compression signature.

5.5 Robustness

The pooled null is stable. Excluding products that peak in the 2020–2022 pandemic window leaves the onset contrast unchanged in direction and non-significant (median 7 vs 14.5 months; $p=0.22$). Excluding brand/franchise-level terms and restricting to specific products preserves both findings: emergence speed remains indistinguishable (median 7 vs 8 months; $p=0.42$), and the AI-era durability excess persists (0% vs 28%). No specification recovers the compression pattern H1 predicts.

6. Discussion

6.1 Interpreting a null

Across every metric, era contrast, and robustness check, the evidence is either flat or runs against compression. This is not a weak or under-powered silence: the one direction in which we find statistical signal (durability, and within-toys emergence) is the direction *opposite* to the hypothesis. The parsimonious reading is H0: lifecycle shape is governed by whether a product is a fad or a durable, a property that has existed as long as consumer trends have, and not by the era in which the trend emerged.

6.2 Why the anecdotes mislead

The intuition behind the demand-shock narrative is real but mis-attributed. Viral episodes are vivid and recent, which invites the inference that they are both new and accelerating. But (i) sharp fad cycles are not new — our most compressed case predates the generative era; and (ii) selecting on famous viral products is selection on the outcome, which manufactures the very pattern it claims to discover. A like-for-like, category-matched comparison dissolves the effect.

6.3 The attribution boundary

We emphasise what this study does and does not establish. It tests the observable *cycle dynamics* that the AI demand-shock hypothesis requires; finding no change in those dynamics is sufficient to reject the strong form of the claim, because compression is a precondition for it. It does *not* attempt to measure the channel share of generative-AI assistants in discovery or conversion, which would require proprietary platform data; that is a separate, and currently under-evidenced, question. Our result is therefore best read as: to the extent AI-mediated discovery is operating, it has not (yet) produced a detectable compression of demand cycles in public search behaviour.

7. Limitations

Several limitations bound the claim and define the path to a confirmatory study. **(i) Proxy.** Google Trends measures search interest, not sales; attention and realised demand can diverge, and a search-to-sell-through validation is needed to promote these findings beyond a proxy. **(ii) Sampling.** The corpus is purposive, not a random draw from a defined frame; despite category matching, residual selection is possible. **(iii) Power.** Group sizes are modest; timing nulls denote no detected difference rather than proven equivalence, and two pre-AI cells were too thin to test. **(iv) Granularity.** Monthly data may understate sub-month spikes; weekly retrieval could sharpen emergence estimates. **(v) Construct.** Some AI-era series are brand- or franchise-level, which inflates measured durability; our brand-excluded robustness check mitigates but does not eliminate this. A confirmatory design would use a random, category-matched, equal-volume sample of specific products, sales-linked where possible, ideally retrieved through a programmatic pipeline.

8. Founder Perspective

I sponsored this study because the demand-shock narrative was shaping real decisions — inventory bets, forecasting investments, planning assumptions — faster than anyone seemed to be checking whether it was true. When a claim starts to move capital, it deserves evidence, not repetition. So we set out to test it the way I would want any consequential assumption tested in my own company: with a falsifiable hypothesis, public data, and a willingness to be proven wrong.

What I take from the result is not that artificial intelligence is unimportant to commerce. It is that one specific, confident, and widely repeated version of the story — that AI has compressed demand cycles — did not survive contact with the data. The trends we all remember as evidence of acceleration turn out, when measured against comparable trends from a decade ago, to be ordinary in their timing. The vivid examples were doing the persuading; the measurement was never done.

For enterprise leaders, I draw three practical conclusions. First, the durable planning skill is still the old one: distinguishing a fad from a lasting shift. That judgement predicted lifecycle shape in our data far better than knowing whether a trend was “AI-era.” Second, be cautious with narratives that arrive already confident and conveniently aligned with what vendors are selling; the compression story was assumed, not demonstrated. Third, the more interesting question is not whether AI made trends faster — our evidence says it did not — but whether AI is changing *where* demand forms and *how* it converts, which is a different claim requiring different, largely proprietary, data that the industry has not yet opened to scrutiny.

I want to be honest about what this work is and is not. It is an exploratory, transparent first pass built on a search-interest proxy and a modest, purposely-assembled sample. It is enough to challenge a strong claim; it is not the last word. We are publishing it as a preprint precisely so that others can pressure-test it. If the compression effect is real, a better-instrumented study should find it. We looked carefully, and we did not.

9. Conclusion

We tested whether artificial intelligence has compressed consumer demand cycles by comparing the lifecycles of pre-AI and AI-era product trends within matched categories. The data do not support the hypothesis: emergence and decay speeds are statistically indistinguishable across eras, while the only significant effects indicate that AI-era trends persist somewhat *longer*. Lifecycle compression is a property of fad products, present in both eras, not a signature of the AI era. For demand planners, the practical implication is that the discipline of distinguishing fads from durables remains more predictive than the era of a trend's emergence. For researchers, we offer a reproducible protocol — and a cautionary catalogue of the data-quality traps — for testing demand-shock claims rigorously rather than anecdotally.

Reproducibility and Data Availability

All series were retrieved from the public Google Trends interface (United States; web search; monthly; 2016-01 to 2026-05). The complete list of products, retrieval settings, and computed lifecycle metrics is provided in Appendix A, enabling exact replication. No data were simulated, imputed, or synthetically generated; every reported statistic is computed directly from the retrieved series. Analysis used standard open-source statistical libraries (Mann–Whitney, Fisher exact, bootstrap resampling).

Author Note on References

This preprint cites only foundational, long-established methodological and theoretical works. The reference list below should be verified and expanded by the authors with current, peer-reviewed literature on AI-mediated commerce, social-media-driven demand, and search-based nowcasting prior to formal submission; specific empirical claims about AI-assistant adoption or commerce-channel share must be supported by primary sources at that stage. No citation here should be taken as vetted for a particular venue's reference requirements without that verification step.

References

- [1] Bass, F. M. (1969). A new product growth for model consumer durables. *Management Science*, 15(5), 215–227.
- [2] Rogers, E. M. (2003). *Diffusion of Innovations* (5th ed.). Free Press.
- [3] Lee, H. L., Padmanabhan, V., & Whang, S. (1997). Information distortion in a supply chain: The bullwhip effect. *Management Science*, 43(4), 546–558.
- [4] Choi, H., & Varian, H. (2012). Predicting the present with Google Trends. *Economic Record*, 88(s1), 2–9.
- [5] Mann, H. B., & Whitney, D. R. (1947). On a test of whether one of two random variables is stochastically larger than the other. *Annals of Mathematical Statistics*, 18(1), 50–60.
- [6] Cliff, N. (1993). Dominance statistics: Ordinal analyses to answer ordinal questions. *Psychological Bulletin*, 114(3), 494–509.
- [7] Efron, B. (1979). Bootstrap methods: Another look at the jackknife. *Annals of Statistics*, 7(1), 1–26.

About XSYDA Research

XSYDA Research is the independent research and thought-leadership initiative of XSYDA Technologies. Its mission is to investigate emerging enterprise and computing-intelligence phenomena through evidence-based analysis, practical experimentation, and real-world validation. XSYDA Research publishes findings as preprints regardless of whether they confirm or contradict prevailing expectations, in the belief that disciplined, transparent inquiry serves enterprise decision-makers better than consensus repeated with confidence.

Author. Srinivasan Subramanian, Founder, XSYDA Technologies.

Published by. XSYDA Research, XSYDA Technologies.

Contact. contact@xsyda.com

Appendix A. Complete Product Corpus and Lifecycle Metrics

All 57 usable series (United States, web search, monthly, 2016–2026). “Onset→peak” and “decay half-life” in months; “censored” denotes no decay to 50% of peak within the window (durable). † marks brand/franchise-level terms flagged in robustness checks. Two pre-AI cells (food, home) retain a single usable specific product and are excluded from within-category tests.

Product	Category	Era	Peak	Onset→peak	Decay half-life
Adidas Gazelle	footwear	new	2024-08	2	2
Adidas Handball Spezial	footwear	new	2024-11	2	4
Adidas Samba	footwear	new	2026-04	44	censored
HOKA †	footwear	new	2024-07	19	11
New Balance 550	footwear	new	2026-05	5	censored
New Balance 9060	footwear	new	2024-12	24	censored
On Cloud †	footwear	new	2024-12	40	censored
Onitsuka Tiger Mexico 66	footwear	new	2026-01	26	censored
Salomon XT-6	footwear	new	2025-09	14	2
Adidas Ultraboost	footwear	old	2016-12	—	17
Allbirds	footwear	old	2026-04	56	1
Nike Air Max 270	footwear	old	2019-12	19	9
Nike Air VaporMax	footwear	old	2018-12	7	22
adidas NMD	footwear	old	2017-02	—	14
adidas Originals YEEZY Boost 3	footwear	old	2016-02	—	1
Dieux †	beauty	new	2025-11	14	censored
Drunk Elephant †	beauty	new	2023-12	55	3
Glow Recipe †	beauty	new	2023-11	7	5
Rare Beauty †	beauty	new	2023-03	15	2
Sol de Janeiro †	beauty	new	2025-12	1	1
Anastasia Beverly Hills †	beauty	old	2016-11	0	6
ColourPop †	beauty	old	2024-01	92	1
Fenty Beauty †	beauty	old	2024-12	20	censored
Glossier †	beauty	old	2022-11	74	4
Huda Beauty †	beauty	old	2025-07	79	2
Kylie Cosmetics †	beauty	old	2016-08	3	1
Chamoy pickle	food	new	2022-08	-4	1
Dubai chocolate	food	new	2025-03	3	7
Owala	food	new	2025-08	12	2
Poppi †	food	new	2026-01	3	censored
Stanley Quencher	food	new	2024-08	5	2
Unicorn Frappuccino	food	old	2017-04	0	1
Caraway †	home	new	2025-11	51	3
Carote †	home	new	2025-12	29	censored
Ember Mug	home	new	2024-12	1	1
Shark FlexStyle	home	new	2024-12	27	1
Instant Pot †	home	old	2020-01	25	5
Bluey (toys) †	toys	new	2024-12	54	1

Product	Category	Era	Peak	Onset→peak	Decay half-life
Labubu	toys	new	2025-06	5	4
Magic Mixies	toys	new	2023-12	6	2
Pop Mart †	toys	new	2026-03	107	1
Smiski	toys	new	2025-12	31	censored
Sonny Angel	toys	new	2025-03	24	7
Squishmallows	toys	new	2022-12	10	5
Beyblade	toys	old	2019-12	15	6
Fidget spinner	toys	old	2017-06	1	1
Fingerlings	toys	old	2018-12	9	1
Hatchimals	toys	old	2016-12	1	1
L.O.L. Surprise	toys	old	2024-12	5	2
Rainbow Loom	toys	old	2016-07	7	4
Shopkins	toys	old	2015-12	0	4
Bounty Paper Towels †	staples	control	2020-04	1	1
Charmin †	staples	control	2020-03	0	1
Clorox †	staples	control	2020-04	46	2
Colgate †	staples	control	2026-04	—	censored
Dawn †	staples	control	2025-07	24	4
Tide †	staples	control	2018-08	—	3

Table A1. Full corpus with computed lifecycle metrics (n=57).