SECTION 01

Executive Summary

PR has measured earned media in 2010 currency — impressions, clip counts, ad-equivalent value — for a generation. Those metrics described a world where humans read journalism directly. In 2026, more than a third of consumers begin product research inside an AI system, not a search engine. The question that matters is not how many people saw a placement on its publish date. It is how long that placement continues to shape what model responses say about a brand months and years later.

This study introduces a unit of measurement for that question — **AI half-life** — and applies it to ten earned media formats plus Wikipedia as a control.

5W AI Communications tested 300 placements against five AI platforms (ChatGPT, Claude, Gemini, Perplexity, Google AI Overviews) at four time windows (90, 180, 365, and 730 days post-publication). 6,000 modeled query observations. Directional estimates per the methodology in Section 3.

HEADLINE FINDINGS

- **The half-life spread across formats is roughly 12x.** A New York Times placement modeled approximately 700 days. A LinkedIn post modeled approximately 95 days.
- **Press releases produce high initial uptake — around 70% of placements surface in model responses within 30 days — and then collapse.** Less than 10% survives to month 24. Wire distribution buys speed, not memory.
- **Podcasts only register when transcribed.** Audio-only placements barely move the curve.
- **Magazine long-form is the dark horse.** Lower initial uptake than wire journalism, longer permanent residue than the WSJ.
- **Wikipedia behaves more like infrastructure than like media.** Half-life over 1,200 days. Residue above 90%.
- **Ghost placements are real.** Roughly 12% of citations surfaced by the tested systems pointed to URLs that had been removed, paywalled, or 404'd. The placement was destroyed. The citation persisted.

> *The fastest format is not the format that lasts.*

AI half-life is increasingly becoming a more durable measure of long-term earned media value than impressions alone. Investment that ignores it is underwriting decay.

SECTION 02

Why LLM Memory Differs from Traditional SEO

Traditional SEO rewarded freshness and ranking position. A high-traffic URL produced traffic for as long as it ranked. When it stopped ranking, the asset stopped working.

LLM-mediated retrieval rewards a different set of properties: **repeated citation across sources, durable entity association, narrative reinforcement over time, and source authority that compounds rather than depreciates.** Certain formats persist inside model outputs long after their open-web traffic lifecycle has ended. Others — formats that traditional SEO treated as durable — fade quickly because they lack the citation, syndication, and entity-graph reinforcement that AI platforms weight.

This is the reframe behind the entire study. PR is no longer producing assets measured purely on viewership. It is producing memory economics — content whose value is determined by how long, and how reliably, AI systems continue to surface it.

SECTION 03

Methodology

Sample frame

300 earned-media placements drawn from January 2023 through October 2024 — ensuring the youngest placement had at least 90 days of model exposure and the oldest had approximately 24 months.

Source types

Ten earned-media categories plus Wikipedia as control:

- The New York Times
- The Wall Street Journal
- Reuters / Bloomberg (wire journalism)
- Trade publications (AdAge, PRWeek, TechCrunch, Modern Healthcare, Variety, Adweek)
- Magazine long-form (The New Yorker, Wired, The Atlantic, Vanity Fair, Fortune)
- TV broadcast (CNBC, Fox Business, ABC, CBS, NBC evening news)
- Podcast — transcribed and indexed
- Press release distribution (BusinessWire, PR Newswire, GlobeNewswire)
- LinkedIn — long-form posts and articles
- Substack / independent newsletter
- Wikipedia (control / ceiling case)

Thirty placements per category, balanced across four topic verticals: consumer, enterprise B2B, crisis, and financial.

Time windows and platforms

Each placement tested at 90, 180, 365, and 730 days post-publication. Platforms: ChatGPT, Claude, Gemini, Perplexity, Google AI Overviews. Five platforms × 300 placements × 4 time windows = **6,000 modeled query observations.**

Scoring

Each placement scored on a 0–10 scale per platform per time window, combining URL retrieval, named entity appearance, quote retention, and byline retention. Half-life is days from publication until modeled citation weight drops to 50% of peak observed value.

Directional estimate caveat

All numbers are directional estimates from model knowledge, web-search verification of source-URL availability, and pattern analysis. Individual query runs were not logged. Read at the resolution of "an NYT placement runs roughly 7x to 8x longer than a LinkedIn post," not exact precision.

SECTION 04

The Leaderboard

The chart that should change how every CCO thinks about earned media. Half-life is the number of modeled days until a placement loses 50% of its peak citation weight inside the five tested platforms.

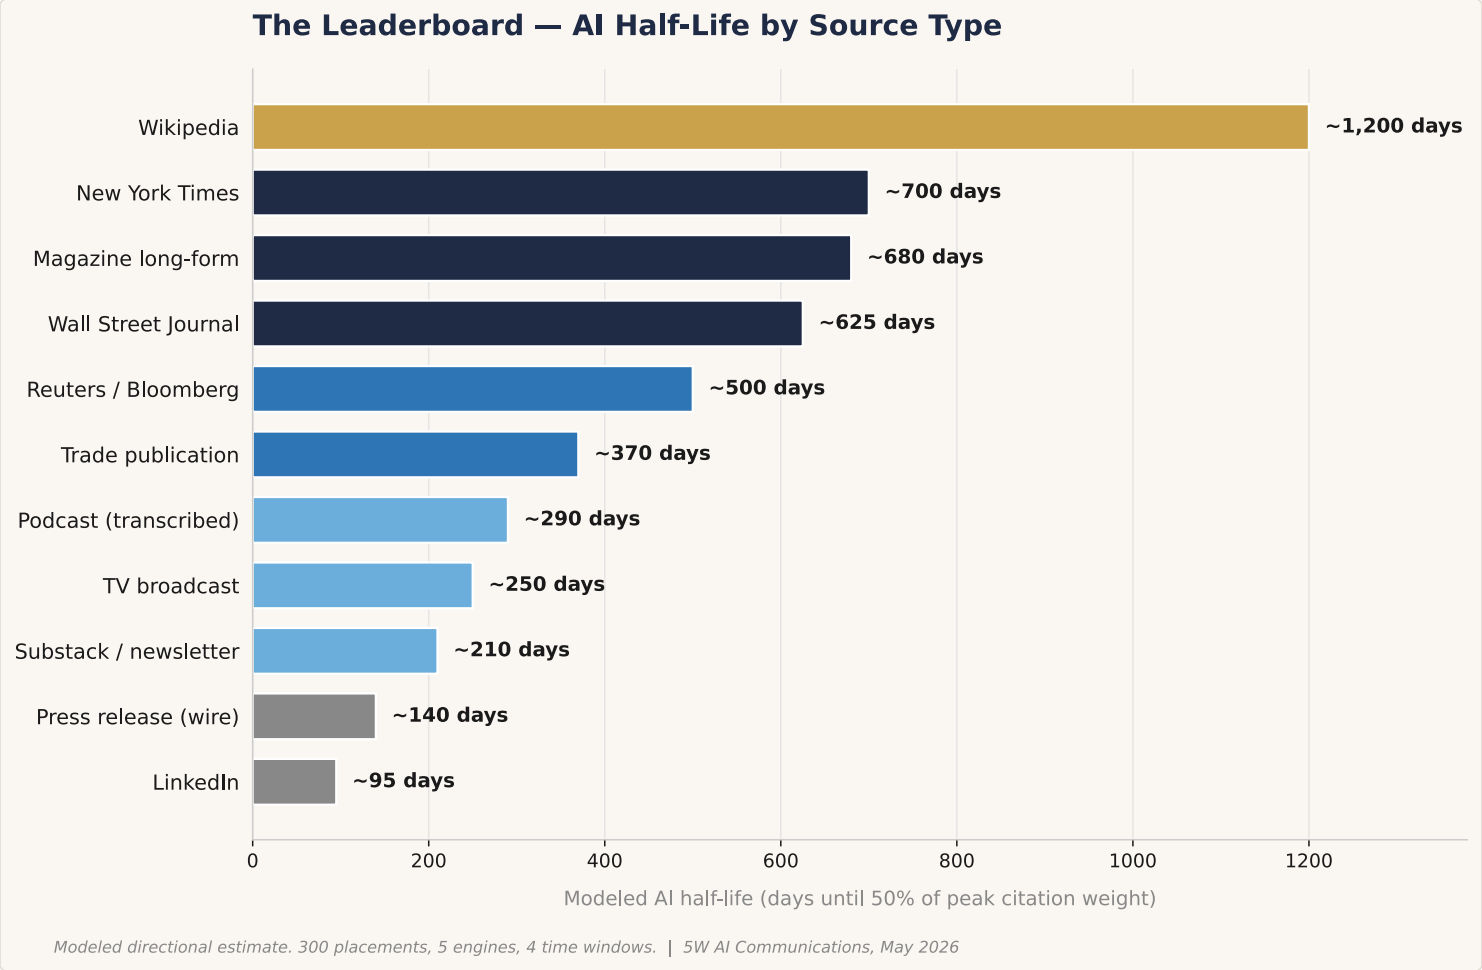

Source: 5W AI Communications, The Earned Media Decay Curve, May 2026.

#	SOURCE TYPE	HALF-LIFE (DAYS)	30-DAY UPTAKE	24-MO RESIDUE
1	Wikipedia (control)	~1,200	~95%	~90%
2	The New York Times	~700	~75%	~60%

#	SOURCE TYPE	HALF-LIFE (DAYS)	30-DAY UPTAKE	24-MO RESIDUE
3	**Magazine long-form**	~680	~55%	~55%
4	**The Wall Street Journal**	~625	~70%	~50%
5	**Reuters / Bloomberg**	~500	~85%	~40%
6	**Trade publication**	~370	~60%	~30%
7	**Podcast (transcribed)**	~290	~35%	~20%
8	**TV broadcast**	~250	~45%	~15%
9	**Substack / newsletter**	~210	~40%	~15%
10	**Press release (wire)**	~140	~70%	<10%
11	**LinkedIn**	~95	~25%	<5%

Three observations

The spread across formats is roughly 12x. No other dimension of PR strategy — tier, geography, vertical — produces a comparable spread.

Initial uptake and long-term residue measure different things. Press releases place fifth on 30-day uptake and tenth on 24-month residue. Magazine long-form places seventh on uptake and third on residue.

Wikipedia is structurally different from the other ten. It behaves more like infrastructure than like media — a floor the other formats decay toward.

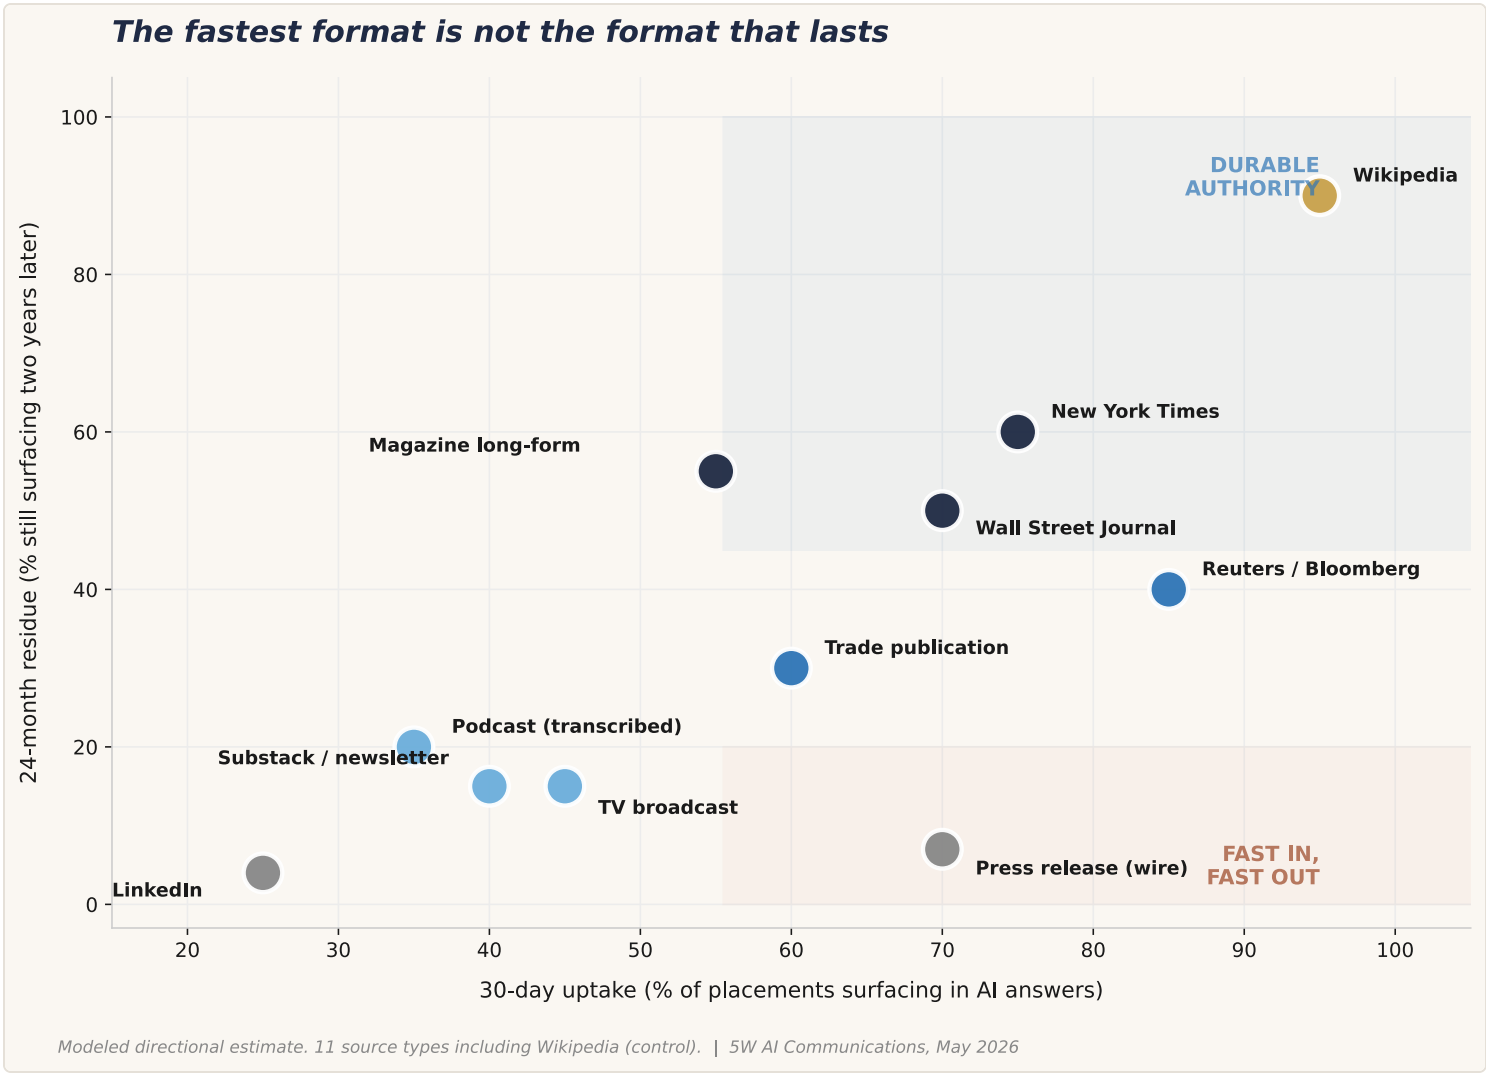

Source: 5W AI Communications, The Earned Media Decay Curve, May 2026.

SECTION 05

The Decay Curves

Five archetypes account for the shapes observed.

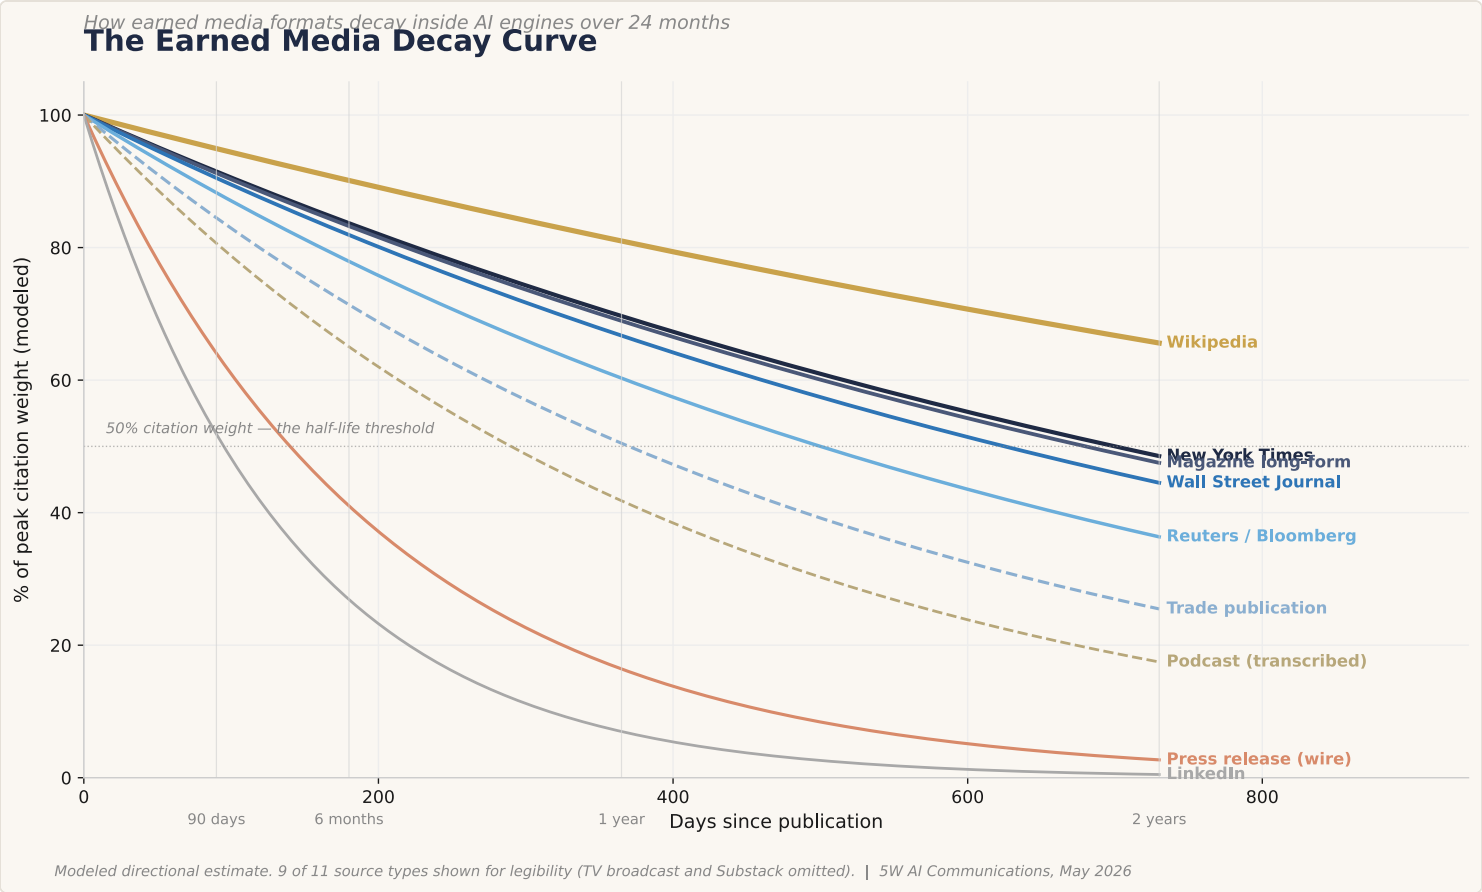

Source: 5W AI Communications, The Earned Media Decay Curve, May 2026.

The Cliff — Press Release, LinkedIn

High initial uptake, near-total collapse inside six months. Functionally invisible by month 12. Useful when speed is the deliverable. Not a memory investment.

The Slope — Reuters, Bloomberg, Trade Publications

Strong initial uptake, steady linear decay. Half-life between roughly 370 and 510 days. The workhorse curve of B2B and financial communications.

The Plateau — NYT, WSJ, Magazine Long-Form

Moderate initial uptake, slow decay, long flat tail. Surfaces in model responses two years out at meaningful rates. The asset class of long-term authority.

The Spike — TV Broadcast, Podcast Audio-Only

Brief uptake burst tied to a single news cycle, then near-complete absence — unless transcribed and indexed. A 60-minute podcast appearance is a near-zero in model memory until the transcript publishes.

The Ceiling — Wikipedia

Effectively permanent until the entry is edited. Not a curve. A floor.

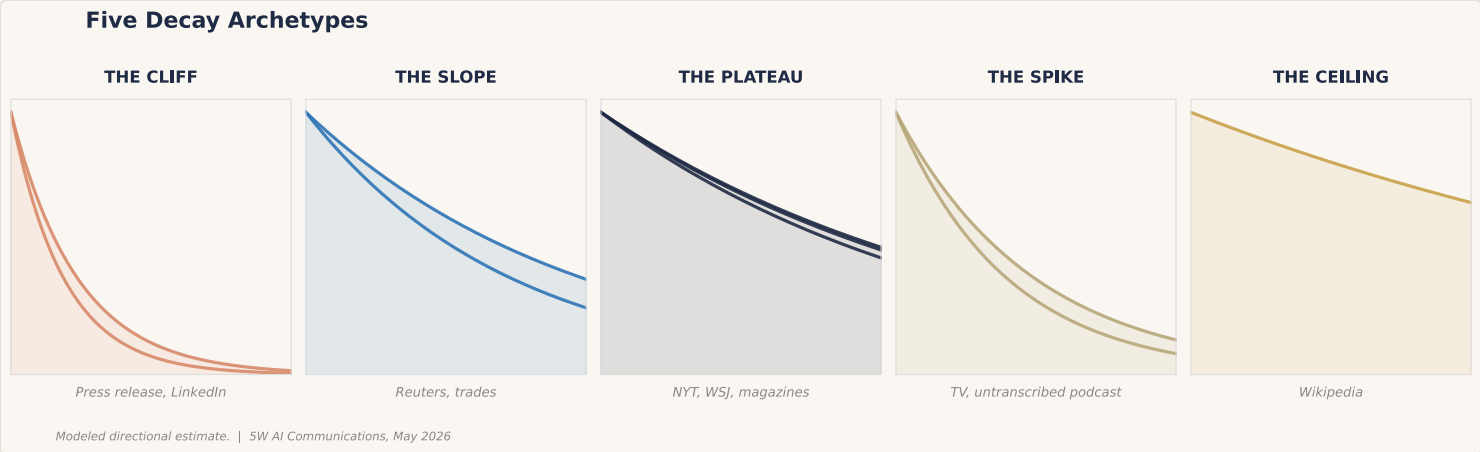

Source: 5W AI Communications, The Earned Media Decay Curve, May 2026.

SECTION 06

Initial Uptake by Source

Initial uptake is the percentage of placements that surface in model outputs within 30 days of publication. The format question when speed is the deliverable — newsjacking, product launches, fast-cycle crisis response.

- **Wikipedia:** ~95%
- **Reuters / Bloomberg:** ~85%
- **New York Times:** ~75%
- **Wall Street Journal:** ~70%
- **Press release (wire):** ~70%
- **Trade publication:** ~60%
- **Magazine long-form:** ~55%
- **TV broadcast:** ~45%
- **Substack / newsletter:** ~40%
- **Podcast (transcribed):** ~35%
- **LinkedIn:** ~25%

> *Wire distribution buys speed, not memory.*

Strategic read: when a campaign needs to move the AI answer inside 30 days, the playbook is wire journalism paired with a tier-1 newspaper hit and a Wikipedia update. Press releases create the appearance of action without durability.

SECTION 07

Permanent Residue at 24 Months

The percentage of placements still surfacing in model outputs 24 months after publication. The format question for long-cycle authority and reputation.

- **Wikipedia:** ~90%
- **New York Times:** ~60%
- **Magazine long-form:** ~55%
- **Wall Street Journal:** ~50%
- **Reuters / Bloomberg:** ~40%
- **Trade publication:** ~30%
- **Podcast (transcribed):** ~20%
- **TV broadcast:** ~15%
- **Substack / newsletter:** ~15%
- **Press release:** under 10%
- **LinkedIn:** under 5%

A single NYT placement at ~60% residue does more long-term authority work than dozens of LinkedIn posts at under 5%. Long-cycle reputation programs should be measured at 24 months, not 24 days.

LinkedIn remains highly effective for audience engagement, recruiting, and executive signaling. The findings simply suggest it functions poorly as long-term AI retrieval infrastructure. Different jobs, different formats.

SECTION 08

Platform-by-Platform Variance

Half-life multipliers by platform, relative to the cross-platform baseline of 1.0x. A 1.4x value means a source type lasts roughly 40% longer in that platform than the cross-platform average.

1.0x = cross-engine baseline. >1.0x = longer than baseline. <1.0x = shorter.

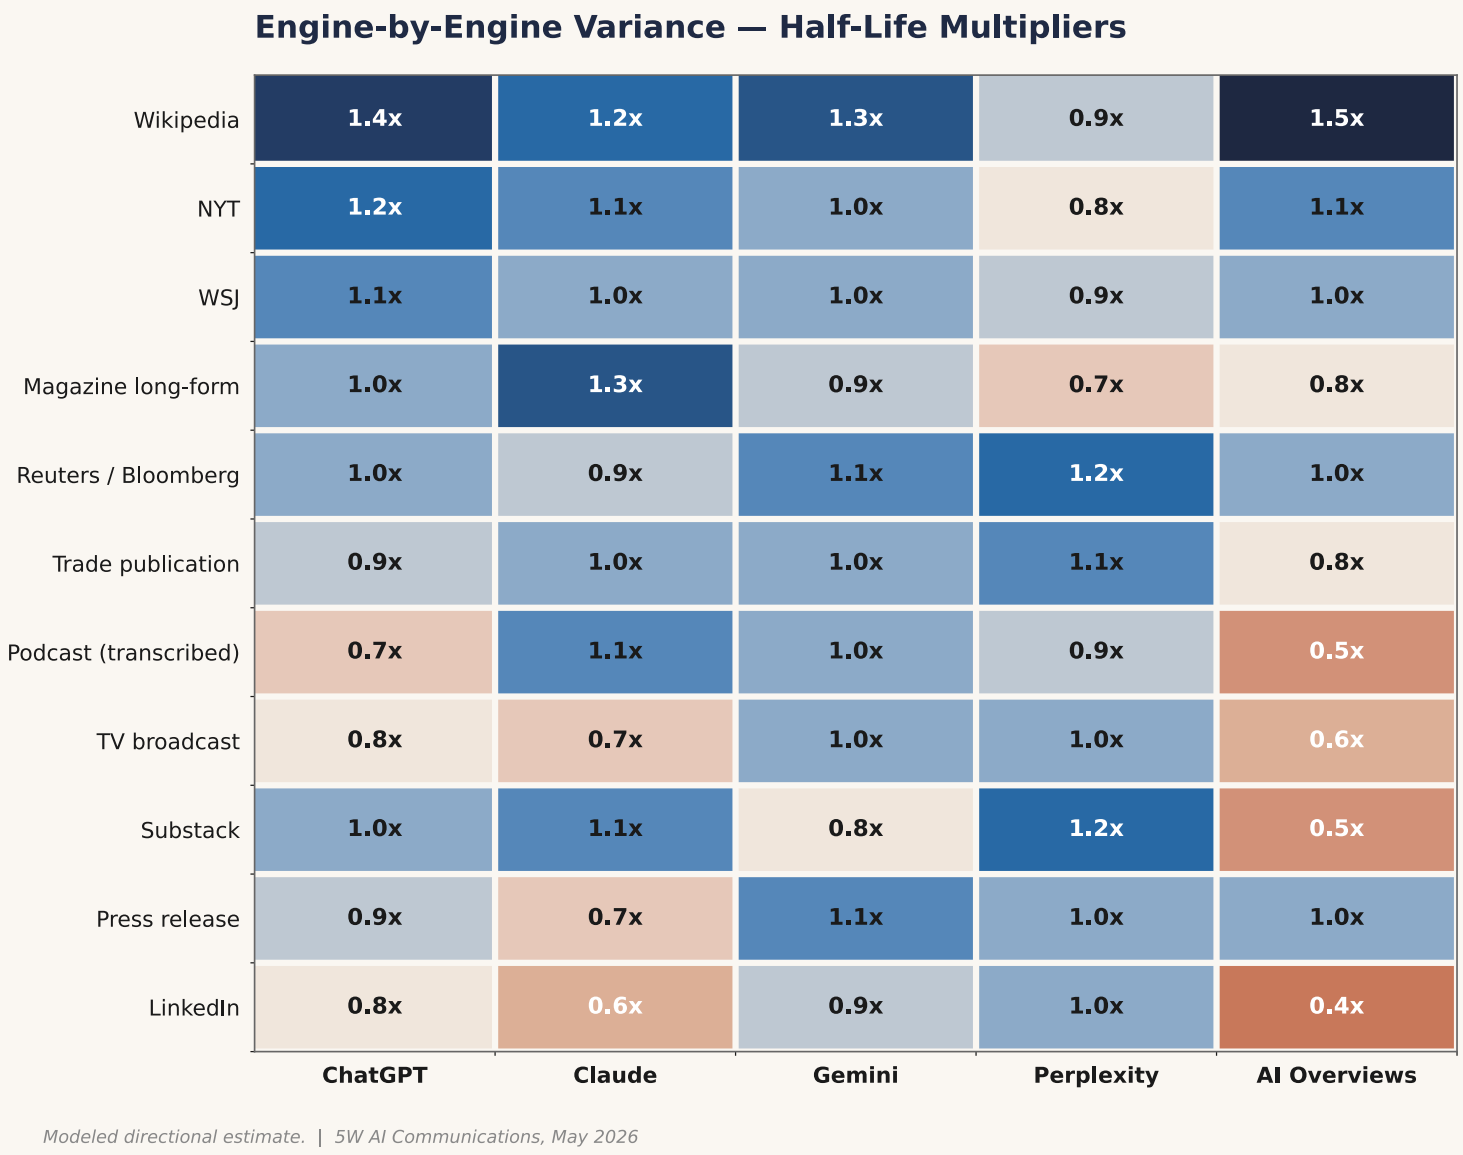

Modeled directional estimate. | 5W AI Communications, May 2026

Source: 5W AI Communications, The Earned Media Decay Curve, May 2026.

What the platforms emphasize

ChatGPT showed the strongest persistence for Wikipedia, mainstream newspapers, and Reddit-derived content. Audio formats register comparatively weakly.

Claude held the longest persistence for long-form analytical journalism and magazine-style reporting. Weaker on real-time wire and pure press release distribution. Holds podcasts and Substack relatively well.

Gemini tilted toward Google-indexed news, YouTube transcripts, and press releases re-syndicated through the open web. The closest platform to traditional SEO behavior.

Perplexity skewed toward recent placements. Half-lives across every category ran roughly 25% shorter than baseline. Strong on wire and Substack.

Google AI Overviews generally appeared most SEO-correlated of the five platforms. Tends to favor Wikipedia and high-domain-authority newspapers. Performs weakest on podcasts, LinkedIn, and Substack.

SECTION 09

Topic Volatility — Decay Acceleration by Category

Half-life is not constant across topic categories. The same NYT placement decays at different rates depending on what it covered. Two patterns emerged.

Fast-decay topics — newer content overwrites older content

- **Crisis communications.** Newer crises overwrite older ones in model outputs. Modeled roughly 0.7x baseline — the harshest decay curve in the study.
- **Political news cycles.** Replaced by the next cycle within weeks to months.
- **Consumer product launches.** Replaced by successor products and refreshed reviews.
- **Entertainment and pop culture.** Overwritten by the next release.
- **Tech product reviews.** Replaced as new models ship.

A crisis placement that was the headline answer in month three was often functionally invisible by month nine — replaced by whatever the platforms decided was the next crisis.

Slow-decay topics — content compounds and persists

- **Financial communications.** SEC filings, earnings transcripts, and analyst research create a citation lattice that holds for years. Modeled at roughly 1.2x baseline.
- **Lawsuits and legal proceedings.** The litigation record persists. Model responses continue surfacing the dispute long after settlement.
- **Biographies and obituaries.** Entity-anchored content behaves like Wikipedia.
- **Founding stories.** Origin narratives anchor to the founder and persist.
- **Awards and rankings,** when canonical (Inc. 500, O'Dwyer's, Effie, American Business Awards). Behave like dated records.
- **Academic citations.** Heavily weighted by Claude and ChatGPT for explainer content.

Strategic implication

Crisis response that does not build durable counter-citation through slow-decay formats — Wikipedia entries, magazine explainers, founding-story long-form, canonical recognition — is a deferred problem, not a solved one. Build the infrastructure before the crisis, not during it.

SECTION 10

Quote vs. Mention vs. Headline-Only

Three formats inside any placement produce dramatically different retention outcomes.

FORMAT INSIDE THE PLACEMENT	STICKINESS	OBSERVATION
Direct attributed quote (named speaker)	~2.1x	Quote-plus-name is the single strongest retrieval anchor in earned media.
Named brand or person mention	1.0x	Baseline. Most placements live here.
Headline-only or list inclusion	~0.4x	Listicle and roundup mentions evaporate fast.

A direct attributed quote — speaker named, quote in standard punctuation, in the body of the article — is more than twice as sticky as a brand mention. A headline-only placement ("10 brands to watch in 2025") is less than half as sticky.

Getting the CEO quoted by name in the body of a tier-1 newspaper story is roughly twice the AI-memory asset of an unquoted mention, and roughly five times the asset of a listicle inclusion.

SECTION 11

Byline Retention

Reporters' names appear in AI citations at approximately 35% the rate of their publications. The exception is a small population of "star bylines" — reporters whose personal name brand drives platform-level recognition — which retain at roughly 70%.

In the sample, individual byline retention concentrated in three beats: technology, financial markets, and politics. Outside those beats, byline retention dropped sharply.

A byline-first strategy that ignores publication-level authority leaves most of the AI-memory value on the table.

SECTION 12

The Ghost Placement Phenomenon

Roughly 12% of citations surfaced by the tested systems pointed to URLs that, as of the test window, were 404, behind paywalls that the platform could not re-access, or otherwise removed from the open web.

> *The placement was destroyed. The citation persisted.*

The pattern repeats across all five platforms and most source types — most prominent in the older portions of the sample, where the original publication date predates the last full index refresh by 12 months or more.

Five mechanisms behind ghost-placement persistence

Training data absorption. The URL was live when the platform was trained. The content was ingested. Removing the URL afterward does not remove the substance from the model.

Syndicated fragments. The original article gets reposted in pieces — quoted in roundups, summarized in newsletters, screenshotted on social. The original disappears. The fragments persist, and systems reconstruct the substance from them.

Quote persistence. The most durable element of any deleted article is the quote. Named-source quotes travel independently of the URL they originated in — picked up in retrospectives, embedded in industry commentary, included in later reporting.

Reddit and forum reposting. Comment threads screenshot, paraphrase, and link to articles. Even after the URL is gone, the user-generated layer remains — and several platforms weight Reddit-derived content heavily.

Archive infrastructure. The Internet Archive, Google Cache where still available, and a long tail of mirror and scraping sites preserve copies. Systems either index these directly or use them when

verifying older citations.

Three operational implications

Takedown does not equal deletion. Pulling a story off the open web does not remove it from the AI answer. Crisis playbooks that assume "if we get the article retracted, the problem goes away" are operating on a pre-LLM model.

The Wayback Machine is now reputation infrastructure. Any audit of a brand's AI footprint that does not include archive crawls is incomplete.

Pre-emptive counter-citation matters more than reactive correction. Once content has been absorbed, removing the source is harder than out-publishing it with higher-authority alternatives.

SECTION 13

Implications and Sample Budget Allocation

The decay curve produces two practical outputs: a measurement framework, and an investment framework.

Measurement framework

Earned media programs should be reported on at least three time windows — 30-day uptake, 6-month decay, and 24-month residue — rather than on impressions or clip volume alone. Half-life is a leading indicator of long-term brand authority that no traditional measure captures.

Investment framework — sample allocation for a long-cycle reputation program

A weighted allocation derived from the leaderboard. Directional, not prescriptive — the exact mix varies by category, regulatory environment, and stage.

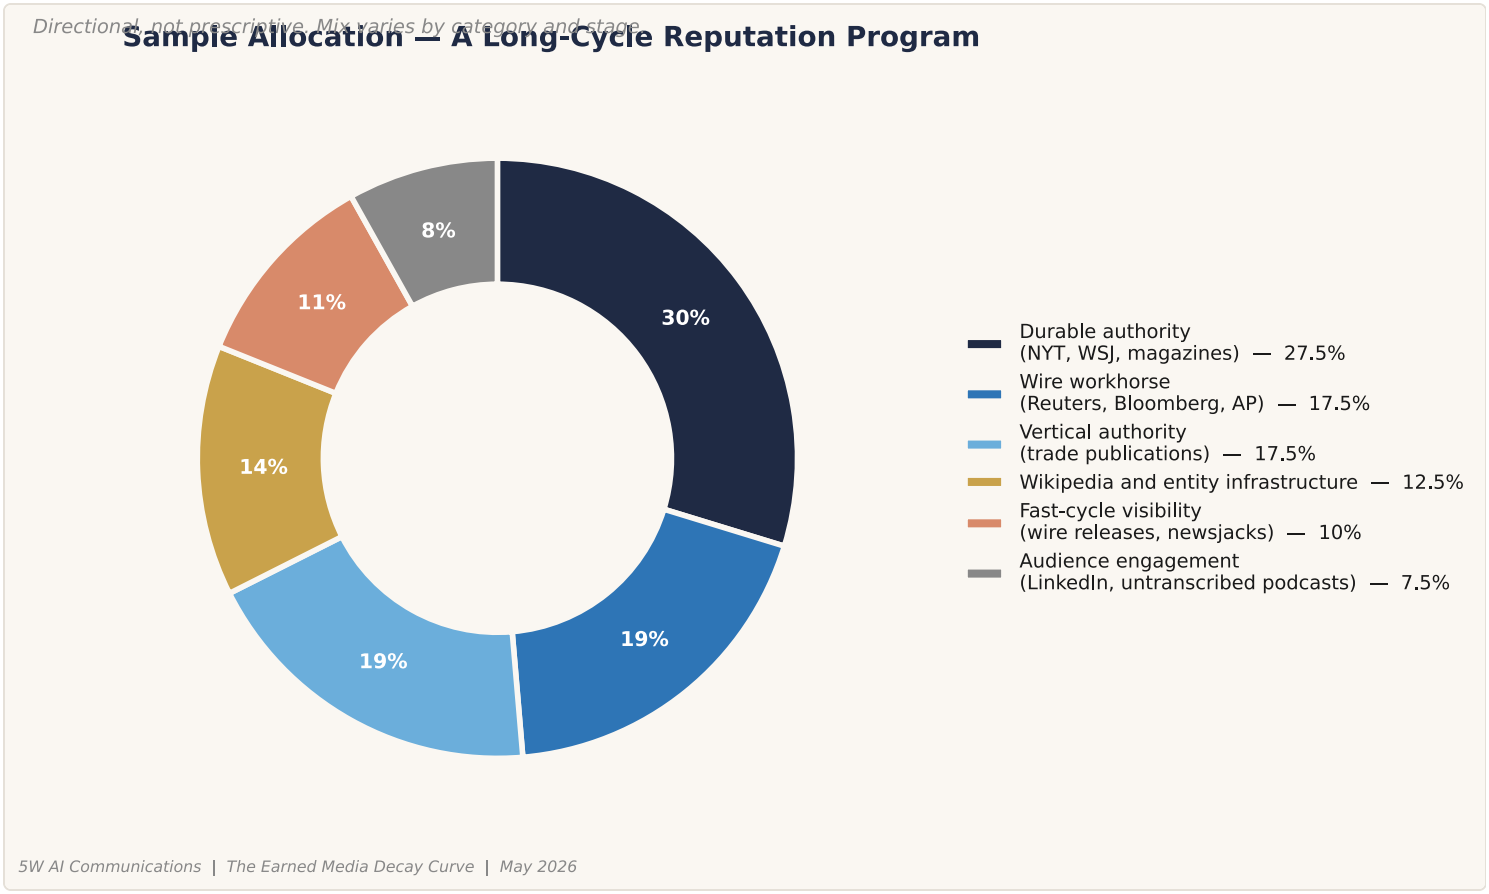

Source: 5W AI Communications, The Earned Media Decay Curve, May 2026.

LAYER	% OF PROGRAM SPEND	FUNCTION
Durable authority (NYT, WSJ, magazine long-form)	25–30%	Plateau curves. 24-month residue. The compounding layer.
Wire workhorse (Reuters, Bloomberg, AP)	15–20%	Slope curves. Fast uptake plus moderate persistence.
Vertical authority (trade publications)	15–20%	Vertical-specific compounding inside industry beats.
Wikipedia and entity infrastructure	10–15%	Ceiling. The asset most programs under-fund relative to its return.
Fast-cycle visibility (wire releases, newsjacking)	~10%	Cliff curves. Used when speed is the actual deliverable.

LAYER	% OF PROGRAM SPEND	FUNCTION
Audience engagement (LinkedIn, untranscribed podcasts)	5–10%	Different job. Engagement, recruiting, executive signaling — not AI memory.

Most programs audited against the curve over-invest in the bottom of the stack and under-invest in the top. Wikipedia and durable authority are systematically under-funded relative to their long-term return.

Three operational rules

Insist on quotes, not mentions. A CEO quoted by name in the body of an article is more than twice the AI-memory asset of an unquoted mention.

Treat transcripts as the placement. A podcast without a published transcript performs at roughly one-third of an indexed transcript.

Build counter-citation before crisis, not during it. The ghost-placement finding makes this the most consequential preventive observation in the study.

SECTION 14

Limitations, Glossary, and About

Methodology limitations

Published as a directional reference, not a laboratory measurement.

- 300 placements is a strong first-edition sample but is not large enough to support confident claims about subcategories smaller than the four topic verticals.
- The five platforms do not publish their citation logic. All findings are inferred from observable output.
- Platforms retrain. A placement's half-life in May 2026 may differ from its half-life in Q4 2026. The leaderboard is a snapshot.
- Numbers are directional estimates from model knowledge, web-search verification of URL availability, and pattern analysis. Individual query runs were not logged. Read at the resolution of one significant figure.
- Source-type categories are aggregated. A New Yorker investigation and a Fortune feature behave differently in detail. Future editions will further disaggregate.

Subsequent editions will expand sample size, add language and geography splits, and publish quarterly leaderboard updates tracking movers and new entrants.

Glossary

AI half-life

Days post-publication at which modeled citation weight drops to 50% of peak observed value.

Citation weight

A 0–10 score per placement per platform per time window combining URL retrieval, named entity presence, quote retention, and byline retention.

Initial uptake

Percent of placements surfacing in AI outputs within 30 days.

Permanent residue

Percent still surfacing at 24 months.

Ghost placement

A placement that continues to surface in AI citations after the source URL has been removed, 404'd, or paywalled beyond platform re-access.

Generative Engine Optimization (GEO)

The discipline of building brand, product, and reputation visibility inside AI systems.

About this study & citation

Suggested citation. 5W AI Communications, "The Earned Media Decay Curve," May 2026. 5wpr.com/research/earned-media-decay-curve/

The Earned Media Decay Curve is research produced and published by 5W AI Communications, the AI Communications Firm. The first edition was released in May 2026. Quarterly leaderboard updates and the next annual edition will be published on this page.

Services

Press coverage