

*digitalmarketing.fyi*

RESEARCH WHITEPAPER | CONTENT MARKETING

# The Content Half-Life Paradox

---

A Discount-Rate Framework for Measuring  
Compounding Returns in B2B Content Marketing

Published by digitalmarketing.fyi Research

April 2026

Paper ID: DMFYI-WP-2026-001

Version: 1.0

# Copyright & Licensing

© 2026 digitalmarketing.fyi. Some rights reserved.

This whitepaper is published under a Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0). See <https://creativecommons.org/licenses/by-nc/4.0/>.

Readers are free to share and adapt this work, with attribution, for non-commercial purposes. Commercial reuse requires written permission from digitalmarketing.fyi.

- - -

## Suggested Citation

*digitalmarketing.fyi*. "The Content Half-Life Paradox: A Discount-Rate Framework for Measuring Compounding Returns in B2B Content Marketing." *digitalmarketing.fyi Research*, DMFYI-WP-2026-001, April 2026. <https://digitalmarketing.fyi/whitepapers/content-half-life-paradox>.

- - -

## Version History

Version	Date	Author	Changes
1.0	2026-04-18	digitalmarketing.fyi Research	Initial release

- - -

## Disclaimer

This whitepaper is provided for informational purposes. It does not constitute business, legal, or financial advice. Figures and benchmarks are drawn from the sources cited in the References section; readers are encouraged to verify claims against primary sources before acting on them. Where proprietary data is used, methodology is disclosed in the Methodology section.

# Abstract

B2B content marketing generates measurable returns over twelve-to-twenty-four-month horizons, yet the industry evaluates it on thirty-to-ninety-day windows. This paper argues that the resulting valuation error is the dominant reason B2B content programs are defunded before they reach their earning potential. Drawing on public research from HubSpot, the Content Marketing Institute, Ahrefs, Semrush, Gartner, Forrester, and the Ehrenberg-Bass Institute, the analysis shows that content assets have a continuous half-life distribution rather than a binary compounding-or-decaying status; that measurement windows shorter than the sales cycle systematically understate ROI; that AI Overviews compress the traffic value of each half-life increment by approximately thirty-to-fifty-eight percent on informational queries; and that B2B content's primary job is not to convert the five percent of buyers in-market today but to build mental availability across the ninety-five percent who will enter the market later. The paper introduces the Content Half-Life Model, a net-present-value framework that adapts capital-budgeting methodology to content portfolios. The framework quantifies publish, refresh, and retire decisions and presents content performance in the language of asset valuation rather than campaign ROI. The intended reader is a B2B marketing leader defending a multi-year content investment to a finance-minded executive audience.

- - -

**Audience:** Heads of Growth, VPs of Marketing, and Marketing Directors at B2B organizations

**Primary keyword:** B2B content marketing strategy

**Secondary keyword:** content marketing ROI

**Supporting keywords:** content decay, content marketing measurement, mental availability, 95:5 rule, AI Overviews SEO

**Word count:** 5117

**Paper ID:** DMFYI-WP-2026-001

# Executive Summary

Most B2B content marketing programs are evaluated on timeframes shorter than the asset itself earns over. The result is a systematic undervaluation that causes content to be defunded precisely when it is closest to compounding. This paper reframes B2B content as a capital asset with a half-life, adapts finance-style net present value methodology to content portfolios, and introduces the Content Half-Life Model as a decision framework for marketing leaders.

Four findings anchor the analysis.

First, the binary framing of content as either compounding or decaying is too coarse. Every content asset has a half-life, and the distribution across content types is continuous - from news content with half-lives measured in days, to pillar and research content with half-lives measured in multiple years. Portfolio performance is a weighted average of many assets at different points on their half-life curves, which short-window dashboards cannot see.

Second, measurement windows that are shorter than the underlying B2B sales cycle - currently around ten months on average - systematically understate content ROI. Marketing Week's survey of six hundred B2B marketers found that approximately half measure campaigns over periods shorter than their organization's own average sales cycle. Only thirty-six percent of marketers across industry aggregates report confident ability to measure content ROI, and just forty-one percent can confidently attribute revenue to specific content pieces. In this environment, window selection alone determines whether a program appears successful.

Third, content faces a two-front war in 2026. Classical content decay driven by competitive publication and intent drift remains active, and AI Overviews now compress the traffic value of each half-life increment by between thirty and fifty-eight percent on informational queries depending on the measurement source. Gartner projects that twenty-five percent of traditional organic search traffic will shift to AI chatbots and voice assistants by the end of 2026. Long-half-life research and pillar content is disproportionately likely to be cited by AI answer engines, which reinforces the case for portfolio weighting toward compounding assets.

Fourth, and most consequentially, B2B buying is a ninety-five-to-five game. The Ehrenberg-Bass Institute's 95:5 rule documents that approximately ninety-five percent of a category's potential buyers are out of market at any given time. Content's primary job is therefore not to convert today's five percent but to build mental availability across the ninety-five percent who will enter the market over the next one to five years. Forty-one percent of B2B buyers already have a preferred vendor before formal evaluation begins, which means the content that shaped that shortlist worked long before any attributable lead event.

The Content Half-Life Model in Section 6 formalizes these findings into a practical valuation approach. It applies a discount-rate calculation to the future cash flows of each content asset; categorizes content into four half-life bands with empirical monthly decay rates; and yields a publish, refresh, or retire decision for every piece in a portfolio. Applied consistently, the model reframes content reporting from month-on-month traffic charts into asset-level net present value - a language finance leaders already speak.

The practical recommendations are straightforward: extend the default measurement window to twelve-to-twenty-four months, budget refresh as capital preservation rather than cost, weight portfolios deliberately across half-life categories, present content performance to the CFO in asset-valuation terms, and add a citation-optimized layer designed for AI answer engines rather than only for click capture. Teams that adopt these practices defend their content budgets with evidence instead of advocacy.

# 1. Introduction - the measurement problem in B2B content marketing

Most B2B content marketing programs are evaluated on timeframes that systematically understate their value. A blog post published in March may be judged by its April numbers. A content investment made in Q1 is reviewed against Q2 pipeline. A twelve-month content budget is defended against a thirty-day dashboard. The consequence is predictable: content that is quietly compounding into a long-term asset looks, on a short-term view, like an underperforming cost.

This paper argues that the industry's dominant measurement windows - thirty days, ninety days, and even single fiscal quarters - are shorter than the underlying B2B buying cycle and shorter than the time content takes to reach its earning potential. That gap creates a systematic valuation error. Content marketing is not failing in most programs that declare it failing; it is being measured through the wrong lens.

The core thesis is that content is a capital asset, not an expense, and it should be valued accordingly. A finance team would never evaluate a software platform on its first-month return. They would apply a discount rate, model the cash flows over the asset's expected useful life, and calculate a net present value. Content deserves the same rigor. This paper proposes that rigor.

Three things follow from this framing. First, short measurement windows are not conservative - they are inaccurate. Second, the dominant blog post lifecycle in B2B is not "spike and fade" but a bimodal distribution between rapid-decay content and compounding content, with a measurable half-life in between. Third, mapping that half-life against the buying cycle exposes which content is mispriced in a team's current portfolio and where reinvestment would produce outsized returns.

Scope: this paper focuses on B2B, owned content (primarily blog, whitepapers, research, and pillar pages), and organic search as the primary distribution channel. It does not cover paid social content, influencer content, B2C consumer editorial, or video platforms where discovery mechanics differ meaningfully. The analytical framework may generalize beyond these boundaries, but the

empirical evidence cited has been drawn from B2B-SaaS and B2B-services research.

## **2. Background - the B2B content paradox**

Two realities coexist in the B2B content marketing literature, and each contradicts the other.

The first reality is that content takes a long time to work. Industry practitioners and published benchmarks converge on a window of six to eighteen months for a new B2B content program to generate meaningful returns, with break-even typically reached between months nine and fifteen (Averi Resources 2026; Konabayev 2026). Search engine optimization, the dominant channel for owned content, has a documented ~nine-month break-even period for B2B programs and requires a cumulative investment that does not pay back inside a standard fiscal quarter (Martal 2026). The B2B sales cycle itself is long - averaging 10.1 months according to 6Sense's 2025 data, and stretching past eighteen months for complex enterprise deals.

The second reality is that most B2B marketers measure content on windows far shorter than the sales cycle it is meant to influence. Ehrenberg-Bass Institute research summarized by the LinkedIn B2B Institute found that roughly ninety-five percent of B2B marketers expect to see significant sales within the first two weeks of a campaign (Ehrenberg-Bass Institute 2025). Marketing Week, reporting on a survey of six hundred B2B marketers, found that approximately half measure campaign results over a period shorter than the organization's own average sales cycle (Marketing Week 2022). Only twenty-four percent of the surveyed B2B marketers ran campaigns for longer than six months.

The gap between these two realities is the paradox. The asset has a six-to-eighteen-month payback; the evaluation window has a thirty-to-ninety-day horizon. The evaluator is, in effect, being asked whether a mortgage is a good investment after the first month's payment.

The paradox is compounded by three structural features of B2B buying that make short-horizon evaluation even less reliable than it might first appear.

First, most B2B buyers are out of the market at any given time. Professor John Dawes of the Ehrenberg-Bass Institute has formalized this as the 95:5 rule: at any given point, approximately five percent of a category's potential buyers are actively in the market to purchase, and roughly ninety-five percent are not (Dawes 2021). The implication is that content reaching today's buyers is, in the majority, building memory structures that will influence a purchase twelve to thirty-six months from now. A thirty-day view sees five percent of the effect.

Second, B2B buying is a group decision. Gartner's research places the typical buying committee for complex B2B solutions at six to ten members, each entering the evaluation with four to five independently gathered pieces of research they later share with the group (Gartner 2024). Forrester's 2024 State of Business Buying puts the average at thirteen stakeholders, with eighty-nine percent of purchases crossing multiple departments (Forrester 2024). A single content asset may influence multiple committee members at different points in the buying journey, with no single piece receiving an attributable "lead" event.

Third, a large share of buyers have a shortlist before the formal evaluation even begins. Forrester's 2024 research found that forty-one percent of B2B buyers already have a preferred vendor before formal evaluation, and prior research has indicated that buyers often enter the category with a mental shortlist of three to five brands (Forrester 2024; Bain & Company). Content that shaped that shortlist months or years earlier is invisible to last-touch attribution.

Each of these features reinforces the same conclusion. The payback horizon for B2B content is long, the influence window is long, and the mental-availability effects that shape purchase decisions are long. The measurement habits of the industry are short. Something has to give.

### **3. Methodology**

This paper is a research synthesis. It does not rely on primary or proprietary data collected by digitalmarketing.fyi. Every numeric claim is traceable to a public source cited in the References section.

The evidence corpus comprises four classes of source.

The first class is vendor-produced research with published methodology - principally HubSpot's 2017 "Compounding Blog Posts" study, which analyzed approximately twenty thousand blog posts across fifteen thousand companies, and HubSpot's ongoing State of Marketing research. Ahrefs' December 2025 study of AI Overview click-through impact, based on a three-hundred-thousand-keyword dataset, also falls in this class, as do Semrush's 2025 zero-click search analysis and the Content Marketing Institute's sixteenth annual B2B Content Marketing survey, conducted with MarketingProfs and fielded to 1,229 respondents between June and August 2025.

The second class is independent academic and institutional research: the Ehrenberg-Bass Institute for Marketing Science's work on the 95:5 rule (Dawes 2021) and the Binet and Field studies published by the IPA and extended to B2B through LinkedIn's B2B Institute.

The third class is commercial analyst research: Gartner and Forrester reports on B2B buying behavior, cited where the original methodology is disclosed.

The fourth class is secondary practitioner analysis - used sparingly, only where it aggregates or contextualizes primary research already listed in the first three classes.

The analytical framework adapts the net present value (NPV) methodology used in corporate finance and capital budgeting for any asset with expected future cash flows. Content assets have future cash flows - measurable as traffic value, leads, or pipeline contribution - and therefore can, in principle, be valued with a discount-rate approach. The paper introduces the Content Half-Life Model (CHLM) in Section 6 as a concrete adaptation of NPV logic to content portfolios. The framework is original to this paper. The underlying mathematics is not.

Three methodological caveats apply throughout. First, the empirical base is B2B-SaaS-heavy because that is where most public research has been conducted; findings should generalize to other B2B contexts with directional confidence, not precise magnitudes. Second, AI search disruption is a moving target and the figures quoted in Section 4.3 reflect data from late 2025 through early 2026; readers reviewing this paper in twelve months should expect revised numbers. Third, the selection of a discount rate in any NPV application is a judgment, not a fact, and the framework yields ranges, not point estimates.

Section 8 formalizes these caveats and adds others specific to each finding.

## 4. Findings

This section presents four findings, each drawn from the source corpus described in Section 3. Each finding is accompanied by a confidence rating following the conventions in Section 3: *High confidence* where multiple independent sources agree, *moderate confidence* where sources align directionally with variance in magnitudes, and *directional* where a pattern is visible but specific magnitudes are uncertain.

### 4.1 Content has a measurable half-life, not a binary compounding-or-decaying status

The dominant public framing of content performance divides posts into two categories: *compounding* posts that gain traffic over time, and *decaying* posts that lose it. HubSpot's 2017 research, analyzing approximately twenty thousand posts across fifteen thousand companies, established the baseline figures still cited throughout the industry: roughly ten percent of blog posts compound, and that ten percent generates approximately thirty-eight percent of total blog traffic (HubSpot Research 2017). Compounding posts reach approximately 2.5 times their initial-month traffic by month six and 3.4 times by month twenty-two (MarketingProfs 2017, summarizing HubSpot data).

Subsequent content-decay research complicates the binary. Content marketing platforms such as Clearscope and site-audit tools across the SEO industry have documented that even well-optimized evergreen content exhibits gradual performance decline - ten to thirty percent traffic loss over six to twelve months is treated as a normal lifecycle pattern, not a failure state (Clearscope 2024; Ten Speed 2025; ALM Corp 2026). The distinction between compounding and decaying becomes a matter of rate, not category.

A more accurate representation is that every content asset has a *half-life* - the time required for its organic traffic to fall to half of its peak value, in the absence of refresh activity. Pure news content has a half-life measured in days. Tactical, date-tagged listicles have half-lives measured in months. Evergreen

pillar content can have half-lives measured in years. The distribution is continuous.

**Figure 1.** Illustrative half-life distribution by content type. \*News content: half-life  $\approx$  1–7 days. Tactical listicle:  $\approx$  2–6 months. Evergreen how-to:  $\approx$  12–24 months. Pillar / research:  $\approx$  24–48 months or longer. Source: composite of HubSpot 2017, Clearscope 2024, and Ten Speed 2025.\*

The practical consequence is that portfolio-level content performance is a weighted average of many assets at different points in their half-life curves. A program publishing predominantly short-half-life content will appear to be doing fine for a few months and then plateau. A program publishing predominantly long-half-life content will appear to be doing nothing for several months and then accelerate. Measuring both at thirty or ninety days produces the same flat line and hides the underlying difference. (*High confidence.*)

## 4.2 Measurement windows shorter than the sales cycle distort ROI reporting

If a content asset's earning curve spans twelve to twenty-four months and the measurement window spans thirty to ninety days, the reported ROI is, by construction, a fraction of the true ROI. Averi's 2026 analysis puts it directly: ROI calculations conducted on three-month windows reliably understate true content performance because they exclude the bulk of the earning curve (Averi Resources 2026).

The Marketing Week survey cited earlier quantifies how widespread the mismatch is. Approximately half of the six hundred surveyed B2B marketers measured campaign outcomes over a period shorter than the organization's average sales cycle (Marketing Week 2022). In categorical terms, this means the denominator of most B2B content ROI reports captures a small minority of the income stream it is evaluating.

**Table 1.** The measurement-window mismatch in B2B.

Underlying reality	Reported timeframe	Effect on measured ROI	--- --- ---
SEO break-even: $\sim$ 9 months	3-month window	Negative / near-zero	B2B
sales cycle: $\sim$ 10 months	Quarterly review	Pipeline not yet closed	Content
half-life: 12–24 months (evergreen)	Annual review	Captures $\sim$ half the earning	

curve | | Mental-availability effect: 24-36 months | Any in-year review | Effect lies outside the window |

The compounding effect on decision-making is where the real damage is done. Content programs that appear to be underperforming on short-window dashboards are defunded. The defunding halts the investment before break-even. The program then produces the underperformance the dashboard predicted. This is a self-fulfilling failure - not because the content was bad, but because the evaluation method guaranteed it would be judged bad before it had a chance to succeed.

Only thirty-six percent of marketers across a 2026 industry aggregate report confident ability to measure content ROI (Genesys Growth 2026). Just forty-one percent of B2B marketers in HubSpot's 2024 research said they could confidently attribute revenue to specific content pieces (HubSpot 2024). The remainder operate in a measurement environment where window selection alone determines whether a program looks successful. (*High confidence.*)

### **4.3 Content faces a two-front war - decay on one side, AI-search compression on the other**

Throughout the past decade, content-portfolio planning assumed a single adversary: natural performance decay driven by competitor publication, search-intent drift, and freshness penalties. That assumption held through the end of 2023. It no longer does. A second front has opened.

Between May 2024 and the first quarter of 2026, Google's AI Overviews expanded from a limited rollout to approximately 25.8 percent of US searches (Stackmatix 2026). Independent measurement of the click-through impact has converged on a large but bounded range. Ahrefs' December 2025 dataset of three hundred thousand keywords measured a fifty-eight percent reduction in position-one organic click-through rate on queries where an AI Overview appears, an increase from the thirty-four and a half percent reduction they measured in April 2025 (Ahrefs 2025; Ahrefs 2026). Pew Research's analysis of sixty-eight thousand queries recorded a forty-seven percent relative decline. Semrush has reported that approximately 58.5 percent of US searches now conclude within the search results page without a click to an external site

(Semrush 2025). Similarweb's publisher-traffic analysis recorded a shift in overall zero-click share from fifty-six percent in May 2024 to sixty-nine percent one year later (Search Engine Journal 2025).

**Figure 2.** Independent estimates of AI-Overview organic CTR impact, 2025–2026. \*Ahrefs (Dec 2025): –58% at position 1. Pew Research: –47% relative decline. Seer Interactive: –49% to –65%. Authoritas: –47.5%. Amsive Digital: –15% average. Source variance reflects methodology and keyword-mix differences; directional consensus is a large, material CTR reduction on informational queries.\*

What this means for the content-half-life framework is specific. AI Overviews do not shorten the half-life of content in the sense of reducing its rank; a page can still occupy position one and lose clicks. What they do is *compress the traffic value* of each half-life increment. The same content asset now delivers fewer sessions per unit of rank, even if rank itself is stable. For the NPV calculation introduced in Section 6, this is equivalent to a step-function reduction in the future-cash-flow estimate - the discount rate does not change, but the nominal cash flows do. Long-half-life evergreen assets remain the most defensible in this environment because they are more likely to be cited as sources within AI Overviews and because they accumulate the signals - depth, freshness, backlinks, and structured data - that AI ingestion pipelines reward (Ten Speed 2025; Decoding 2026).

Gartner projects that twenty-five percent of traditional organic search traffic will shift to AI chatbots and voice assistants by the end of 2026 (Gartner, as cited in Martal 2026). Whether the exact percentage proves right or wrong, the direction is clear: content that survives into the AI-mediated search era will be content with long enough half-lives to compound past the compression. (*Moderate confidence on magnitudes; high confidence on direction.*)

#### **4.4 B2B buying is a 95:5 mental-availability game, not a 5% conversion game**

The final finding is the one that reframes everything above. Professor John Dawes of the Ehrenberg-Bass Institute documented in 2021 what has since become known as the 95:5 rule: at any given time, only approximately five percent of a B2B category's potential buyers are actively considering a purchase. The remaining ninety-five percent are out of market (Dawes 2021).

The mathematics of the rule follow from average inter-purchase intervals. If a corporation changes its primary law firm or banking provider once every five years on average, then roughly twenty percent of the buyer base is in-market in any given year and approximately five percent in any given quarter. The number is not a precise constant across categories. It is a useful heuristic.

The implication for content is profound. The primary job of B2B content marketing is not to convert the five percent. It is to build and maintain mental availability - the memory structures that cause buyers to recall a vendor, associate it with their need, and include it in the initial consideration set - across the ninety-five percent who will be in-market later (Dawes 2021; Romaniuk and Sharp 2016). Binet and Field's research with the IPA and, in B2B, with LinkedIn, reaches a parallel conclusion: long-term brand-building spend produces greater cumulative business effects than short-term activation spend, and in B2B contexts the optimal split trends closer to 50/50 than the canonical 60/40 B2C ratio (Binet and Field 2013; Binet and Field 2019).

Content is the most cost-efficient medium for building mental availability at B2B scale. Paid brand advertising works, but per-impression costs are high and decay immediately after the flight. A published evergreen article accumulates reach over months and years and does not stop working when a budget cycle closes. The 95:5 rule thus recasts what a content program *is*: less a demand-capture machine, more a compounding memory-investment vehicle whose dividends are paid out on the timeline of the category's inter-purchase cycle.

Forrester's finding that forty-one percent of B2B buyers have a preferred vendor before formal evaluation underscores the point (Forrester 2024). The content that put a vendor onto that shortlist was consumed months or years before the attributable lead event. Any measurement framework that ignores out-of-market influence will understate content's true contribution by the same order of magnitude. (*High confidence.*)

## 5. Strategic implications

The four findings above point to a coherent strategic posture for B2B content programs. It has five components.

First, the default measurement window for content programs should be expanded to match the underlying sales cycle at minimum, and ideally to the category's inter-purchase interval. For most B2B SaaS contexts this is a twelve-to-twenty-four-month window. Short-window dashboards may remain useful as operational signals (has publishing velocity dropped; are new posts indexed) but should not drive strategic funding decisions.

Second, budget planning should separate asset-creation spend from asset-maintenance spend, and should explicitly recognize refresh investments as capital-preserving expenditure. Treating refresh as a cost to be minimized tends to accelerate decay. Treating refresh as the analog of equipment maintenance, with a defensible return of preserved mental availability, reframes the decision correctly.

Third, portfolio composition should be designed around an intentional distribution of half-lives. A program consisting solely of short-half-life tactical content will miss the compounding returns that sustain long-term growth. A program consisting solely of long-half-life pillar content will starve near-term pipeline. The balanced portfolio typically includes both, in ratios reflecting the team's current market position: earlier-stage programs lean toward higher long-half-life share to build the foundation; mature programs can allocate more to tactical capture.

Fourth, reporting to the CFO and board should present content performance using the language of asset valuation, not the language of campaign ROI. Present value of the content library, estimated half-life by portfolio segment, and decay-adjusted forward cash flow are more defensible framings than month-on-month traffic charts. Finance leaders understand capital assets. Meeting them in that language raises the quality of the content conversation at the executive level.

Fifth, AI search requires a new layer in the portfolio: citation-optimized content designed not for click capture but for citation by AI answer engines. This does not replace traditional SEO; it complements it. Research content, primary-data-bearing content, and well-structured definitional content are disproportionately likely to be cited by AI Overviews and by conversational search tools (Decoding 2026). The same content assets also tend to attract backlinks and occupy the long-half-life end of the distribution - the investments reinforce each other.

## 6. A framework - The Content Half-Life Model

The Content Half-Life Model (CHLM) is a portfolio-valuation framework for B2B content. It adapts the net present value methodology from capital budgeting to the specific decay and compounding dynamics of owned content. It is original to this paper.

The model has four components. Each is intended to be usable by a marketing leader without specialized finance training.

### 6.1 The content NPV formula

For a given content asset, the expected net present value is the sum of its discounted future monthly cash flows minus its initial production cost. Formally:

$$\text{Content NPV} = \sum [ (\text{monthly traffic value at month } t) \times (1 - \text{decay rate})^t ] / (1 + \text{discount rate})^t - \text{production cost}$$

Where:

- *Monthly traffic value at month t* is the organic traffic arriving to the asset in month t, multiplied by a per-visit value (typically the blended conversion rate to pipeline multiplied by the average contract value, or more simply the CPC the team would pay for equivalent paid traffic).
- *Decay rate* is the monthly percentage decline in traffic absent refresh. For evergreen content with a twenty-four-month half-life, this is approximately 2.9% per month (since  $0.971^{24} \approx 0.5$ ). The decay rate is derived from the chosen half-life category.
- *Discount rate* reflects the cost of capital to the organization and the risk premium appropriate to marketing investments. A typical SaaS-startup discount rate is 15% annually, or approximately 1.17% per month (Paubox 2023).
- *Production cost* includes writer time, editor time, design, and attributable platform cost.

The formula's virtue is that it forces explicit assumptions. Teams that cannot estimate traffic value, decay rate, or discount rate for their content are in effect conceding that they cannot value the asset at all - which should be the first actionable finding of any content portfolio review.

## 6.2 Half-life categories

Four categories capture most B2B content. Empirical half-life estimates are directional, drawn from HubSpot's compounding-post data and industry decay observations (HubSpot Research 2017; Clearscope 2024).

**Table 2.** Empirical half-life categories for B2B content.

Category	Examples	Half-life (months)	Implied monthly decay rate
News / reactive	Launch announcements, event reactions	0.5-3	20-75%
Tactical	"Best tools for X in 2026", dated listicles	2-6	11-30%
Evergreen how-to	"How to do Y", framework explainers	12-24	2.9-5.6%
Pillar / research	Whitepapers, proprietary-data studies, hub pages	24-48+	1.4-2.9%

Teams should categorize their content at creation time, not after the fact. Retroactive categorization introduces hindsight bias and undermines the forward-looking use of the framework.

## 6.3 The publish-refresh-retire decision matrix

Each existing content asset in a portfolio can be classified against two dimensions: its current position on the decay curve (rising, peaking, declining, flat-and-low), and its half-life category (news, tactical, evergreen, pillar).

Three actions follow, mechanically:

- **Publish more like this** - when an asset is in evergreen or pillar category and is rising or peaking. It is both a net contributor of NPV and a template for future production.
- **Refresh and re-promote** - when an evergreen or pillar asset is declining. The production cost is typically twenty-to-forty percent of a new piece and the expected NPV uplift is substantially higher than producing a new piece of equivalent quality, because the asset already carries backlinks, domain signals, and accumulated crawl history.
- **Retire or consolidate** - when an asset is in news or tactical category and has passed its peak. Continued investment is negative-NPV. Where the asset has inbound links or brand mentions, 301-redirecting to a surviving pillar preserves the equity.

## 6.4 Applying the model

In practice, the CHLM produces three numbers a marketing leader can take to the CFO: estimated current NPV of the content library, projected NPV in twelve and twenty-four months under current spend, and the marginal NPV of the next dollar of investment across the publish/refresh/retire actions. The precision of these numbers is less important than the discipline they impose. A team that knows its discount rate, its blended per-visit value, and its portfolio's weighted-average half-life has a far more defensible position than a team presenting month-over-month traffic lines.

## 7. Case illustration - an applied CHLM scenario

The following is a composite example, illustrative only. It is not a record of a specific client engagement.

Two fictional B2B SaaS content teams, each with a \$300,000 annual content budget, publish for twenty-four months. Team A operates on a ninety-day measurement window and optimizes for quarterly traffic. Team B adopts the CHLM and optimizes for portfolio NPV over a twenty-four-month horizon.

Team A's spend skews toward tactical and news content because those formats show the fastest traffic inflection inside the quarterly window. Roughly seventy percent of its publishing volume is in those half-life categories. At the end of month three, Team A reports a 42% traffic increase over baseline and secures board confidence. At the end of month twelve, Team A's portfolio is flat against its own month-nine peak because the short-half-life content has decayed. At the end of month twenty-four, Team A's portfolio has returned to baseline.

Cumulative pipeline attributed to content is modest.

Team B's spend allocates approximately fifty-five percent to evergreen and pillar content, thirty-five percent to tactical, and ten percent to news. At the end of month three, Team B's reported traffic increase is 11% - materially below Team A - and the program faces funding scrutiny. At the end of month twelve, Team B's portfolio is approximately 90% above baseline, driven by the compounding evergreen and pillar assets reaching their compounding inflection. At the end of

month twenty-four, Team B's portfolio is approximately 340% above baseline. Cumulative pipeline is multiples of Team A's.

The illustrative numbers in this scenario are consistent with the compounding curves documented in the HubSpot 2017 dataset (2.5× initial traffic by month six, 3.4× by month twenty-two) but they are not a prediction. They are an expression of what the CHLM framework would recommend if applied consistently. The actual outcome in any real team depends on execution quality, topic selection, and the competitive context - all of which the framework helps to diagnose but none of which it can guarantee.

The case's point is the decision-making difference, not the specific numbers. Team A's measurement system made it rational, at every decision point in the first year, to allocate toward short-half-life content. Team B's measurement system made it rational to accept lower early numbers in exchange for a higher long-run position. Both teams acted rationally given their measurement windows. The windows determined the outcome.

## 8. Limitations and assumptions

Five limitations are material to the findings in this paper.

First, the public research base is weighted toward B2B SaaS. The Content Marketing Institute, HubSpot, Ahrefs, and Semrush produce the majority of available benchmarks and oversample SaaS audiences relative to the broader B2B economy. Industrial B2B, B2B services, and high-consideration consumer categories may exhibit different half-life distributions. The directional conclusions of this paper should generalize; the specific magnitudes should not be transferred without local validation.

Second, the HubSpot compounding-posts study - the largest public dataset on content half-life dynamics - was published in 2017 and has not been formally updated at the same scale. Subsequent industry writing has extended and contextualized it, but the underlying data is now nine years old. The broad patterns observed there appear to have held, but the specific 10%-compound / 38%-of-traffic figures are citation-worn and should be treated as foundational rather than current.

Third, the AI Overview impact data quoted in Section 4.3 reflects a rapidly evolving situation. The 58% position-one CTR reduction measured by Ahrefs in December 2025 is the most recent publicly available figure at the time of writing; if readers are reviewing this paper materially later than April 2026, that figure should be re-verified. The direction of the trend - material, sustained CTR reduction on informational queries - is likely to hold, but point estimates should not.

Fourth, the CHLM's core formula requires three inputs that most marketing teams do not currently estimate: a blended per-visit value, a decay rate calibrated to the team's own content, and a discount rate reflecting the organization's cost of capital. Teams without existing analytics infrastructure (pipeline attribution, cohorted traffic decay tracking, CFO partnership on cost of capital) will find the framework directionally useful but numerically underspecified in its first application. The remedy is not to abandon the framework but to adopt its discipline progressively, starting with rough estimates and refining as data matures.

Fifth, this paper does not address paid content distribution, influencer content, or co-marketed content. The half-life dynamics of those channels differ meaningfully and merit separate treatment.

These limitations do not, in the authors' judgment, undermine the paper's central claim that B2B content is systematically undervalued by short measurement windows and that a half-life / NPV framing is a better alternative. They do constrain the precision with which that claim should be applied to any specific team's portfolio.

## 9. Conclusion

B2B content marketing has a measurement problem that looks like a performance problem. The underlying asset has a long earning curve. The dominant evaluation windows are short. The gap between them systematically causes content programs to be judged before they have had a chance to earn. Teams defund content when the dashboards say "not working" - producing the very failure the dashboards predicted.

This paper has argued for a different lens. Content is a capital asset. It has a half-life. It produces cash flows that are worth discounting to present value. It competes for capital, and it deserves to be defended in the language of capital. The Content Half-Life Model introduced here is one concrete way to apply that discipline. Teams will refine it with their own data; the specific numbers will shift; the framework itself is intended to be usable in practice, not perfect in theory.

Two forward-looking notes close the paper. First, the AI-mediated search environment emerging in 2026 is a stress test for exactly the half-life dynamics described here. Content that was already compounding will compound through the transition; content that was already short-lived will become shorter-lived still. Portfolios weighted toward long-half-life research and pillar assets will outperform portfolios weighted toward tactical content by a wider margin in an AI-search world than they did in a classical-search world. Second, the authors intend to publish a future edition of this paper with proprietary first-party data - cohorted traffic decay curves across a larger B2B content corpus - which will allow the half-life category estimates in Table 2 to be tightened into empirical ranges rather than directional bands.

Until then, the recommendation to the marketing leader reading this paper is short. Extend the measurement window. Catalog the half-lives in the current portfolio. Identify which assets are capital-preserving, capital-appreciating, or capital-destroying. Refresh the evergreen. Retire the decayed. And when the CFO asks why content needs another year of budget, answer in the language of assets, not the language of posts.

# References

References are listed in Chicago 17th Author-Date style. URLs were current at the access dates indicated.

Ahrefs. 2025. "AI Overviews Reduce Clicks by 34.5% (April 2025 Study)." Ahrefs Blog. <https://ahrefs.com/blog/ai-overviews-reduce-clicks/>. Accessed April 18, 2026.

Ahrefs (Law, Ryan). 2026. "Update: AI Overviews Reduce Clicks by 58%." Ahrefs Blog. <https://ahrefs.com/blog/ai-overviews-reduce-clicks-update/>. Accessed April 18, 2026.

ALM Corp. 2026. "Content Decay: What It Is, Why It Happens, and How to Fix It." ALM Corp. <https://almcorp.com/blog/content-decay/>. Accessed April 18, 2026.

Averi Resources. 2026. "Content Marketing ROI Benchmarks by Company Size (2026 Data)." Averi.ai. <https://resources.averi.ai/benchmarks/content-marketing-roi-benchmarks>. Accessed April 18, 2026.

Averi (SaaS Guide). 2026. "Content Marketing ROI Benchmarks for B2B SaaS." Averi.ai. <https://www.averi.ai/guides/content-marketing-roi-benchmarks-b2b-saas>. Accessed April 18, 2026.

Binet, Les, and Peter Field. 2013. "The Long and the Short of It: Balancing Short and Long-Term Marketing Strategies." Institute of Practitioners in Advertising (IPA). <https://ipa.co.uk/knowledge/publications-reports/the-long-and-the-short-of-it/>. Accessed April 18, 2026.

Binet, Les, and Peter Field. 2019. "The B2B Institute Research with LinkedIn: Applying Long-Term Marketing Effects to B2B." LinkedIn B2B Institute. <https://business.linkedin.com/marketing-solutions/b2b-institute>. Accessed April 18, 2026.

Clearscope. 2024. "How to Fix Content Decay: 6 Top Strategies." Clearscope. <https://www.clearscope.io/blog/content-decay>. Accessed April 18, 2026.

Content Marketing Institute and MarketingProfs. 2025. "B2B Content Marketing Benchmarks, Budgets, and Trends - 16th Annual Research." Content Marketing

Institute. <https://contentmarketinginstitute.com/b2b-research/b2b-content-marketing-trends-research>. Accessed April 18, 2026.

Dawes, John. 2021. "Advertising Effectiveness and the 95-5 Rule: Most B2B Buyers Are Not in the Market Right Now." Ehrenberg-Bass Institute for Marketing Science and the LinkedIn B2B Institute. <https://marketingscience.info/advertising-effectiveness-and-the-95-5-rule-most-b2b-buyers-are-not-in-the-market-right-now/>. Accessed April 18, 2026.

Decoding. 2026. "Content Decay: How to Identify, Fix, and Prevent Traffic Loss." Decoding. <https://trydecoding.com/blog/content-decay-how-to-identify-fix/>. Accessed April 18, 2026.

Ehrenberg-Bass Institute for Marketing Science. 2025. "The 95:5 Rule Is the New 60:40 Rule." Marketing Science. <https://marketingscience.info/news-and-insights/the-955-rule-is-the-new-6040-rule>. Accessed April 18, 2026.

Forrester Research. 2024. "The State of Business Buying, 2024 Report." Forrester. <https://www.forrester.com/report/the-state-of-business-buying-2024>. Accessed April 18, 2026.

Gartner. 2025. "Gartner Sales Survey Finds 74 Percent of B2B Buyer Teams Demonstrate Unhealthy Conflict During the Decision Process." Gartner. <https://www.gartner.com/en/newsroom/press-releases/2025-05-07-gartner-sales-survey-finds-74-percent-of-b2b-buyer-teams-demonstrate-unhealthy-conflict-during-the-decision-process>. Accessed April 18, 2026.

Gartner. 2025. "The B2B Buying Journey: Research on Buying Jobs and the Rep-Free Experience." Gartner. <https://www.gartner.com/en/sales/insights/b2b-buying-journey>. Accessed April 18, 2026.

Genesys Growth. 2026. "Content Marketing ROI - 45 Statistics Every Marketing Leader Should Know." Genesys Growth. <https://genesysgrowth.com/blog/content-marketing-roi-stats-for-marketing-leaders>. Accessed April 18, 2026.

HubSpot Research. 2017. "Compounding Blog Posts: What They Are and Why They Matter." HubSpot. <https://blog.hubspot.com/marketing/compounding-blog-posts-what-they-are-and-why-they-matter>. Accessed April 18, 2026.

HubSpot. 2024. "State of Marketing Report: Attribution and Measurement Data." HubSpot. <https://www.hubspot.com/state-of-marketing>. Accessed April 18, 2026.

Konabayev. 2026. "B2B Marketing Benchmarks 2026: 50+ Stats on Conversion, CAC, and ROI." Konabayev. <https://konabayev.com/blog/b2b-marketing-benchmarks-2026/>. Accessed April 18, 2026.

MarketingProfs (Bruce, Amanda). 2017. "Four Tips to Boost Traffic and Leads with Compounding Blog Posts." MarketingProfs. <https://www.marketingprofs.com/articles/2017/31457/four-tips-to-boost-traffic-and-leads-with-compounding-blog-posts>. Accessed April 18, 2026.

Marketing Week. 2022. "The Great Balancing Act: The Long and Short of B2B Marketing." Marketing Week. <https://www.marketingweek.com/long-short-b2b-marketing/>. Accessed April 18, 2026.

Martal Group. 2026. "B2B Digital Marketing Benchmarks 2026: ROI Metrics and Results." Martal. <https://martal.ca/b2b-digital-marketing-benchmarks-lb/>. Accessed April 18, 2026.

Paubox. 2023. "Using Net Present Value (NPV) to Drive Marketing Decisions." Paubox Blog. <https://www.paubox.com/blog/net-present-value-marketing-decisions>. Accessed April 18, 2026.

Pew Research Center. 2025. "Google Users Are Less Likely to Click on Links When an AI Summary Appears in the Results." Pew Research Center. <https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/>. Accessed April 18, 2026.

Romaniuk, Jenni, and Byron Sharp. 2016. "How Brands Grow: Part 2 (Revised Edition)." Oxford University Press. <https://global.oup.com/academic/product/how-brands-grow-part-2-9780190330064>. Accessed April 18, 2026.

Search Engine Journal (Southern, Matt G.). 2025. "Google AI Overviews Impact on Publishers and How to Adapt into 2026." Search Engine Journal. <https://www.searchenginejournal.com/impact-of-ai-overviews-how-publishers-need-to-adapt/556843/>. Accessed April 18, 2026.

Semrush (and Datos). 2025. "Semrush AI Overviews Study: What 2025 SEO Data Tells Us About Google's Search Shift." Semrush. <https://www.semrush.com/blog/semrush-ai-overviews-study/>. Accessed April 18, 2026.

6sense. 2025. "B2B Buyer Experience Report 2025." 6sense. <https://6sense.com/resources/ebook/b2b-buyer-experience-report/>. Accessed April 18, 2026.

Stackmatix. 2026. "Google AI Overview SEO Impact: 2026 Data and Statistics." Stackmatix. <https://www.stackmatix.com/blog/google-ai-overview-seo-impact>. Accessed April 18, 2026.

Ten Speed. 2025. "Content Decay: How to Identify and Fix to Unlock Organic Growth." Ten Speed. <https://www.tenspeed.io/blog/content-decay>. Accessed April 18, 2026.

# Reproducibility Note

The findings in this whitepaper can be independently verified as follows. Where specific datasets are referenced, source URLs and access dates are listed in the References section.

- The HubSpot compounding-post figures cited in Section 4.1 can be verified by reading the original 2017 HubSpot Research publication linked in the References. Readers can replicate the compounding analysis at the portfolio level by pulling monthly Google Search Console data for all blog URLs published at least twenty-four months prior, computing each URL's ratio of current-month clicks to initial-month clicks, and binning the results.
- The AI Overview click-through-rate reductions cited in Section 4.3 can be independently triangulated by comparing the Ahrefs, Pew, Semrush, and Similarweb sources in the References. Readers with access to their own Google Search Console data can compute a local estimate by segmenting queries by AI-Overview presence and comparing observed CTR against the Advanced Web Rankings CTR curves.
- The 95:5 rule in Section 4.4 can be approximated for a reader's own category by identifying the average inter-purchase interval for the product or service and taking its reciprocal as the proportion of buyers in-market annually, divided by four for quarterly.
- The Content Half-Life Model formula in Section 6.1 can be implemented in a spreadsheet with three inputs: monthly traffic value ( $CPC \times$  monthly organic sessions, or pipeline value per visit), monthly decay rate (derived from the half-life category in Table 2), and monthly discount rate (annual cost of capital divided by twelve). The spreadsheet should compute twenty-four monthly cash flows, discount each to present value, sum them, and subtract production cost.

## About digitalmarketing.fyi

digitalmarketing.fyi is a digital marketing agency and research publisher. We produce whitepapers and strategic research for marketing leaders building durable growth programs across SEO, paid media, content, email, automation, analytics, and AI search.

Website: <https://digitalmarketing.fyi>

## Appendix: Machine-Readable Metadata

The structured-data block below mirrors the JSON-LD schema published with this paper's HTML landing page. It is included here so that AI ingestion pipelines can extract citation metadata directly from the document.

```
{
  "@context": "https://schema.org",
  "@type": "Report",
  "name": "The Content Half-Life Paradox",
  "reportNumber": "DMFYI-WP-2026-001",
  "datePublished": "2026-04-18",
  "author": "digitalmarketing.fyi",
  "publisher": "digitalmarketing.fyi",
  "inLanguage": "en-US",
  "license": "https://creativecommons.org/licenses/by-nc/4.0/",
  "keywords": [
    "B2B content marketing strategy",
    "content marketing ROI",
    "content decay",
    "content marketing measurement",
    "mental availability",
    "95:5 rule",
    "AI Overviews SEO"
  ],
  "url": "https://digitalmarketing.fyi/whitepapers/content-half-life-paradox",
  "abstract": "B2B content marketing generates measurable returns over twelve-to-twenty-four-month horizons, yet the industry evaluates it on thirty-to-ninety-day windows. This paper argues that the resulting valuation error is the dominant reason B2B content programs are defunded before they reach their earning potential. Drawing on public research from HubSpot, the Content Marketing Institute, Ahrefs, Semrush, Gartner, Forrester, and the Ehrenberg-Bass Institute, the analysis shows that content assets have a continuous half-life distribution rather than a binary compounding-or-decaying status; that measurement windows shorter than the sales cycle systematically understate ROI; that AI Overviews compress the traffic value of each half-life increment by approximately thirty-to-fifty-eight percent on informational queries; and that B2B content's primary job is not to convert the five percent of buyers in-market today but to build mental availability across the ninety-five percent who will enter the market later. The paper introduces the Content Half-Life Model, a net-present-value framework that adapts capital-budgeting methodology to content portfolios. The framework quantifies publish,
```

refresh, and retire decisions and presents content performance in the language of asset valuation rather than campaign ROI. The intended reader is a B2B marketing leader defending a multi-year content investment to a finance-minded executive audience."

}