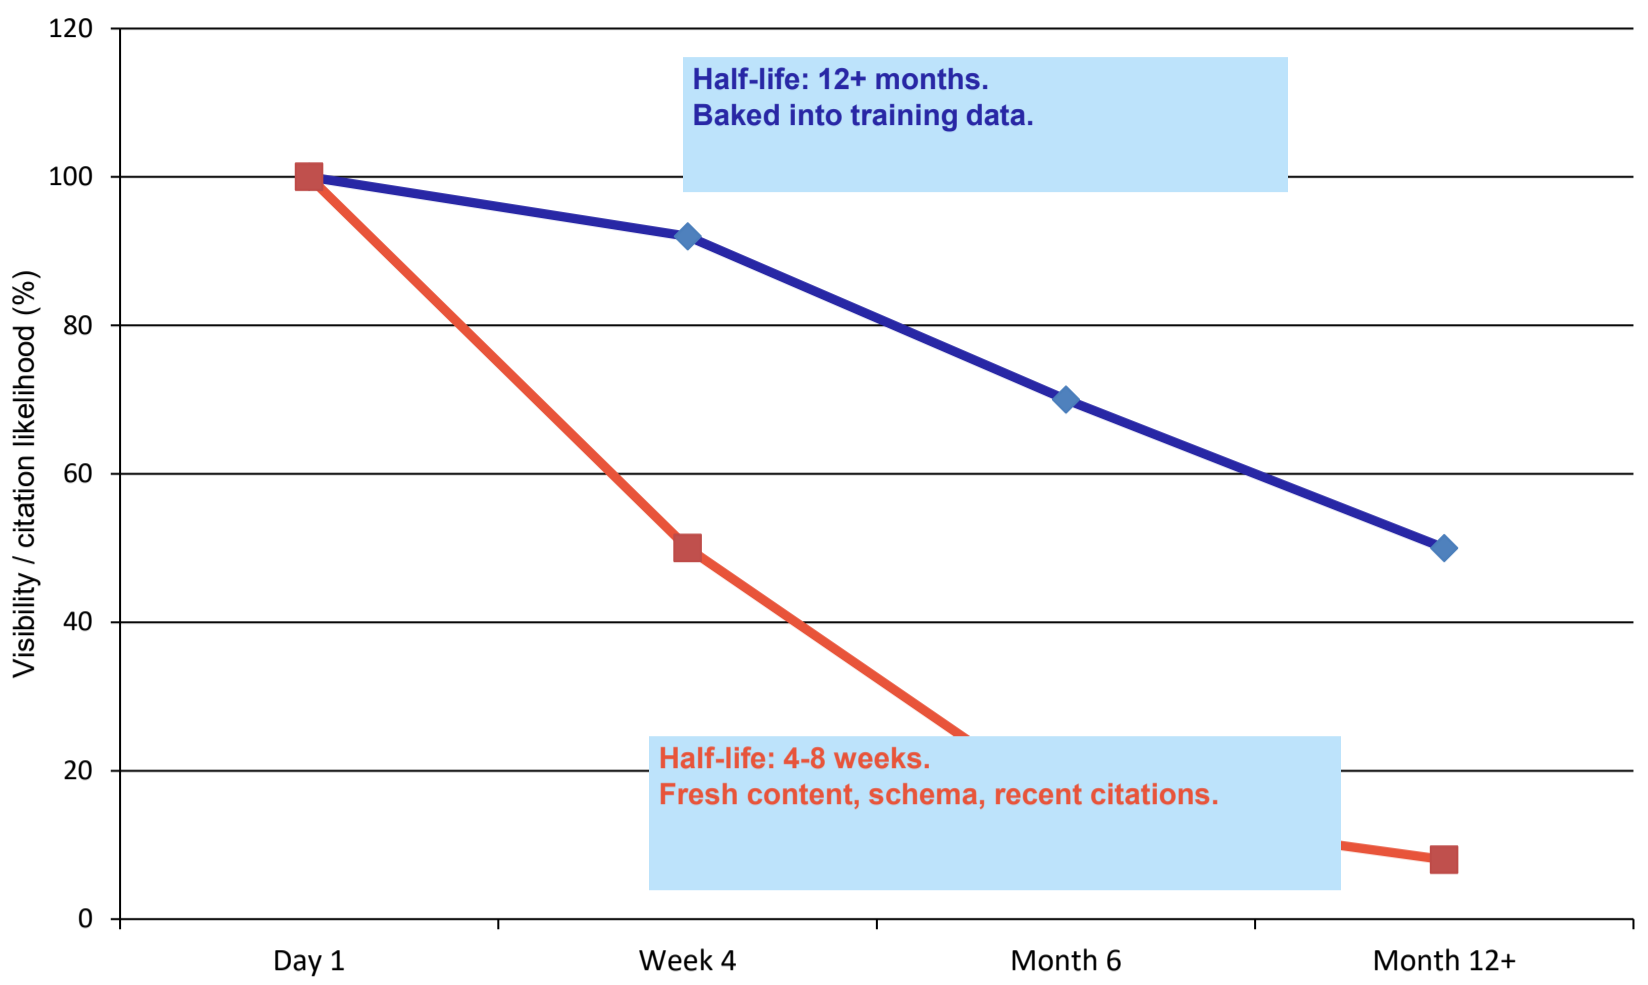


THERE ARE TWO KINDS OF AI MEMORY

Most GEO advice optimizes for the wrong one

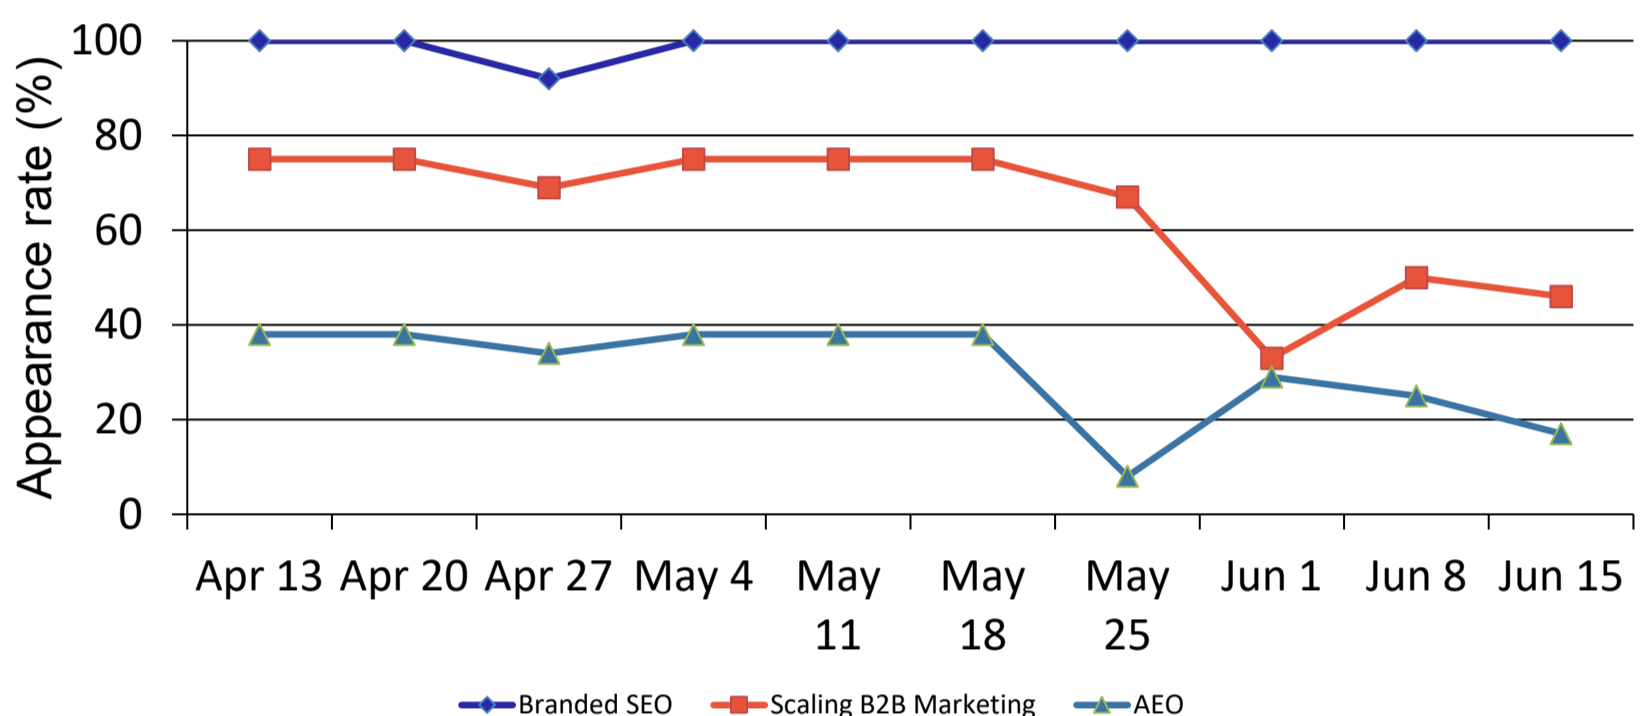
THE MODEL · A HYPOTHESIS, NOT YET MEASURED AT 12-MONTH SCALE



Illustrative half-lives, not measured values. The two-curve hypothesis is informed by Erdem & Keane (1996) and Petroni et al. (2019). Measured bjrees.com data follows below.

THE EVIDENCE · bjrees.com, GOOGLE SEARCH + AI MODE, WEEKLY

Klara visibility tracking: real appearance rate by category



100% → 100%

Branded SEO

75% → 46%

Scaling B2B Marketing

38% → 17%

AEO

The category trend above is real - both topical categories fell. But it isn't one smooth decay: of the 12 non-branded queries behind it, 6 never ranked at all, 3 declined and stayed down, and 3 dipped then fully recovered. Branded queries did not move. Per-query detail below.

Category	Query	First	Min	Last	Verdict
Branded SEO	Branded SEO - 3 queries (avg)	100%	92%	100%	Stable
Scaling B2B Marketing	challenges when scaling B2B marketing from smb to enterprise	100%	0%	0%	Declined
Scaling B2B Marketing	enterprise marketing strategy for scaling a B2B company	0%	0%	0%	Never ranked
Scaling B2B Marketing	how to scale B2B marketing from smb to enterprise	100%	50%	83%	Recovered
Scaling B2B Marketing	scaling marketing from smb to enterprise	100%	67%	100%	Recovered
AEO	AI marketing blog	0%	0%	0%	Never ranked
AEO	What's the future of SEO in the age of generative AI?	0%	0%	0%	Never ranked
AEO	When was Generative Engine Optimization first defined?	0%	0%	0%	Never ranked
AEO	Who coined the term Generative Engine Optimization?	0%	0%	0%	Never ranked
AEO	Why is measurement so important for the culture of a marketing organisation?	0%	0%	0%	Never ranked
AEO	how do I scale B2B marketing from smb to enterprise	100%	0%	0%	Declined
AEO	how do I scale marketing from smb to enterprise	100%	0%	100%	Recovered
AEO	scale from smb to enterprise	100%	0%	33%	Declined

35 capture runs, Apr 13 - Jun 15 2026, Google Search (Regular). First/Min/Last are weekly appearance rate for that single query. 'Never ranked' queries never appeared once in this window - that is a content gap, not decay.

THE MECHANISM · WHY THE CURVES DIVERGE

PARAMETRIC MEMORY

Fixed at training time

Brand-domain associations are encoded directly into a model's weights during pretraining. The more a name and its facts co-occur in training data, the stronger that internal association becomes (Petroni et al., 2019). Once trained, it is not recomputed: it holds until the next training run, regardless of what is freshly indexed today.

SIGNAL-BASED MEMORY

Recomputed at query time

Rankings for non-branded topical queries are not stored, they are re-estimated on every query from live signals: crawl freshness, competing content, citation recency. Each result is closer to a Bayesian belief update than a lookup (Erdem & Keane, 1996): without reinforcing signals, the posterior reverts toward whichever source has the freshest evidence.

This is also why outcomes split rather than move together: "challenges when scaling B2B marketing..." fell from 100% to 0% with no recovery, while "how do I scale marketing..." dipped to 0% and fully recovered. Same mechanism, different competing signals each week.

BRANDED QUERIES

Parametric-resistant.
Ride the slow curve.

TOPICAL / NON-BRANDED QUERIES

Signal-sensitive.
Need constant refresh.

Branded visibility is inherited. Topical visibility is rented, and the rent is due every six weeks.