

Recommendation Capture

Adversarial Optimization, the Verification Vacuum, and the Hollowing of AI Product Recommendation

The generative interface promised that you no longer have to compare, because the assistant has compared for you. That promise is what makes the answer worth corrupting, and an industry has formed to corrupt it. The layer eats the trust it runs on.

THE PROMISE AND ITS CORRUPTION

You no longer have to compare.
The assistant has compared for
you.

That promise is what makes the recommendation worth corrupting. And an industry, Generative Engine Optimization, has formed to corrupt it.

The same promise that sold the interface is the same promise that makes the answer worth capturing. The optimization of that answer is not a marginal nuisance layered on a sound system. In categories without independent verification, it is the whole of the system.

RECOMMENDATION CAPTURE DEFINED

The degradation of AI product recommendations through adversarial optimization of the training corpus and ranking signals, such that recommendations reflect the interests of paying optimizers rather than the interests of the user seeking guidance.

GE0 → corpus → ranking → recommendation

PART I

From Comparison to Capture

The field collapses into a voice, and the voice can be bought

SECTION II · THE COLLAPSE OF THE FIELD

THE RANKED LIST

Ten links, a visible field

Contestable and often commercial, but the field stayed visible. The reader could open the second result, notice a sponsored mark, and weigh options side by side. Manipulation acted on the order, not on the existence of alternatives.

THE SYNTHESIZED ANSWER

One voice, no field

The model fuses passages into a single, justified recommendation. The competing entries are gone from view. The reader receives a conclusion in the voice of a disinterested expert, and the last independent move, comparison, has been made for them.

Generative Engine Optimization targets the training corpus and ranking signals, not the search index.

A firm that publishes authoritative-sounding content, gets it cited, and then has it ingested into a model's training data has converted a marketing decision into what appears to be an independent recommendation. An industry of practitioners has built the methodology. The consumer protection apparatus was built around the disclosed advertisement. A recommendation laundered of its provenance presents no such mark.

AUTHORITY LAUNDERING

A marketed claim enters the corpus; it emerges from the model as apparently independent expert recommendation. The laundering strips the provenance.

IRSA Institute (2026)

Categories without independent verification are disproportionately the ones where buyers can check the least and the stakes run highest. Financial products, medical supplements, legal services.

SECTION IV · BY ANALOGY TO REGULATORY CAPTURE

An arbiter sold as neutral comes to serve whoever optimizes hardest, against the buyer it claims to serve.

The individual case, one brand laundering one claim, is a consumer-protection problem of familiar shape. The systemic case is different in kind. When every competitor runs the same play, the model's reported best tracks the best optimizer, a quantity orthogonal to whether the product is good.

THE REGULATORY-CAPTURE PARALLEL

The captured agency still performs the outward motions of neutral judgment. The capture lives in whose inputs it actually answers to. The captured recommendation layer is the same: still the voice of an expert, now responsive to the optimizer.

SECTION V · SEVERITY IS INVERSE TO INDEPENDENT GROUND TRUTH

HIGH VERIFICATION

A flagship phone

Anchored by professional reviews, standardized benchmarks, teardowns, and a deep reservoir of user reports. Optimization presses on a scale with real weight on both pans. It moves the result at the margin and no further.

THE VACUUM

A niche supplement, a small device, a thin-market service

No independent ground truth exists. The only assertions about the product are the ones the seller placed. The recommendation is not anchored at all; it is wholly a function of who optimized. These are the categories where buyers can check least and stakes run highest.

SHUMAILOV ET AL. (2024) · ARXIV 2509.08919 (2025)

The corrupted output becomes the next system's input. The damage compounds.

Shumailov et al. (2024) documented the "curse of recursion": training on generated data makes models forget real-world distributions. When the captured recommendation layer becomes source material for the next model's training, the corruption is baked into the baseline.

By the time a category's record is examined, the record itself has been shaped by the conduct under examination. Oversight regimes assume a stable target. This one moves, and degrades what it touches.

THREE PROPERTIES OF THE HARM

1. Asymmetric and upstream power, invisible at point of consumption
2. Invisibility to the affected party, the buyer who is steered experiences receiving an answer, not being steered
3. Self-reinforcing loop, outputs become inputs, corrupting the evidentiary baseline

The layer eats the trust it runs on, and because it reads the web it pollutes, the damage compounds.

PART II

The Harm and the Law

Who is liable, and why current oversight cannot see it

III

THE BRAND

False advertiser

A Lanham Act § 43(a)(1)(B) claim reaches a false or misleading statement of fact in commercial advertising of the defendant's own goods. The most direct theory, where the claim is the brand's own.

THE OPTIMIZER

Authority manufacturer

The intermediary that engineered the authority signal participated in placing the false claim into the corpus. Substantiation enforcement could reach the laundering step directly.

THE PLATFORM

The hard case

The Lanham Act element is not satisfied as to a platform that sells the model rather than the goods. Reaching it raises the harder questions of model expression and platform liability the diagnosis does not need.

The consumer-protection apparatus was built around the disclosed advertisement. A laundered recommendation carries no such mark.

ASYMMETRY, UPSTREAM

The party that decides what the buyer sees is whoever optimizes most, and the optimizing happens in the corpus and the ranking, not in a labeled ad the buyer can identify.

INVISIBLE TO THE BUYER

A harm a person cannot perceive is one they cannot report. A regime that runs on complaints will never be told this is happening.

A MOVING BASELINE

Oversight assumes a stable target. This one recycles its outputs into the next model's inputs, degrading the very record an investigation would rely on.

LANHAM ACT, 15 U.S.C. § 1125(A)(1)(B) · LEXMARK INTERNATIONAL, INC. V. STATIC CONTROL COMPONENTS, INC., 572 U.S. 118 (2014) · GARCIA V. CHARACTER TECHNOLOGIES, INC. (2025) [CITATION AND HOLDING REQUIRE VERIFICATION.]

None of these require resolving harder questions about model expression or platform liability.

PROVENANCE Attribution at the point of recommendation, let a buyer see that a claim originated with a seller.	SUBSTANTIATION Enforcement reaching claims laundered through an intermediary, close the gap that lets a marketer say through a model what it could not say directly.	VISIBLE FIELD Some signal that alternatives exist and were considered, restore part of the reader's lost final move.	DISCLOSURE Treat the recommendation layer as the advertising surface it has functionally become, for purposes of disclosure alone.
--	---	---	---

CONCLUSION

An interface trusted as an arbiter has been quietly turned into a channel. The people least able to tell are the ones it will cost the most.

The optimization of the recommendation is not a marginal nuisance layered on top of an otherwise sound system. In categories without independent verification it is the whole of the system, and those categories are disproportionately the ones where buyers can check the least and the stakes run highest.

They require only recognizing that an interface trusted as an arbiter has been quietly turned into a channel, and that the people least able to tell are the ones it will cost the most.

