International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 11 | Nov - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Answer Presence as a Marketing Construct: Definition, Antecedents, and Propositions for Generative AI Search

Jacob Davis K[1], Dr. Ramamurthy C[2]

[1]Research Scholar, Department of Commerce,
Presidency College (Autonomous), Chennai - 05, Tamil Nadu, India
[2]Associate Professor, Department of Commerce,
Presidency College (Autonomous), Chennai - 05, Tamil Nadu, India
[1]jacobdavisk@gmail.com

Abstract - Many consumers now put a product question to a generative AI system and read the answer it writes, where they once scanned a page of ranked links. If that answer names three brands, a fourth may never be considered. Firms have started to practice generative engine optimization (GEO): changing content so that answer engines mention the brand, cite its pages, and describe it well. Marketing research has a detailed account of search engine optimization and paid search, but it does not explain how a brand gets into a machine-written answer or what being there does to choice. We define answer presence as a construct with four dimensions: inclusion, prominence, framing, and attribution. We then argue that a brand can reach an answer by two routes. In the parametric pathway the model draws on what it learned in training. In the retrieval pathway it reads documents fetched when the question is asked. The two routes reward different investments. Using research on consideration sets, search costs, persuasion knowledge, and brand signaling, together with the first computer science studies of GEO, we derive nine propositions. Among them: large brands have a wider lead in answers drawn from memory than in answers drawn from retrieval; third-party coverage matters more for evaluative questions than for factual ones; and answer presence can raise choice without raising visits to the brand's website. We close with implications for managers and regulators and a research agenda.

Key Words: Generative engine optimization, AI search, Large language models, Brand visibility, Consideration sets, Search marketing.

I. INTRODUCTION

Since the late 1990s, a brand that wanted to be found online had to rank. The consumer typed a query, read a list of links, and opened a few. Marketing research on this setting is extensive. It has measured how profitable each ad position is (Agarwal et al., 2011), estimated the causal effect of rank on search and purchase (Narayanan & Kalyanam, 2015; Ursu, 2018), analyzed the auctions that price sponsored links (Edelman et al., 2007), and modeled search engine optimization (SEO) as a strategy that interacts with paid search (Berman & Katona, 2013). Across this work, firms compete for positions on a list, and the consumer decides which result to open.

ChatGPT, Perplexity, Gemini, and the AI summaries that Google shows above its results reply to a query in prose. Ask one of them for the best running shoes for flat feet and you get a paragraph that names a few models and says why each one fits, sometimes with sources. Aggarwal et al. (2024) call these systems generative engines. A large share of users read the answer and go no further. Pew Research Center tracked the browsing of 900 US adults and found that people clicked a conventional search result on 8% of Google visits that included an AI summary and on 15% of visits that did not. They clicked a link inside the summary on about 1% of visits (Chapekis & Lieb, 2025). When the decision is made on the answer page, the brands the answer names become the options, and a brand it leaves out may never be weighed.

Aggarwal et al. (2024) named the response generative engine optimization (GEO) and defined it as improving a source's visibility inside the responses of generative engines. In their experiments, simple edits to a source document, such as adding statistics, quotations, or references, raised its visibility in answers by as much as 40%. Later work found that hidden text on a product page can move a product to the top of a model's recommendations (Kumar & Lakkaraju, 2024) and that crafted pages can lead production engines to promote one product while running down its competitors (Nestaas et al., 2024). One commercial tool that tracks brand mentions in AI answers, Ranqo, supplied the data for the largest brand-level audit so far (Kumar, 2026).

Marketing research on AI has covered strategy (Huang & Rust, 2021), the consumer's experience of AI (Puntoni et al., 2021), how the discipline may change (Davenport et al., 2020), the use of language models as research instruments (Goli & Singh, 2024; Li et al., 2024), and the relationships that form when both firms and customers use AI agents (Pattabhiramaiah et al., 2026). We have found no marketing study that explains how a brand comes to be named in a machine-written answer, which brands that process favors, or how the mention affects the person reading it. The computer science studies of GEO measure visibility and stop there. They do not link it to consideration or choice, and they do not ask what happens to answers once many brands optimize at once.

Our first contribution is a construct. We define answer presence as the extent and character of a brand's appearance in generative answers to the questions that matter in its category.

It has four dimensions: inclusion, prominence, framing, and attribution. They need separate measurement because they move separately. A brand can appear in most answers and be described unfavorably in them, or have its pages cited often while competitors get the recommendation. A single visibility score cannot show either pattern.

The second is a distinction between two pathways into an answer. In the parametric pathway the model answers from what it stored in its weights during training. In the retrieval pathway it fetches documents when the question arrives and writes from them (Lewis et al., 2020). Kandpal et al. (2023) and Mallen et al. (2023) found that language models recall facts about often-mentioned entities much better than facts about rarely mentioned ones, and Mallen et al. found that retrieval closes much of that gap. Put in marketing terms, the parametric pathway works like accumulated brand salience and favors large brands. The retrieval pathway works more like search, and smaller brands can enter through it, so the two call for different marketing investments.

The third is a set of nine propositions linking the antecedents of answer presence to consumer and market outcomes. They draw on consideration-set research (Hauser & Wernerfelt, 1990; Nedungadi, 1990), the economics of search (Nelson, 1970; Stigler, 1961), the persuasion knowledge model (Friestad & Wright, 1994), and brand signaling (Erdem & Swait, 1998). Two of the predictions matter most for practice, in our view. Answer presence can raise a brand's share of choices without raising its web traffic, so traffic-based reports will understate what GEO does. And the payoff from manipulative tactics should shrink as rivals copy them, while the payoff from building third-party coverage should hold up.

II. LITERATURE REVIEW

A. What research on search marketing established

The starting point of search marketing research is that position on a results page changes behavior. Ghose et al. (2014) found that a product's rank on a search engine affects both whether people click on it and whether they buy. Agarwal et al. (2011) found that higher ad positions attract more clicks while conversion rates do not drop in proportion, so the top slot is not always the most profitable one. Narayanan and Kalyanam (2015) used a regression discontinuity design to estimate causal position effects in search advertising and found that their size differs across advertisers and keywords. Ursu (2018) used random rankings on a travel website and showed that rank changed which hotels people looked at but did not change purchase once a hotel had been looked at, so in her data rank affected choice through attention alone.

SEO is the unpaid route to a better position. Berman and Katona (2013) modeled it as a costly investment that improves a site's organic rank and interacts with the bidding for sponsored links. The investment was aimed at a list, and the consumer still read the list and chose from it.

B. How generative engines answer

A generative engine writes its answer in prose. It can draw on two kinds of knowledge to do so, and the model in this article is built on the difference between them.

The first is parametric memory. In training, a language model absorbs statistical patterns from a very large body of text, and part of what it absorbs behaves like factual knowledge. Petroni et al. (2019) showed that pretrained language models can answer factual questions with no external text to consult. This memory is uneven. Kandpal et al. (2023) found that whether a model could answer a question depended on how many training documents mentioned the entities in it. Mallen et al. (2023) found that models recalled facts about popular entities well and facts about less popular ones badly, and that larger models improved far more on the popular ones.

The second is retrieval. A retrieval-augmented system collects documents relevant to the question when it is asked and generates its answer from them (Lewis et al., 2020). Perplexity and Google's AI summaries work this way, and general assistants do the same when they browse the web. Retrieval changes what can appear in an answer. In Mallen et al.'s (2023) tests, retrieval improved accuracy most on questions about less popular entities, which the models could not answer from memory.

Neither pathway is reliable. Language models produce fluent false statements, usually called hallucinations (Ji et al., 2023), and retrieval reduces the problem without removing it. Liu et al. (2023) audited four generative search engines and found that only 51.5% of generated sentences were fully supported by the citations given, and only 74.5% of citations supported the sentence they were attached to. Venkit et al. (2024) found inaccurate citations and overconfident answers in three commercial answer engines.

C. What the first GEO studies show

Aggarwal et al. (2024) built a benchmark of 10,000 queries and tested nine ways of rewriting a source. Adding quotations, statistics, or references to credible sources raised the source's visibility in answers by up to 40%, and by up to 37% on Perplexity. Keyword stuffing, an old SEO tactic, brought little or no gain, and a more persuasive tone brought no significant gain. The gains depended on where a source started. With the cite-sources method, visibility rose by 115% for the source ranked fifth in the underlying search results and fell by 30% for the source ranked first. The authors concluded that GEO helped lower-ranked websites most.

Three studies tested deliberate manipulation. Kumar and Lakkaraju (2024) inserted an optimized text sequence into a product page within a catalog of fictitious coffee machines and made a seldom-recommended machine more likely to be named the top pick. Pfrommer et al. (2024) injected text into product websites, changed the order in which conversational search engines recommended the products, and found that the attacks carried over to Perplexity. Nestaas et al. (2024) ran what they call preference manipulation attacks against Bing and Perplexity. They describe the result as a prisoner's dilemma: each seller gains by attacking, and if every seller attacks, answers get worse for all users.

Two large audits disagree about sources. Chen et al. (2025) compared AI search services with Google across product categories and languages. The AI services relied heavily on earned media, meaning third-party and editorial sources, and much less than Google on brand-owned and social content. The authors also report a bias toward large brands. Kumar (2026) analyzed more than 100,000 responses about more than 100 brands and found that about 78% of cited sources were corporate websites. The same study found a steep gradient by brand size. On a first tracking run, global household names appeared in 73% of relevant answers, mid-market brands in 44%, and niche brands in 11%. Whether a brand was described positively or negatively changed about 6.7 times as often across runs as whether it was mentioned. Both studies are preprints. The model below proposes one way to reconcile their findings on owned and earned sources, and Table I compares GEO with SEO.

TABLE I. HOW GENERATIVE ENGINE OPTIMIZATION DIFFERS FROM SEARCH ENGINE OPTIMIZATION

Feature	SEO	GEO
What the consumer sees	A ranked list of links	A written answer that names a few options, sometimes with citations
Unit of competition	Position on the results page	Inclusion, prominence, framing, and attribution in the answer
Where the brand competes	The search index	Model memory (parametric pathway) and documents fetched at query time (retrieval pathway)
Main levers	Query keywords, site structure, links from other sites	Content that is easy to extract and verify; third-party coverage built up over time
Keyword tactics	Matching query keywords is a standard practice	Keyword stuffing gives little or no gain (Aggarwal et al., 2024)
Link to traffic	Higher rank brings more clicks (Ghose et al., 2014)	Users click a link inside AI summaries on about 1% of visits (Chapekis & Lieb, 2025)
Variation	One ranked list for a query at a given time and place	Answers vary across runs, engines, and wording (Chen et al., 2025; Kumar, 2026)

Note. SEO = search engine optimization; GEO = generative engine optimization.

D. Theoretical foundations

1) Search costs and the short list

Stigler (1961) modeled consumer search as costly: a consumer keeps searching until the expected gain from one more inquiry falls below its cost. Nelson (1970) separated search goods, whose quality can be judged before purchase, from experience goods, whose quality is learned only by using them. Generative engines lower the cost of search a great deal. In two experiments by Spatharioti et al. (2025), participants using LLM-based search finished decision tasks faster, ran fewer but more complex queries, and found the experience more satisfying than participants using conventional search.

Cheaper search ought to widen consideration, yet an answer engine hands the consumer a short list. The consumer saves effort and in return accepts the engine's selection. Search theory predicts that the trade is most attractive when the selection is good enough to make further search look pointless. In Spatharioti et al.'s (2025) experiments, people using LLM-based search decided as accurately as people using conventional search when the model was right, and relied on it too much when it was wrong. If consumers accept the engine's selection whether or not it is accurate, brands have reason to compete for a place in it.

2) Consideration without evaluation

Nedungadi (1990) distinguished two ways marketing can change choice. One changes how consumers evaluate a brand. The other changes whether the brand is retrieved at the moment of choice, and it works even when evaluations do not move. Generative answers operate through the second. A consumer who reads an answer naming brands A, B, and C picks among those three. Brand D may be liked just as much, and it still loses because it is not on the list. Inclusion has value through this mechanism even when the answer contains no persuasion at all.

3) Persuasion knowledge and synthesized answers

According to the persuasion knowledge model, consumers learn to recognize persuasion attempts, and once they recognize one they discount the message and change how they respond (Friestad & Wright, 1994). Recognition depends on cues. Wojdynski and Evans (2016) found that many readers did not recognize native advertising as advertising, and that the placement and wording of the disclosure changed how many did.

A generative answer has very few cues. The engine blends many sources into its own voice and does not mark which statements came from material written to influence it. A shopper who would discount a claim in a sponsored link or on a seller's page may accept the same claim after the engine has absorbed and restated it. People also often give algorithmic advice more weight than the same advice from a person (Logg et al., 2019), although they lose confidence in an algorithm quickly once they see it err (Dietvorst et al., 2015). Taken together, these findings suggest that persuasion carried inside a neutral-sounding answer will often go unnoticed and be believed.

4) Signals and manipulation

Erdem and Swait (1998) argued that brands signal quality when buyers have less information than sellers, and that a signal is credible when it is costly to fake. Third-party coverage meets that test, since a firm cannot easily buy a favorable review in a respected outlet. Content a brand writes about itself is cheap, so it tells buyers little.

Research on online reviews shows firms faking signals when faking pays. Mayzlin et al. (2014) found more signs of promotional reviews among hotels whose ownership made faking cheaper and more profitable. Luca and Zervas (2016) found that restaurants on Yelp with weak or falling reputations were more likely to commit review fraud. The Federal Trade Commission (2024) has since banned buying and selling fake consumer reviews and testimonials. GEO adds a new incentive of the same kind, since text written for machines to read can carry claims that no human visitor ever sees.

5) Engines as platforms

Answer engines are platforms. They serve consumers who ask questions and, more and more, advertisers who want to be in the answers. Two-sided platforms set prices and rules to keep both sides participating (Rochet & Tirole, 2003). A platform that also sells its own products may favor them, and whether it should be allowed to is contested (Hagiu et al., 2022). The engine decides which sources to trust, how to treat optimized content, and whether to sell placement inside answers, and each of those decisions changes what GEO earns. The model therefore treats the engine as a strategic actor.

III. METHODOLOGY

The article is conceptual in MacInnis's (2011) sense: it proposes a construct, places it among existing ones, and derives propositions for empirical work to test. We collected no data. The argument rests on three bodies of research. Marketing and economics research on search, rankings, consideration, and persuasion supplies the mechanisms. Computer science research on how language models store and retrieve knowledge describes what engines can say. A small, quickly growing set of studies measures GEO directly. Several of those studies are preprints without peer review, and two of them disagree about which sources engines prefer. We mark preprints when we use them, and we use them only to motivate propositions.

The scope is answer engines that consumers use for product and service decisions. That includes general assistants such as ChatGPT and Claude, search-native engines such as Perplexity, and AI summaries inside conventional search. AI agents that buy on the consumer's behalf, without the consumer reading an answer, are outside the main model. We return to them under boundary conditions, because an agent that buys from whichever brands an engine names converts answer presence straight into sales.

IV. RESULTS: A MODEL OF ANSWER PRESENCE

This study is conceptual, so its results are a construct and a model with nine testable propositions.

A. Definition and dimensions of answer presence

Answer presence is the extent and character of a brand's appearance in the answers generative engines give to the queries relevant to its category. The definition covers a population of queries because one answer tells you little. Engines are stochastic: the same question can get a different answer on a second run, on another engine, or after a small change of wording (Chen et al., 2025; Kumar, 2026). Answer presence is therefore a distribution over queries, runs, and engines, and measuring it requires repeated sampling.

It has four dimensions. Inclusion is whether the answer names the brand. Prominence is how much of the answer the brand takes up and where it sits, for instance whether it is named first and recommended or mentioned last as an alternative. The position-adjusted word count used by Aggarwal et al. (2024) measures something close to prominence. Framing is what the answer says about the brand, including its valence and the attributes it attaches to the brand. Attribution is whether the answer cites or links to the brand's own content.

The four dimensions can diverge. An answer can cite a brand's website for a technical specification and recommend a competitor. A brand can appear in nearly every answer in its category and be called the expensive choice in all of them. The dimensions also differ in stability. Kumar (2026) found framing far less stable across runs than inclusion. A firm that tracks inclusion alone will miss a framing problem. A firm that tracks citations alone may confuse being used as a source with being recommended.

B. Relation to existing constructs

Answer presence overlaps in part with three established constructs. The nearest is brand salience, the propensity of a brand to come to mind in buying situations (Romaniuk & Sharp, 2004). Salience lives in consumers' memories. In the parametric pathway, answer presence lives in a model's memory, and that memory is built from what people and organizations have written about the brand. The two should be correlated, because much of that writing reflects the same market position that makes a brand salient to shoppers. They can still diverge. A store's private label may be bought every week and written about rarely, which would make it salient for its shoppers and nearly absent from a model's memory.

The consideration set is the second. Consumers choose among a few brands they consider, and the size of that set depends on how costly it is to evaluate more options (Hauser & Wernerfelt, 1990). Nedungadi (1990) showed that making a brand easier to recall can change choice while evaluations of the brand stay the same. An answer that names three brands does that recall for the consumer, so answer inclusion feeds the consideration set directly.

Share of voice, a brand's share of the communication in its category, is the third. Answer presence differs from it in two ways: no brand can buy space in an organic answer, and a third party, the engine, writes the text. In this respect answer presence is closer to earned media. Stephen and Galak (2012) found that both traditional and social earned media raised sales in a microlending marketplace. Unlike a journalist, though, an engine writes a separate piece for each reader and can be swayed by material written for it.

C. Overview of the model

Fig. 1 shows the model. Answer presence depends on three groups of antecedents. Brand inputs are the volume and character of third-party coverage and the form of the brand's

own content. Engine design covers whether the engine answers from memory or from retrieval and which sources it prefers. Query type is whether the consumer asks for an evaluation or a fact. Answer presence then shapes consumer response: consideration and choice on one side, visits to the brand's site on the other. Persuasion knowledge and the consumer's willingness to search beyond the answer moderate that link. Competition among optimizing brands and the engine's countermeasures change the written record that later answers draw on.

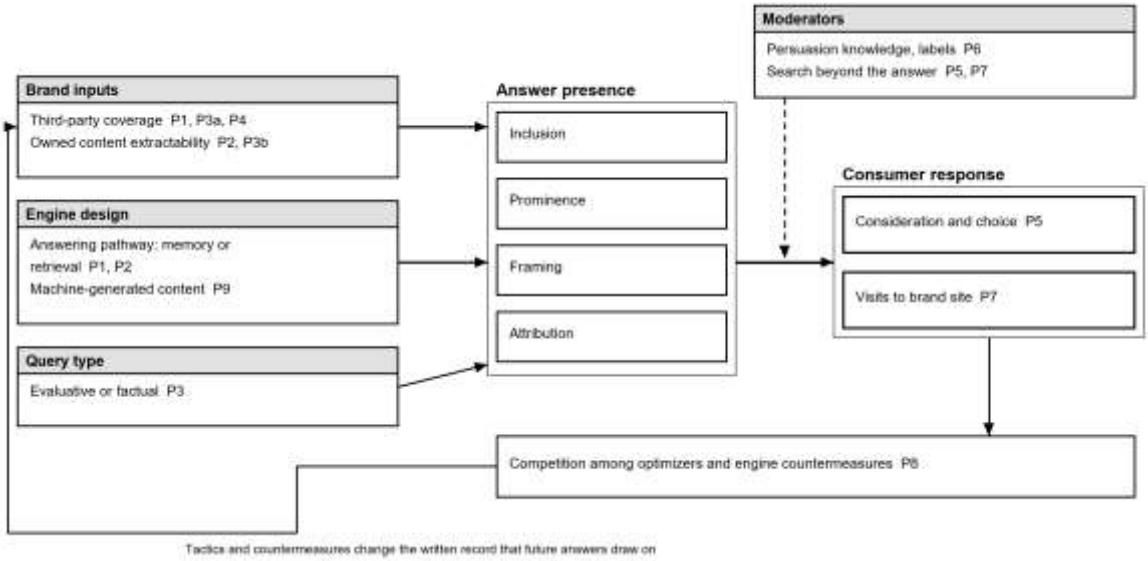


Fig. 1. A model of answer presence in generative engines. Proposition numbers show where each proposition acts; Table II gives the predicted effects. The dashed arrow marks moderators of the link between answer presence and consumer response. The feedback line shows that tactics and countermeasures change the written record later answers draw on.

D. Two pathways into the answer

1) The parametric pathway

An engine answering without retrieval can name only brands it learned about in training, and it learned about each in proportion to how much had been written about it. Kandpal et al. (2023) found that a model's accuracy on a question rose with the number of training documents related to the entities in the question. Mallen et al. (2023) found accuracy falling steeply for less popular entities and larger models improving mainly on popular ones. Brand names are entities like any other. A brand covered in many reviews, news stories, forum threads, and comparison pages will be recalled, and a brand with a thin written record will not.

The process resembles the way brand salience builds in consumers' memories through repeated exposure (Keller, 1993; Romaniuk & Sharp, 2004), with one difference. A model's memory is fixed when training ends and changes only when a new model is trained. Investment in the written record therefore pays off late and in jumps. The pathway also widens existing size differences, because large brands generate more text. Chen et al. (2025) report a big-brand bias in AI search, and Kumar's (2026) inclusion rates of 73%, 44%, and 11% for global, mid-market, and niche brands fit that account.

2) The retrieval pathway

When an engine retrieves documents, a brand can enter the answer through those documents even if the model knows little about it. Mallen et al. (2023) found that retrieval helped most with exactly the entities memory handled worst. Retrieval also opens the answer to content optimization. The gains Aggarwal et al. (2024) measured came from editing the documents an engine reads, and they were largest for low-ranked sources while the top-ranked source lost ground. Our interpretation is that engines favor sources that are easy to quote and look credible. A writer can add those features to any page with modest effort, whereas the inbound links and domain history behind conventional rank take years to accumulate. If so, the two pathways spread advantage differently.

> *P1: The gap in answer inclusion between brands with large and small written records is wider when engines answer from parametric memory than when they answer from retrieved documents.*

> *P2: In the retrieval pathway, content that is easy to extract and verify (specific statistics, attributed quotations, and references to credible sources) increases a source's prominence in answers, and the increase is larger for sources that rank low in the engine's underlying retrieval than for sources that rank high.*

For managers, P1 and P2 point small brands toward the retrieval pathway, where content work can show results within weeks, and toward engines and query types that trigger retrieval. Large brands hold a lead in model memory that they never paid for directly, and that lead is weakest where retrieval is used. If more answers come to rely on retrieval, as they do whenever a general assistant browses the web, P1 also predicts that the brand-size gradient Kumar (2026) observed will become flatter.

E. Source type and query type

Chen et al. (2025) found that AI search favors earned media, while Kumar (2026) found that most citations go to corporate websites. We propose that both can be right, because the source an engine relies on depends on what it is asked. Some queries ask for a judgment: which running shoe suits flat feet, which project management tool fits a small agency. Others ask for a fact about a named product: the battery capacity of a particular phone, or whether a software plan includes single sign-on.

Signaling theory predicts that these two kinds of query draw on different sources. For a factual question, the brand's own page is the authoritative record, and an engine trying to be accurate should cite it. For an evaluative question, the brand's page is an interested party. Its claims cost nothing to make and, in Erdem and Swait's (1998) terms, carry little signal. An engine tuned to give answers users find credible should lean on third-party sources, which are more expensive to fake. If the two audits sampled different mixes of queries, their results do not conflict.

> *P3: Query type moderates which sources raise answer presence. (a) For evaluative and comparative queries, third-party coverage raises inclusion and favorable framing more than an equal volume of brand-owned content. (b) For factual queries about a named brand or product, brand-owned content is the main source of attribution.*

If P3 holds, owned content and earned coverage do different jobs. Owned content earns citations, which help accuracy and attribution, and earned coverage earns recommendations. A firm that judges GEO by its citation count will put too much into owned content.

F. The stability of framing

Inclusion depends on whether the engine recalls or retrieves the brand. Framing depends on what the engine then says, a later step with more room to vary. A model trained on mixed coverage of a brand can produce a favorable description on one run and an unfavorable one on the next, and a small change in the wording of a question can change which attributes the engine treats as relevant. Kumar (2026) reported that framing changed about 6.7 times as often as inclusion across runs.

Framing should be most stable when the written record about a brand is consistent. If third parties keep describing a brand in terms of one attribute, such as durability or low price, the engine has little contrary material and should describe it the same way across runs. This mirrors the role Keller (1993) gives to strong and unique associations in brand equity, applied to the record a model learns from.

> *P4: Across repeated runs, engines, and paraphrases of a query, a brand's framing varies more than its inclusion. The variance of framing is lower for brands whose third-party coverage consistently associates them with the same attributes.*

For measurement, P4 implies that one audit run says little about framing and that a change in framing between two audits may be noise. Estimating framing takes many runs per query.

G. From answer to choice

1) Consideration and choice

Consideration-set research predicts that inclusion affects choice even with evaluations held constant (Nedungadi, 1990). A generative answer is an unusually strong recall cue: it arrives at the moment of decision and lists only a few options. The effect should be larger when the consumer does not look further. Pew's data show that most people who see an AI summary do not click through (Chapekis & Lieb, 2025), and Spatharioti et al. (2025) found that LLM-based search led people to run fewer queries. For a consumer who stops at the answer, the brands it names are the whole consideration set.

> *P5: Inclusion in a generative answer raises a brand's probability of being considered and chosen, even when consumers' evaluations of the brand are held constant. The effect is stronger when consumers do not search beyond the answer.*

An experiment can test P5 by varying which brands an answer names, keeping what it says about each brand fixed, and measuring consideration, choice, and brand attitude.

2) Persuasion knowledge

A brand claim that reaches a consumer through a generative answer arrives in the engine's voice. The reader cannot tell whether a sentence such as "the X model is widely regarded as the most durable" came from an independent review or from a page written for engines to absorb. The persuasion knowledge model predicts that consumers discount persuasion they notice and do not discount persuasion they miss (Friestad & Wright, 1994). Wojdynski and Evans (2016) found that disclosures can restore recognition, depending on where they appear and how they are worded.

> *P6: Consumers activate less persuasion knowledge in response to brand claims embedded in a generative answer than in response to the same claims in a sponsored listing or on the brand's own page. Labeling sponsored or brand-supplied material inside the answer increases activation and reduces the claims' effect on brand evaluations.*

Under P6, GEO may persuade more per exposure than conventional advertising does. The advantage depends on consumers not knowing what was optimized, so it could disappear once engines label such material, and it raises regulatory questions taken up in the discussion.

3) Visibility and traffic

In conventional search, rank and traffic rise together because higher positions get more clicks (Ghose et al., 2014). SEO budgets rest on that link: firms pay for rank to get visits and judge SEO by the visits and the sales that follow. Generative answers weaken the link. If the answer settles the question, there is little reason to click, and Chapekis and Lieb (2025) found clicks on links inside AI summaries on about 1% of visits that showed one.

A brand can therefore benefit from answer presence without receiving visits. The consumer reads the recommendation, closes the engine, and buys later from a retailer, a marketplace,

or a store, and nothing in the brand's web analytics connects that sale to the answer.

> *P7: The association between a brand's answer presence and the traffic its website receives is weaker than the association between its organic search rank and that traffic. As a result, traffic-based measures understate the effect of answer presence on choice, and the understatement is largest for brands sold mainly through intermediaries.*

We consider P7 the proposition with the largest consequences for practice. A firm that judges GEO by referral traffic will conclude it does not pay and will underinvest. Answer presence itself, and its relation to brand-level outcomes, is a better guide.

H. Competition among optimizers

1) Manipulation and its decline

The propositions so far describe one brand optimizing against an engine that does not change. In real categories many brands optimize at once and the engine reacts. Two kinds of GEO tactic should be separated. Manipulative tactics, such as hidden text sequences (Kumar & Lakkaraju, 2024) and injected instructions (Nestaas et al., 2024; Pfrommer et al., 2024), exploit weaknesses in how engines read pages, and they leave the product and its documentation unchanged. Coverage-building tactics, such as earning reviews and press coverage, add to the written record in ways that also inform human readers.

Two forces erode the returns to manipulation. One is competition: Nestaas et al. (2024) showed that when every seller attacks, each seller's gain shrinks while answers deteriorate. The other is the engine. An engine that loses users to manipulated answers has every reason to filter manipulative content. Aggarwal et al.'s (2024) finding that keyword stuffing does little in generative engines shows one old tactic that has already lost its value. Coverage-building tactics also face competition, but an engine has no motive to filter them, since they make its answers better, so their returns should fall more slowly.

> *P8: The gains in answer presence from manipulative tactics decline as more competitors in a category adopt them and as engines deploy countermeasures. The gains from building third-party coverage decline less over the same period.*

2) Homogenization

Another feedback loop runs through the text itself. Brands that write for engines increasingly write with language models. Dai et al. (2024) found that neural retrieval models rank text produced by language models above human-written text of similar relevance, a pattern they call source bias, and their analysis suggests that generated text is easier to match because its meaning is more concentrated. If engines favor generated text, brands that use it gain presence, more brands adopt it, and the text an engine reads about a category converges. Answers assembled from convergent text should describe brands in convergent terms. Sharma et al. (2024) found a related effect among users: people searching with LLM-powered conversational systems queried in a more one-sided way than people using conventional search.

> *P9: As the share of machine-generated content in a category's written record grows, the attributes that answers associate with different brands in that category become more similar.*

P9 bears on brand differentiation. When engines describe the brands in a category in the same terms, their choice among those brands will turn on something other than distinctive associations: price, availability, or the engine's own commercial arrangements.

Table II lists the propositions.

TABLE II. SUMMARY OF PROPOSITIONS

No.	Focus	Predicted effect	Main basis
P1	Answering pathway	The inclusion gap between brands with large and small written records is wider for answers drawn from memory than for answers drawn from retrieval	Kandpal et al. (2023); Mallen et al. (2023)
P2	Content extractability	Statistics, quotations, and citations raise prominence in the retrieval pathway, more for low-ranked sources	Aggarwal et al. (2024)
P3a	Evaluative queries	Third-party coverage raises inclusion and favorable framing more than owned content	Erdem & Swait (1998); Chen et al. (2025)
P3b	Factual queries	Owned content is the main source of attribution	Kumar (2026)
P4	Framing stability	Framing varies more than inclusion, less so when third-party coverage is consistent	Keller (1993); Kumar (2026)
P5	Consideration and choice	Inclusion raises consideration and choice with evaluations held constant, more so when consumers do not search further	Nedungadi (1990); Spatharioti et al. (2025)
P6	Persuasion knowledge	Claims inside answers activate less persuasion knowledge; labels restore it	Friestad & Wright (1994); Wojdynski & Evans (2016)
P7	Traffic	Answer presence is weakly linked to site visits; traffic measures understate its effect on choice	Ghose et al. (2014); Chapekis & Lieb (2025)
P8	Competition	Gains from manipulation decline with adoption and engine countermeasures; gains from coverage decline less	Nestaas et al. (2024); Aggarwal et al. (2024)
P9	Homogenization	A larger share of machine-generated content makes the attributes answers assign to brands more similar	Dai et al. (2024); Sharma et al. (2024)

Note. P1 and P2 concern the pathway into the answer, P3 and P4 the sources and stability of answer presence, P5 to P7 the consumer, and P8 and P9 competition.

I. Boundary conditions

The model should hold most strongly when four conditions are met. First, the purchase turns on attributes that can be described in text before buying. Following Nelson (1970), answers should matter more for search goods than for experience goods, where people rely more on their own trial and on advice from people they know. Second, the consumer has no strong prior preference. A loyal customer tends to name the brand in the query, which turns an evaluative query into a factual one and changes the sources P3 predicts. Third, the stakes are low enough that the consumer does not check the answer. For expensive or risky purchases, people may treat the answer as a starting point, which weakens P5 and P7.

The fourth condition concerns who reads the answer. AI agents that search and buy on a consumer's behalf are moving into use (Pattabhiramaiah et al., 2026). An agent that buys from whichever brands an engine names has no persuasion knowledge to activate, and answer presence becomes sales with no human step in between. In that setting P6 does not apply, and we expect P5 and P7 to hold more strongly.

J. Applying the model to three categories

Consider three hypothetical firms. The first is a small direct-to-consumer maker of running shoes, well reviewed by its own customers, with little press coverage and a short history. A consumer asks a general assistant, with browsing off, which shoes are best for flat feet. The answer comes from memory, and P1 predicts that this brand will not be named, since the model learned little about it. Put the same question to a search-native engine and the answer comes from retrieved pages, where the brand can appear if its pages or reviews are fetched. P2 predicts that pages with specific, checkable information, such as arch support stated in millimeters, a named podiatrist quoted, or a link to an independent lab test, will raise its prominence. The query is evaluative, so P3a predicts that coverage on running websites and in independent reviews will help the brand more than its own pages. Its best investment is earning that coverage. Its best measure is how often it appears in answers on retrieval-based engines, because referral visits will understate the effect (P7).

The second, a software firm selling project management tools to small businesses, faces two kinds of question. Some are factual: does the basic plan include time tracking, and what does it cost per user. Others are evaluative: which tool suits a ten-person design agency. P3 predicts that the firm's documentation and pricing pages will be the main cited sources for factual questions, so those pages need to be correct and easy to extract. For evaluative questions the firm depends on comparison articles and reviews on software marketplaces. P4 predicts that its framing in evaluative answers will vary from run to run unless that coverage describes it consistently. If some answers call it the simple option and others the feature-heavy one, its position is unstable, and the firm will only see this if it tracks framing over many runs.

The third is a regional hotel competing with chain hotels in the same city. Travelers ask engines where to stay, an evaluative question that rests on reviews. P3a predicts that the number of reviews and the consistency of what they say will drive inclusion and framing. This is also where the incentive to fake third-party content is strongest. Mayzlin et al. (2014) found more signs of promotional positive reviews for independent hotels, for which faking is cheaper, than for hotels in large chains, and Luca and Zervas (2016) found that businesses with weak reputations were more likely to commit review fraud. GEO raises the payoff to such fraud, since a planted review can now shape the answer a traveler reads as well as the review page. P8 predicts that engines will learn to discount such content, and the Federal Trade Commission's (2024) rule makes it unlawful. For this hotel, the reliable route into answers is slow: service good enough to earn consistent, genuine reviews.

V. DISCUSSION

A. Theoretical contributions

The first theoretical contribution is the answer presence construct. In ranked search, visibility was a single ordered position. Presence in an answer has four dimensions that behave differently and need different evidence. Separating them explains why a brand can be cited often and recommended rarely, and why its framing can swing while its inclusion stays steady. The construct also connects to brand salience, the consideration set, and earned media without reducing to any of them.

The second is the two-pathway account. Marketing research has mostly treated large language models as one technology. The computer science distinction between parametric and retrieval answering (Lewis et al., 2020; Mallen et al., 2023) has marketing consequences: it predicts where size advantages are largest (P1), where content work pays (P2), and how the brand-size gradient should change as engines move toward retrieval. It also shows how GEO differs from SEO. In the parametric pathway, a brand competes for a place in the model's memory, which is built from everything written about it and updated only when the model is retrained, a form of competition that has no counterpart in SEO.

The third is the link from GEO to choice and competition. The computer science studies take visibility as the outcome. Here it is an input to consideration and choice (P5), moderated by persuasion knowledge (P6) and partly invisible in traffic data (P7). GEO is also treated as a game among brands and against the engine (P8, P9), with predictions about which tactics will last.

B. Implications for managers

The propositions suggest five things for marketing managers. They are untested and should be treated as hypotheses.

First, measure answer presence directly, on all four dimensions, with repeated sampling. A firm can assemble a panel of the queries that matter in its category, covering evaluative and factual types, run each query many times on the engines its customers use, and code inclusion, prominence,

framing, and attribution. P4 implies that framing estimates from a handful of runs are unreliable.

Second, stop judging GEO by referral traffic. Under P7, answer presence affects choice through channels web analytics do not see. Survey measures of consideration, branded search volume, and sales across all channels are better guides. A firm can also estimate GEO's effect with controlled tests that vary GEO effort across markets or product lines.

Third, divide effort by query type. Owned pages should be accurate and easy to extract, since they are the main source for factual queries (P3b) and since the features Aggarwal et al. (2024) tested raise prominence in retrieval (P2). For evaluative queries a firm needs third-party coverage (P3a), such as reviews in respected outlets and genuine customer reviews. That work belongs to public relations and customer experience more than to the web team.

Fourth, small brands should concentrate on retrieval. P1 implies that they cannot close the memory gap quickly, and P2 implies that they gain more than large brands from content work in the retrieval pathway. Large brands should not count on their memory advantage lasting as engines retrieve more.

Fifth, avoid manipulative tactics. P8 predicts that their returns will fall, and the downside is large: an engine that detects manipulation can exclude a brand's content, and the tactics may break consumer protection law.

For engine providers, the model implies that answer quality depends on how optimized content is treated. Engines that reward extractable, verifiable content and discount manipulation should give better answers. Engines that sell placement inside answers will face the trade-off Rochet and Tirole (2003) describe for two-sided platforms, between advertiser revenue and the participation of users who want neutral answers.

C. Implications for public policy

The model raises three policy questions. The first is disclosure. P6 predicts that consumers do not recognize optimized or sponsored material inside a synthesized answer. If engines sell placement in answers, existing rules on advertising disclosure should apply, and Wojdynski and Evans's (2016) findings on native advertising can inform how labels are designed. Organic optimization is harder to regulate because nobody pays for the placement. Hidden text written to steer engines, though, has no purpose for human readers, and we see no case for tolerating it.

The second is fake third-party content. P3 predicts that third-party coverage drives recommendations, which creates an incentive to fake it with invented reviews and paid articles presented as independent. The Federal Trade Commission's (2024) rule on consumer reviews and testimonials already bans fake reviews. Regulators should consider whether content written for engines falls within it.

The third is concentration. P1 predicts that answers drawn from memory favor large brands. In categories where engines answer mostly from memory, they may entrench incumbents more than ranked search did. Competition authorities could monitor how answer presence is distributed across brands of different sizes.

D. Research agenda

Table III sets out a research agenda built on the model. Four priorities stand out to us.

The first is measurement. Each dimension of answer presence needs a validated measure with known reliability across runs and engines. Li et al. (2024) and Goli and Singh (2024) show how marketing researchers can test the validity of language model outputs, and their methods can be adapted to answers about brands.

P1 predicts steeper inclusion gradients where engines answer from memory, which researchers can check by running the same queries on engines and modes that differ in retrieval use. An engine adding or removing web browsing creates a natural experiment.

On the consumer side, P5 to P7 call for experiments that vary the content of answers and measure consideration, choice, persuasion knowledge, and clicks, together with field data linking answer presence to brand-level sales.

P8 and P9 need longitudinal data on GEO tactics and engine countermeasures. Archives of answers, of the kind commercial tracking tools already collect, would let researchers follow answer presence as a category becomes more heavily optimized.

TABLE III. RESEARCH AGENDA

Component	Research questions	Suggested methods
Measurement	How reliably can inclusion, prominence, framing, and attribution be measured across runs, engines, and paraphrases? How many runs give a stable estimate?	Repeated query audits; validity checks against human coding (Li et al., 2024)
Pathways (P1, P2)	Is the brand-size gradient in inclusion steeper for answers drawn from memory? Which content features raise prominence in retrieval?	Comparisons across engines and browsing modes; controlled content experiments
Sources and stability (P3, P4)	Do evaluative and factual queries draw on different sources? Does consistent coverage stabilize framing?	Query-panel audits coded by source type; longitudinal tracking
Consumer response (P5 to P7)	Does inclusion change choice with evaluations held constant? Do consumers recognize optimized content? How much of the effect bypasses web traffic?	Experiments with controlled answers; disclosure experiments; linking audit data to brand sales
Competition (P8, P9)	How fast do manipulative tactics lose value? Does machine-generated content make brands sound alike in answers?	Archived answers over time; text-similarity analysis of answers

Component	Research questions	Suggested methods
Policy and welfare	Do answers concentrate attention on large brands more than ranked search did? Which disclosure formats work inside answers?	Concentration measures by category; disclosure experiments

Note. Proposition numbers refer to Table II.

VI. CONCLUSION

A. Limitations

The article has five main limitations. It is conceptual, and none of the propositions has been tested. The GEO evidence is young: several studies we rely on are preprints, two of them conflict, and engines change faster than journals publish. P3 is our attempt to reconcile the conflicting audits, and it may be wrong.

The model treats engines as one class. They differ in how they retrieve and which sources they trust, and those differences may matter as much as the two pathways. Chen et al. (2025) found that AI search services differed significantly in the domains they drew on and in how sensitive they were to phrasing.

The model also says little about the query. Consumers write queries, and their wording shapes the answer. Sharma et al. (2024) found that people using conversational search wrote more one-sided queries than people using conventional search. A fuller model would treat query wording as something brand marketing can influence as well as an input to the engine.

Finally, the model takes the engine's business model as given. Answer engines are adding advertising and agents that buy, and both will change what the engine wants and therefore what GEO earns. Some of the propositions may not survive those changes.

B. Priorities for testing

We would test P3 first. It is the proposition we are least sure of, because it rests on a reconciliation of two preprint audits that disagree. It also settles a budget question, whether firms should spend on their own pages or on earning coverage elsewhere. A single audit that codes queries as evaluative or factual and records the source type behind each citation would confirm it or overturn it. P7 is a close second. If answer presence moves sales without moving traffic, GEO reports built on referral traffic measure the wrong thing, and firms that rely on them will misjudge what their content work achieves.

DECLARATIONS

Data availability. No new data were created or analyzed in this study. The article is conceptual and draws only on published sources listed in the references.

Ethics approval. This conceptual study involved no human participants, animals, or personal data, so ethics approval was not required.

Competing interests. The authors have no competing interests to declare that are relevant to the content of this article.

Funding. The authors did not receive support from any organization for the submitted work.

Use of generative AI. The authors used Claude (Anthropic), a large language model, to help search for and verify sources, draft and edit text, prepare the figure, revise the prose for clarity, and format the manuscript. The authors reviewed and revised all content, checked every reference against its DOI or publisher record, and take full responsibility for the article. Claude is not an author.

REFERENCES

Agarwal, A., Hosanagar, K., & Smith, M. D. (2011). Location, location, location: An analysis of profitability of position in online advertising markets. *Journal of Marketing Research, 48*(6), 1057–1073. https://doi.org/10.1509/jmr.08.0468

Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K., & Deshpande, A. (2024). GEO: Generative engine optimization. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (pp. 5–16). Association for Computing Machinery. https://doi.org/10.1145/3637528.3671900

Berman, R., & Katona, Z. (2013). The role of search engine optimization in search marketing. *Marketing Science, 32*(4), 644–651. https://doi.org/10.1287/mksc.2013.0783

Chapekis, A., & Lieb, A. (2025). *Google users are less likely to click on links when an AI summary appears in the results* [Pew Research Center short read]. 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/

Chen, M., Wang, X., Chen, K., & Koudas, N. (2025). *Generative engine optimization: How to dominate AI search* (arXiv:2509.08919). arXiv. https://doi.org/10.48550/arXiv.2509.08919

Dai, S., Zhou, Y., Pang, L., Liu, W., Hu, X., Liu, Y., Zhang, X., Wang, G., & Xu, J. (2024). Neural retrievers are biased towards LLM-generated content. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (pp. 526–537). Association for Computing Machinery. https://doi.org/10.1145/3637528.3671882

Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science, 48*(1), 24–42. https://doi.org/10.1007/s11747-019-00696-0

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General, 144*(1), 114–126. https://doi.org/10.1037/xge0000033

Edelman, B., Ostrovsky, M., & Schwarz, M. (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review, 97*(1), 242–259. https://doi.org/10.1257/aer.97.1.242

Erdem, T., & Swait, J. (1998). Brand equity as a signaling phenomenon. *Journal of Consumer Psychology, 7*(2), 131–157. https://doi.org/10.1207/s15327663jcp0702_02

Federal Trade Commission. (2024). *Trade regulation rule on the use of consumer reviews and testimonials* [Final rule, 89 Fed. Reg. 68034]. https://www.federalregister.gov/documents/2024/08/22/2024-18519/trade-regulation-rule-on-the-use-of-consumer-reviews-and-testimonials

Friestad, M., & Wright, P. (1994). The persuasion knowledge model: How people cope with persuasion attempts. *Journal of Consumer Research, 21*(1), 1–31. https://doi.org/10.1086/209380

Ghose, A., Ipeirotis, P. G., & Li, B. (2014). Examining the impact of ranking on consumer behavior and search engine revenue.

Management Science, 60(7), 1632–1654. https://doi.org/10.1287/mnsc.2013.1828

Goli, A., & Singh, A. (2024). Frontiers: Can large language models capture human preferences? *Marketing Science, 43*(4), 709–722. https://doi.org/10.1287/mksc.2023.0306

Hagiu, A., Teh, T.-H., & Wright, J. (2022). Should platforms be allowed to sell on their own marketplaces? *RAND Journal of Economics, 53*(2), 297–327. https://doi.org/10.1111/1756-2171.12408

Hauser, J. R., & Wernerfelt, B. (1990). An evaluation cost model of consideration sets. *Journal of Consumer Research, 16*(4), 393–408. https://doi.org/10.1086/209225

Huang, M.-H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. *Journal of the Academy of Marketing Science, 49*(1), 30–50. https://doi.org/10.1007/s11747-020-00749-9

Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys, 55*(12), 1–38. https://doi.org/10.1145/3571730

Kandpal, N., Deng, H., Roberts, A., Wallace, E., & Raffel, C. (2023). Large language models struggle to learn long-tail knowledge. In *Proceedings of the 40th International Conference on Machine Learning* (Vol. 202, pp. 15696–15707). PMLR. https://proceedings.mlr.press/v202/kandpal23a.html

Keller, K. L. (1993). Conceptualizing, measuring, and managing customer-based brand equity. *Journal of Marketing, 57*(1), 1–22. https://doi.org/10.1177/002224299305700101

Kumar, A., & Lakkaraju, H. (2024). *Manipulating large language models to increase product visibility* (arXiv:2404.07981). arXiv. https://doi.org/10.48550/arXiv.2404.07981

Kumar, P. (2026). *Generative engine optimization at scale: Measuring brand visibility across AI search engines* (arXiv:2606.20065). arXiv. https://doi.org/10.48550/arXiv.2606.20065

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems* (Vol. 33, pp. 9459–9474). Curran Associates. https://doi.org/10.48550/arXiv.2005.11401

Li, P., Castelo, N., Katona, Z., & Sarvary, M. (2024). Frontiers: Determining the validity of large language models for automated perceptual analysis. *Marketing Science, 43*(2), 254–266. https://doi.org/10.1287/mksc.2023.0454

Liu, N. F., Zhang, T., & Liang, P. (2023). Evaluating verifiability in generative search engines. In *Findings of the Association for Computational Linguistics: EMNLP 2023* (pp. 7001–7025). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.findings-emnlp.467

Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes, 151*, 90–103. https://doi.org/10.1016/j.obhdp.2018.12.005

Luca, M., & Zervas, G. (2016). Fake it till you make it: Reputation, competition, and Yelp review fraud. *Management Science, 62*(12), 3412–3427. https://doi.org/10.1287/mnsc.2015.2304

MacInnis, D. J. (2011). A framework for conceptual contributions in marketing. *Journal of Marketing, 75*(4), 136–154. https://doi.org/10.1509/jmkg.75.4.136

Mallen, A., Asai, A., Zhong, V., Das, R., Khashabi, D., & Hajishirzi, H. (2023). When not to trust language models: Investigating effectiveness of parametric and non-parametric memories. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 9802–9822). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.acl-long.546

Mayzlin, D., Dover, Y., & Chevalier, J. (2014). Promotional reviews: An empirical investigation of online review manipulation. *American Economic Review, 104*(8), 2421–2455. https://doi.org/10.1257/aer.104.8.2421

Narayanan, S., & Kalyanam, K. (2015). Position effects in search advertising and their moderators: A regression discontinuity approach. *Marketing Science, 34*(3), 388–407. https://doi.org/10.1287/mksc.2014.0893

Nedungadi, P. (1990). Recall and consumer consideration sets: Influencing choice without altering brand evaluations. *Journal of Consumer Research, 17*(3), 263–276. https://doi.org/10.1086/208556

Nelson, P. (1970). Information and consumer behavior. *Journal of Political Economy, 78*(2), 311–329. https://doi.org/10.1086/259630

Nestaas, F., Debenedetti, E., & Tramèr, F. (2024). *Adversarial search engine optimization for large language models* (arXiv:2406.18382). arXiv. https://doi.org/10.48550/arXiv.2406.18382

Pattabhiramaiah, A., Mende, M., Gao, H., & Sridhar, S. (2026). Agentic AI and the new architecture of marketing relationships. *Journal of Marketing*. Advance online publication. https://doi.org/10.1177/00222429261481116

Petroni, F., Rocktäschel, T., Riedel, S., Lewis, P., Bakhtin, A., Wu, Y., & Miller, A. (2019). Language models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing* (pp. 2463–2473). Association for Computational Linguistics. https://doi.org/10.18653/v1/D19-1250

Pfrommer, S., Bai, Y., Gautam, T., & Sojoudi, S. (2024). *Ranking manipulation for conversational search engines* (arXiv:2406.03589). arXiv. https://doi.org/10.48550/arXiv.2406.03589

Puntoni, S., Reczek, R. W., Giesler, M., & Botti, S. (2021). Consumers and artificial intelligence: An experiential perspective. *Journal of Marketing, 85*(1), 131–151. https://doi.org/10.1177/0022242920953847

Rochet, J.-C., & Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association, 1*(4), 990–1029. https://doi.org/10.1162/154247603322493212

Romaniuk, J., & Sharp, B. (2004). Conceptualizing and measuring brand salience. *Marketing Theory, 4*(4), 327–342. https://doi.org/10.1177/1470593104047643

Sharma, N., Liao, Q. V., & Xiao, Z. (2024). Generative echo chamber? Effect of LLM-powered search systems on diverse information seeking. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (pp. 1–17). Association for Computing Machinery. https://doi.org/10.1145/3613904.3642459

Spatharioti, S. E., Rothschild, D., Goldstein, D. G., & Hofman, J. M. (2025). Effects of LLM-based search on decision making: Speed, accuracy, and overreliance. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (pp. 1–15). Association for Computing Machinery. https://doi.org/10.1145/3706598.3714082

Stephen, A. T., & Galak, J. (2012). The effects of traditional and social earned media on sales: A study of a microlending marketplace. *Journal of Marketing Research, 49*(5), 624–639. https://doi.org/10.1509/jmr.09.0401

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy, 69*(3), 213–225. https://doi.org/10.1086/258464

Ursu, R. M. (2018). The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science, 37*(4), 530–552. https://doi.org/10.1287/mksc.2017.1072

Venkit, P. N., Laban, P., Zhou, Y., Mao, Y., & Wu, C.-S. (2024). *Search engines in an AI era: The false promise of factual and verifiable source-cited responses* (arXiv:2410.22349). arXiv. https://doi.org/10.48550/arXiv.2410.22349

Wojdynski, B. W., & Evans, N. J. (2016). Going native: Effects of disclosure position and language on the recognition and evaluation of online native advertising. *Journal of Advertising, 45*(2), 157–168. https://doi.org/10.1080/00913367.2015.1115380