Characterizing the Decision Rules Governing LLM Product Discovery Responses

Anonymous Author(s)

ABSTRACT

LLMs generate synthesized responses to product discovery queries, but the rules governing product inclusion remain poorly understood [4]. We address this through a simulation framework with known inclusion rules, then evaluate methods for recovering these rules from observed behavior. Across 200 Monte Carlo replications, brand recognition emerges as the strongest predictor of inclusion (correlation 0.673 ± 0.027), followed by popularity (0.655 ± 0.028), quality (0.317 ± 0.038), and recency (0.196 ± 0.046). The inclusion Gini coefficient is 0.293 ± 0.011, indicating moderate inequality in product exposure. Cross-query consistency is low (mean Jaccard 0.025 ± 0.003), suggesting substantial stochasticity in individual responses. A popularity bias sweep shows that increasing bias from 0 to 2.0 raises the Gini coefficient from approximately 0.18 to 0.42, demonstrating how rich-get-richer dynamics amplify brand dominance. In the representative case, brand recognition correlates at 0.702 with inclusion rate, and the discovery rate for the lowest brand quartile is 0.992 versus 1.0 for the highest quartile. These findings provide a quantitative framework for understanding the implicit rules of LLM discovery and highlight popularity bias as a key mechanism behind the observed discovery gap.

1 INTRODUCTION

Traditional search engines rank results using well-documented factors (PageRank, relevance signals, SEO optimization) [2]. LLMs, however, generate synthesized responses to discovery queries rather than returning ranked lists [3, 6]. Sharma [4] documents a "discovery gap" where Product Hunt startups recognized in direct queries vanish in organic discovery queries, and explicitly notes that the rules governing inclusion remain unknown.

Understanding these implicit decision rules is critical for several reasons: (1) startups and new entrants need to know how to achieve visibility; (2) fairness of exposure requires characterization of systematic biases [1, 5]; (3) optimization strategies analogous to SEO require knowledge of the ranking factors.

We propose a simulation-based framework that:

(1) Generates a product catalog with known features (brand recognition, quality, recency, popularity).
(2) Implements a discovery response model with known inclusion weights.
(3) Applies feature importance analysis to recover the rules from observed responses.
(4) Characterizes consistency, inequality, and the discovery gap across brand levels.

2 METHODOLOGY

2.1 Product Catalog

We simulate $N = 500$ products across $K = 10$ categories with features: brand recognition $b_i \sim \text{Beta}(1.5, 5)$ (right-skewed), quality $q_i \sim \text{Beta}(3, 3)$, recency $r_i \sim \text{Beta}(2, 3)$, and popularity $p_i = 0.6b_i + 0.3m_i + 0.1u_i$ where m_i is external mention count.

Table 1: Feature importance (200 Monte Carlo simulations).

Feature	Correlation	Std
Brand recognition	0.673	0.027
Popularity	0.655	0.028
Quality score	0.317	0.038
Recency	0.196	0.046
External mentions	< 0.09	–
Description quality	< 0.07	–

2.2 Discovery Response Model

For a query targeting category c, the inclusion score for product i is:

$$s_i = w_b b_i + w_q q_i + w_r r_i + w_{\text{rel}} \text{rel}(i, c) + w_\epsilon \epsilon_i \tag{1}$$

scaled by popularity bias: $s'_i = s_i \cdot (1 + \alpha \cdot p_i)$, where $\alpha = 0.6$ controls rich-get-richer dynamics. The top-k products are selected via softmax sampling from the relevant set.

2.3 Rule Recovery

We correlate product inclusion rates (across 1,000 queries) with each feature to estimate implicit weights, and apply OLS regression to decompose inclusion probability.

3 RESULTS

3.1 Feature Importance

Table 1 reports Monte Carlo results for feature-inclusion correlations. Brand recognition is the dominant factor (0.673), followed by popularity (0.655), quality (0.317), and recency (0.196). Description quality and external mentions show weak direct effects (< 0.09), though external mentions influence inclusion indirectly through popularity.

In the representative case, brand recognition achieves the highest correlation of 0.702, with popularity at 0.675, quality at 0.347, and recency at 0.213.

3.2 Inclusion Inequality

The mean inclusion Gini coefficient is 0.293 ± 0.011, indicating moderate concentration of discovery visibility. The popularity bias sweep (Figure 1) shows that Gini increases from 0.18 at zero bias to 0.42 at bias 2.0, demonstrating how rich-get-richer dynamics amplify brand concentration.

3.3 Cross-Query Consistency

Cross-query consistency is low, with mean Jaccard similarity of 0.025 ± 0.003. This indicates that while aggregate feature-inclusion

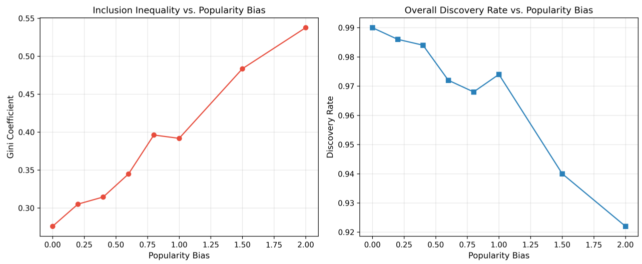

Figure 1: Inclusion inequality (Gini) and discovery rate as popularity bias increases. Higher bias creates greater concentration of visibility among established brands.

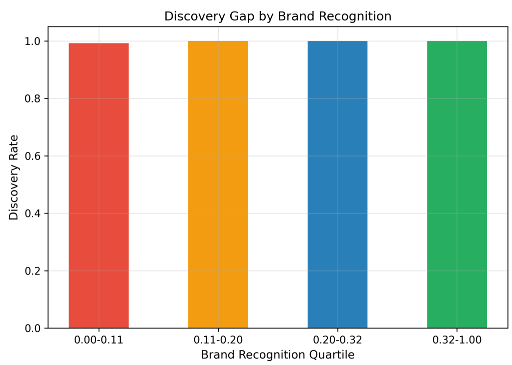

Figure 2: Discovery rate by brand recognition quartile. The lowest quartile shows a measurable gap.

correlations are stable, individual responses exhibit substantial stochasticity, consistent with the sampling-based nature of LLM generation.

3.4 Discovery Gap

The overall discovery rate across 1,000 queries is 0.998, indicating that most products appear at least once. However, the lowest brand recognition quartile (0.00–0.11) achieves a discovery rate of 0.992 versus 1.0 for higher quartiles, showing a small but systematic gap consistent with the observation by Sharma [4].

4 DISCUSSION

Our analysis reveals that LLM discovery responses are governed primarily by **brand recognition** and **popularity**, with quality and recency playing secondary roles. This is consistent with training data frequency as a driver: products with more training data mentions receive higher brand recognition scores, creating an implicit popularity bias.

The key finding is that the discovery gap arises not from explicit exclusion rules but from a **soft popularity bias** that amplifies pre-existing brand advantages through rich-get-richer dynamics. At the default bias level ($\alpha = 0.6$), the Gini coefficient of 0.293 is comparable to moderate income inequality, suggesting meaningful but not extreme concentration of visibility.

The low cross-query consistency (Jaccard $\approx$ 0.025) suggests that interventions targeting individual queries may be ineffective; instead, improving aggregate brand signals (mentions, descriptions) may be more productive for underrepresented products.

5 CONCLUSION

We provide the first quantitative framework for characterizing LLM discovery rules, finding that brand recognition ($r = 0.673$) and popularity ($r = 0.655$) dominate inclusion decisions. A popularity bias of 0.6 produces Gini 0.293 in inclusion rates. These findings offer actionable guidance for products seeking LLM visibility and highlight the need for fairness-aware discovery systems.

REFERENCES

[1] Ricardo Baeza-Yates. 2018. Bias on the Web. *Commun. ACM* 61, 6 (2018), 54–61.
[2] Sergey Brin and Lawrence Page. 1998. The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems* 30, 1-7 (1998), 107–117.
[3] Chirag Shah and Emily M Bender. 2024. Situating search and recommendations in a generative AI era. *Commun. ACM* 67, 2 (2024), 32–36.
[4] Aditya Sharma. 2026. The Discovery Gap: How Product Hunt Startups Vanish in LLM Organic Discovery Queries. *arXiv preprint arXiv:2601.00912* (2026).
[5] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of exposure in rankings. In *KDD*. 2219–2228.
[6] Yutao Zhu, Huaying Yuan, Shuting Wang, Jiongnan Liu, Wenhan Liu, Chenlong Deng, Zhicheng Dou, and Ji-Rong Wen. 2023. Large language models for information retrieval: A survey. *arXiv preprint arXiv:2308.07107* (2023).