



GROWFER RESEARCH · EXPLORATORY COMPANION ANALYSIS

On-site pricing **not** **surfaced as an** **answer.**

The engines found the sites and often saw the prices. So why did they answer “no”?

Public research edition

Revision 2

Exploratory



PUBLIC RESEARCH EDITION

This document is a public, redacted edition of an internal Growfer Research report. Business identities, domains, URLs, row-level excerpts, and internal repository references have been removed or anonymized. Aggregate findings, methodology, statistical summaries, limitations, and research conclusions are preserved from the cited report revision (Revision 2, October 7, 2026). Pattern examples are anonymized cases retaining only the attributes necessary for the methodological point.

STATUS: EXPLORATORY COMPANION ANALYSIS

This document examines one slice of the study's frozen record with mechanical proxies defined below. Its figures did **not** receive the bidirectional audit treatment applied to the main study's audited figure set; they are reproducible from the frozen study record by deterministic, read-only analysis, and every proxy and caveat is labeled where it applies. Conclusions here are interpretations of the primary study's data, not additions to its audited result set.

01 – PURPOSE

Purpose and relationship to the main study

The main report established that businesses with pricing verified on their own websites were nonetheless answered "no pricing visible" about half the time (Gemini surfaced 52.2%, Perplexity 48.4% of the 647-business on-site cohort). This companion asks the follow-up: **for the on-site businesses the engines did not surface, was there a pattern to how the pricing was presented on-site?**

The slice examined: the 647 on-site-cohort businesses, specifically the **183 Gemini "no" rows**, the **334 Perplexity "no" rows**, and the **143 rows both engines answered "no."**

This is the mirror image of the main report's crawler-invisible finding. There: the study's crawl-and-audit process found no pricing, yet engines surfaced it (123 Gemini / 98 Perplexity businesses). Here: the audit verified pricing on the business's own site, yet engines declined to surface it. The two error surfaces point at the same conclusion from opposite directions — what a crawler can verify and what an engine will answer are different quantities.

02 – DENSITY

Presentation-density gradient

Each on-site business carries a count of distinct on-site pricing signals from the crawl (quoted-price occurrences and pricing structures on its rendered pages). Visibility by signal density:

PRICING SIGNALS	N	GEMINI YES	GEMINI NO	GEMINI REFUSED	PERPLEXITY YES	PERPLEXITY NO
1–2 (thin)	94	31 (33%)	39	24	27 (29%)	67
3–5	94	35 (37%)	39	20	31 (33%)	63
6–10	95	43 (45%)	38	14	38 (40%)	57
>10 (dense)	364	229 (63%)	67	68	217 (60%)	147

Roughly a **2× visibility gradient** from thinnest to densest presentation — unusually monotonic, and it appears on both engines rather than depending on one model's behavior.

Interpretation caution. Density may act partly as a proxy for something broader: a site with many pricing signals probably has a dedicated pricing architecture, explicit service names, repeated contextual language, and structured offerings. **Density may be measuring semantic completeness as much as repetition**, and this analysis cannot separate the two (Section 9 records the retest that would). The gradient should not be translated into "publish the same price more often and AI will find you."

03 — PLACEMENT

Pricing placement

Where the pricing sat on the site, by a mechanical location proxy (a dollar-amount pattern in a rendered page's visible-text head, or a pricing-indicative URL path; caveats in Section 7):

LOCATION	N	GEMINI YES	PERPLEXITY YES
Homepage (in visible-text head)	0	—	—
Subpage only, dollar amounts visible in text head	338	194 (57%)	179 (53%)
Pricing-indicative URL, no dollar amounts in text head	231	102 (44%)	97 (42%)
No pricing page detected by the proxy	78	42 (54%)	37 (47%)

Zero of 647 businesses showed dollar amounts in the rendered homepage's text head; **the pricing signals captured by this analysis overwhelmingly appeared on subpages**. Pages with **visible dollar amounts** outperformed pricing-looking URLs without them by ~13pp (Gemini) / ~11pp (Perplexity).

Misread guard. Do **not** summarize this as "put prices on the homepage." The homepage-zero result is partly an artifact of the text-window proxy (Section 7): the homepage's opening rendered text is typically navigation and marketing copy regardless of where prices sit deeper in the page. What the subpage-versus-URL-only gradient supports is that *visible dollar amounts on the relevant pages* associate with better engine outcomes — not any homepage-placement prescription, which this data does not test.

04 — DECOMPOSITION

Retrieval versus semantic responsiveness

The decomposition that reframes the whole phenomenon — what the engines' own answers reveal about *why* they said no (full answer texts, mechanical regex classification; caveats in Section 7):

	GEMINI (183 NO-ROWS)	PERPLEXITY (334 NO-ROWS)
Cited the business's OWN domain anyway	154 (84%)	303 (91%)
"No" answer text contains price figures anyway	86 (47%)	170 (51%)
Plan / membership / bundle language in answer	54 (30%)	56 (17%)
Mentions a different metro than queried	7 (4%)	0 (0%)
True retrieval miss (no price seen AND own domain not cited)	21 (11%)	18 (5%)

Rows both engines answered "no": **143** — at least one engine cited the business's own domain in **133 (93%)**, and at least one "no" answer still contained price figures in **85 (59%)**. Both-no rows concentrate in consultation-pricing categories: dentists (37), pet groomers (27), restaurants (22), property management (13), medical spas (12).

The dominant failure mode is not "the engine couldn't find the site." It is "the engine found the site, saw prices, and judged — usually reasonably — that they did not answer the specific question asked." The branded query asks for a standalone price for one category-typical service; the sites often publish bundles, plans, ranges, product prices, training fees, or other-location pricing instead. The fixed verdict rule classifies those as "no" — measurement-correct, but the bucket's composition matters enormously for interpretation. **Most failures to surface an answer were not simple discovery failures.**

PART 02

Patterns in the negative cases

Anonymous pattern classes, the data-quality findings they exposed, and the mechanical proxies behind every figure.



Growfer

05 — PATTERNS

Anonymous pattern examples

Pattern classes observed in the record (counts = candidate rows matching each pattern; examples anonymized, retaining only the attributes that carry the lesson):

A — Dense subpage pricing, both engines "no," own domain cited (29 candidates).

- A **facial-treatments business classified under massage spas** (21 pricing signals; audited quotes for its actual services): both engines answered no, both citing its own site — one explicitly: it "does not appear to offer 60-minute massage services," the site's services being facial treatments. The engines were right; the category label was wrong.
- A **property-management company** (45 signals; its pages carry abundant per-unit asking rents): both engines answered no, both citing its own domain. The frozen query for this category — "*What does {businessName} charge for monthly rent for a listed unit?*" — is ambiguous for a marketplace-style business: one engine read it as asking the company's management fee (a standard fee or percentage of rent charged to owners), the other as a standard rent the company itself publishes; on either reading, both judged the abundant per-unit asking rents non-responsive to what the company "charges."

B — Thin presentation, both no (33 candidates).

- A **yoga studio** (2 signals; audited quotes were four-figure teacher-*training* prices): both no, both citing its own site — training prices visible, no monthly membership price published. Right answer, wrong offering.
- A **dental practice** published a **\$400 annual in-house plan** covering cleanings, exams, and X-rays — and no standalone cleaning or crown price. Both engines found and quoted the plan, and both correctly declined to treat it as the standalone price of the queried procedures. Bundle-versus-standalone strictness.

C — Range/consultation-style pricing, both no (7 candidates).

- A **salon whose website published pricing for its location in a different metro** than the study cell queried: both engines answered no — one noting explicitly that the site's prices were for the other city's location.
- A **salon whose on-site shop sold products at published prices** (the audited quotes were product prices): both no — one engine: the shop section "only features hair products for sale"; no service price list existed.

D — Divergent engines (96 candidates where Gemini said yes, Perplexity no, and Perplexity cited the site itself).

- A salon with standalone service prices on its own services pages: Gemini answered **yes** citing only the business's Square-hosted site (never its own domain); Perplexity answered **no** despite citing the business's own domain. Same underlying truth, different retrieval paths, opposite verdicts.

E — Promotional/range-format pricing (14 candidates).

- A salon publishing range-format pricing ("From \$..." tiers and custom-quote language): one engine answered yes; the other no — it found a single styled-service price but no price for the queried haircut or color services. One engine's answer also noted the business had closed or rebranded: freshness beyond the study window.

F — Contrast: the visible profile (108 candidates where dense subpage pricing produced both-yes).

- A salon with 103 pricing signals and a dedicated service-price-list page carrying standalone prices for individual services: both engines answered **yes**, both citing the business's own price list and quoting its figures verbatim. **The profile that answers well: dense, standalone, dedicated pricing page.** One contrast case is a profile, not a formula — the Section 2 caution applies.

06 — DATA QUALITY

Data-quality findings

Three findings surfaced by reading the negative cases — valuable for what they say about measurement, not embarrassing to the study:

1. **Market-label mismatches.** A small share of "no" answers referenced a different metro than queried (~4% of Gemini no-rows; none of Perplexity's): the sampled market label and the website's actual content disagreed (a business's site priced its location in another city). Cohort classification is unaffected (pricing location, not geography, defined cohorts), but the fixed query's market token can point an engine at the wrong city's offering.
2. **Product-versus-service pricing in broad signal detection.** The crawl's pricing-signal detection counted quoted prices wherever they appeared — including on-site shops selling products. Some on-site-classified businesses publish prices that are real but are not prices *for the services the query asks about*. A simplistic pricing detector would overcount "published service pricing."

3. **Category-label mismatch.** Some businesses' sampled categories did not match their actual offerings (the facial-treatments example above). When the fixed per-category service query meets a mislabeled business, a correct "no" is the only possible answer.

Implication for future study design: all three are neutralized by establishing the **verified offering first** — auditing what the business actually sells, at what prices, in which market — and querying *that*, instead of querying a category lexicon's guess. Section 9 makes this the core of the Study 2 design.

07 — CAVEATS

Caveats (mechanical proxies, labeled)

- **Location proxy:** dollar-pattern in a rendered page's *visible-text head* (a window onto the page's opening text, not the full page) or a pricing-indicative URL path. The homepage-zero result is partly a text-window artifact; the subpage-versus-URL-only gradient is not.
- **"Price figures in answer"** is a regex over the full answer text; it cannot guarantee the price belongs to the business (a third-party price in the same answer counts). Directional, strongest when combined with the own-domain citation rates.
- **Own-domain citation** uses a last-two-label domain approximation.
- **Pattern examples are class picks, not random samples** (first matching rows in frozen-population order; candidate counts shown per class).
- **Texture subset:** audited-quote texture analysis covers only the 77 on-site businesses with quoted-price audits; within it, promotional-language rows skew Perplexity-no (39% yes vs the 51% quoted-subset baseline) — directional at these sample sizes.
- This document's figures are **exploratory**: reproducible from the frozen record, but outside the main study's bidirectional figure audit.

SECTION 4 RECAP · MOST "NO" ANSWERS WEREN'T DISCOVERY FAILURES

84–91%

of "no" rows cited own domain

Gemini 84% · Perplexity 91%

~1 in 2

"no" answers quoted price figures anyway

47% Gemini · 51% Perplexity

5–11%

true retrieval misses

Perplexity 18 · Gemini 21

PART 03

From visibility to answerability

Where the failures sit in the pipeline — and what that means for the next study.

08 — SYNTHESIS

Answerability synthesis

DEFINITION

Answerability is the degree to which a business's public information can support a correct answer to a real customer question.

The failure pipeline for a branded pricing query:

retrieval → extraction → semantic matching → answer/verdict

The Section 4 decomposition locates the mass. **Retrieval failure is comparatively rare** (true misses 5–11%; the engines cited the business's own domain in 84–91% of no-rows). **Extraction can fail** — thin signals, subpage placement, no visible dollar amounts (Sections 2–3). But a large fraction of failures occur at **semantic matching**: prices were seen, often quoted, and reasonably judged non-responsive to the question actually asked. Training fees versus a membership question; an annual plan versus standalone procedure prices; per-unit asking rents judged non-responsive to what a property-management company itself "charges"; product prices versus service prices; another city's pricing versus the queried market; a mislabeled category versus the business's real offering. These are **semantic correspondence** problems, not presentation problems: no price-page architecture makes facial pricing responsive to a massage query, and one property-management example carried abundant pricing (45 signals) that both engines judged non-responsive to the queried service.

The both-no subset sharpens this into the formulation that generalizes beyond this cohort:

KEY INSIGHT

Presence in the information environment does not guarantee that the information environment can answer the question.

A business owner does not ultimately care whether an engine crawled a pricing page; they care whether someone asking "how much does X cost at Business Y?" gets a useful answer. The deeper model emerging from the two documents together: **retrieval is probabilistic, extraction is imperfect, semantic matching is selective, and the structure of the information environment changes the odds at every stage.** That is why repeated visibility measurements are noisy while something stable sits

underneath — the environment's quality and shape.

Terminology. "Invisible" is an operational shorthand this analysis makes semantically awkward: many of these businesses were seen, price-quoted, and rejected as non-responsive by the engines. Preferred downstream terms: **"not surfaced as an answer"** or **"failed the branded pricing query."**

09 — STUDY 2

Implications for Study 2

1. **Verified offering first.** Study 1's flow: *business* → *category grammar* → *fixed query* → *compare with the environment*. Study 2's flow: **business** → **audited offering** → **known price basket** → **targeted query** → **measured answer**. Capturing a structured basket of the business's actual prices (for its actual services, in its actual market) *before* engine measurement turns accuracy from a descriptive subset into a powered measurement — and neutralizes the Section 6 label/taxonomy findings by construction, since the query targets the verified offering rather than the lexicon's guess.
2. **Density-disentanglement retest.** Re-run the Section 2 gradient after excluding clearly non-responsive pricing: product-only pricing; training/course pricing where a membership was queried; unrelated listing prices; wrong-market pricing; wrong-category matches. **If something resembling 30% → 60% survives those exclusions, that is much stronger evidence that information richness itself increases AI answerability;** until then, density and semantic relevance remain entangled.

10 — THESIS

Relationship to the broader research thesis

This companion pushes in the same direction as the main report's interpretation, from the opposite error surface. The diagnostic question is not "did an AI find a price?" It is closer to:

THE DIAGNOSTIC QUESTION

Does a business's online information environment contain enough specific, consistent, semantically appropriate information for an AI system to answer an actual customer question?

That framing explains why repeated measurements of "visibility" can be noisy while leaving something stable underneath — the quality and shape of the information environment — and it makes the system's object of study considerably less dependent on any particular engine's current behavior. As a research interpretation (not a product-performance claim): **audit the information environment's answerability, not merely a visibility score.**

Growfer Research — public research edition of the exploratory companion analysis. The primary study report ("The Split-Pricing Study: Final Report," public research edition) carries the audited figure set; this document's figures are exploratory and labeled as such throughout.



CLOSING THOUGHT

Audit answerability, not merely a visibility score.

NEXT STEP

Read the primary study.

The Split-Pricing Study: Final Report carries the audited figure set this analysis builds on.

[Primary study report · Rev 6](#)

143

on-site businesses
both engines
answered “no”

93%

of those had own
domain cited