WORKING PAPER

The record you no longer write

Chris LaColla

31 July 2026

ABSTRACT

A travel business used to control the first description a customer read of it. That description now gets assembled, from material the operator wrote across several ownership cycles, from a news cycle the operator does not control, and from public forums that are being deliberately manipulated. The assembled answer is becoming the operating description of the business, and it is produced without the brand in the room. This paper sets out how the assembly works, why the property that wins it is the most internally consistent one, and why continuous publishing is the only input an operator reliably holds.

The hotel of record is now assembled

Model-generated summaries have started behaving like a new form of property record. They gather descriptions, reviews and language written across multiple ownership cycles and present the result as a single current account of the property.

That produces a strange outcome for anyone who has just spent money on the asset. Renovations, repositionings and brand changes do not replace earlier material. They sit alongside it, and the earlier material is frequently easier to find and easier to summarize, because it has had longer to accumulate links, reviews and repetition. The hotel that exists today competes for attention with earlier versions of itself.

The consequence runs straight into capital planning. Improving the physical product no longer guarantees that the improvement propagates into how the market understands the property. Return on a renovation now depends on a second thing, which is whether the change becomes legible upstream of the booking moment. Most renovation budgets contain no line for that.

This is a measurable version of a familiar effect. The billboard research established that visibility on a third-party surface raises a hotel's own-site bookings, and that most direct bookers have passed through an intermediary first. The assembled summary is the same mechanism with a shorter path: it is now often the only third-party surface the shopper sees.

The platform can override the brand outright

The sharper version of this is a summary that contradicts the brand at the top of its own branded search.

Searching for a major United States carrier on Instagram in early 2026 returned a Meta AI summary before any brand content, and the summary assembled recent coverage: weak demand on specific routes, a vote of no confidence in the chief executive, and management job cuts. That material had not traditionally appeared on social platforms at all. These surfaces were a brand's highlight reel. The news cycle is now pulled directly into the search flow, and it can override brand identity with negative sentiment at the exact moment of highest purchase intent, which is when someone searches for you by name.

Set aside the obvious difficulty of converting a shopper under a headline about leadership failure. The structural point is that the summary has become the primary context for the user, and the algorithm decided that a list of crises was the most relevant thing to show. When the platform controls the narrative on branded search, content strategy becomes a secondary instrument. This is a category of reputational risk with no established mechanism for brand control, and no service level behind it.

It is already on your own name

The Instagram case is not an outlier, and the search data puts a number on how routine it has become. Roughly 68% of United States Google searches now resolve without sending traffic to any site. Where an AI Overview appears, that figure rises to about 83%. And AI Overviews appear on roughly 89% of branded search results, which is to say: when someone searches for your property or your brand by name, there is close to a nine in ten chance an assembled answer sits above anything you wrote. Ahrefs measures the click cost at up to 58% off the position-one organic click-through rate where an Overview appears. Branded queries carry the one genuinely encouraging exception in that data, gaining about 19% in click-through, which fits the reading that someone searching your name has already decided and the Overview confirms them. That exception holds only while the summary is favorable. The carrier case above is what it looks like when the summary is not.

The other number worth holding is that brands cited inside an Overview earn around 35% more organic clicks than uncited competitors. Citation, not ranking, is the unit that pays now.

So the surface where a brand historically had the most control, its own name, is the surface where the assembled answer is most likely to appear.

What the models are eating

The grounding material matters as much as the assembly, and its quality is deteriorating in a specific and traceable way.

Every major model is heavily grounded in forum content, and Reddit in particular. There has been a noticeable recent shift on that platform toward astroturfing. Rather than transparently marketing, companies are using the forum-to-model connection to engineer their way into relevance. The pattern is consistent enough to recognize. An account claims its business was failing until it found one specifically named tool. Other accounts arrive to praise the same tool in oddly similar marketing language. More accounts arrive to plug adjacent products. A few humans call it out and report it, and the thread survives anyway because moderation is overwhelmed.

Those fabricated conversations become training and retrieval material, and the answers built on them inherit the fabrication.

The useful qualifier is that this concentrates in business technology forums, which are more tolerant of promotional content and less strictly moderated. Travel communities remain comparatively reliable for genuine human input. That is a real advantage for this sector and it is not guaranteed to persist, since the incentive to manipulate scales with how much commercial weight the assembled answer carries.

Consistency beats quality

The ranking logic that follows from all of this is uncomfortable for operators who compete on product. Discovery mediated by models does not advantage the best property. It advantages the most internally consistent one. A brand whose pricing, availability, positioning and merchandising drift independently will lose visibility regardless of how good the underlying product is, because an assembled answer can only be as coherent as the material available to assemble.

The platforms already work this way, which is worth noticing because it is the same shape. Expedia Group discloses that lodging ranking runs on a weighted mix of price competitiveness, availability, reviews, booking history, cancellation behavior and compensation when offers are similar. Several of those inputs are coupled: raise prices and velocity drops, tighten cancellation terms and conversion drops, open availability and cancellation risk rises. Operators optimize one variable at a time while the system evaluates the interaction. Visibility losses then feel sudden and unexplained, even though every individual decision looked defensible.

An assembled answer is a less legible version of the same thing. Nobody publishes its weights.

The corollary is that stating the obvious has become an infrastructure decision. When Oceania Cruises formalised an adults-only position, the segmentation itself was old news, since the brand had operated that way for years. What changed is that every system describing it now had the audience written down. Where a brand leaves that ambiguous, the model hedges and surfaces it in answers written for somebody else, which buys impressions in front of people the product was never for and the review scores that follow. Explicit audience definition and consistent language decide how these systems read a brand long before anyone buys tooling to influence them.

Credibility now defaults to disbelief

There is a second-order effect on everything an operator publishes to defend itself.

Consider a business testimonial section where the headshots are uniformly polished, the claims cite large specific numbers, and every logo is a placeholder. Those testimonials may be entirely truthful. There is room for plausible explanation: the operator may be protecting client confidentiality, or may simply not have replaced the template assets.

That ambiguity is the problem. When generated material can convincingly produce the appearance of truth, polished and generic social proof stops reading as evidence and starts reading as decoration. The default response to a claim that looks templated is skepticism, whether or not the claim is accurate. Social proof presented as a design element rather than a factual record now costs more in credibility than the effort it saved.

The practical rule is that specificity and verifiability are the only remaining signals. A named client with a checkable outcome carries weight that a stronger unnamed claim does not.

What an operator can actually do

Four things, in descending order of how much control the operator has over them.

Publish continuously. This is the only input that is reliably yours. Feeds reward accounts that publish consistently: Buffer's analysis of more than 100,000 accounts over 26 weeks found those posting at least weekly for the majority of that period earned roughly five times the engagement per post of inconsistent accounts. Assembled answers are built from whatever material exists, so a brand that goes quiet is described by its gaps and by other people's material.

Make the current version of the property easier to find than the previous one. After a renovation or a repositioning, the change has to be published, dated and repeated often enough to outweigh the accumulated older material. This is a budget line, and it belongs in the capital plan.

Hold pricing, availability, positioning and merchandising in agreement. Consistency is the ranking input you control most directly and the one most often broken by internal structure, since these four things usually sit with four different owners.

Make claims checkable. Name the client, cite the figure, link the source. In an environment primed for disbelief, a verifiable modest claim outperforms an unverifiable strong one.

Limitations

The Instagram observation is a single search on a single platform at a point in time, and these summaries are personalized and non-deterministic, so it illustrates the mechanism without measuring it. The astroturfing pattern comes from sustained observation of the platform, with no quantified audit behind it, and the claim that travel communities remain cleaner is an impression I have not tested systematically. Model grounding sources are not disclosed by any major provider, so the path from a forum thread to an assembled answer is inferred from behavior. The Buffer analysis measures engagement, and the step from engagement to commercial outcome is not established here. The zero-click and Overview-prevalence figures come from third-party measurement panels, not from Google, they are US-weighted, and they move quickly enough that any specific percentage here should be treated as a reading from mid-2026. The 89% branded-query figure in particular is a share of sampled queries and will vary by vertical; I have not seen it broken out for travel specifically. The consistency argument predicts that legibility beats quality in assembled discovery; the experiment that would test it, holding product constant while varying only cross-channel consistency, is one I have not run.

SOURCES

1. Meta AI summaries in Instagram search, observed February 2026
2. Buffer, analysis of over 100,000 accounts across 26 weeks
3. Expedia Group, disclosed ranking factors for lodging search
4. Anderson & Han, The Billboard Effect: Still Alive and Well, Cornell SHA (2017)
5. Chris LaColla, The rising cost of ambiguity in cruise branding, Cruise Weekly, 21 January 2026 https://cruiseweekly.com.au/news/the-rising-cost-of-ambiguity-in-cruise-branding/86322
6. Ahrefs, organic click-through rate under AI Overviews (2026)
7. Similarweb, zero-click share of US Google searches (2026) https://www.similarweb.com/blog/marketing/geo/zero-click-marketing/
8. Adobe Digital Insights, Q3 2026 AI Traffic Trends Report https://business.adobe.com/resources/sdk/.q3-ai-traffic-trends-report/q3-2026-ai-sourced-traffic-insights.pdf

HOW THIS WAS MADE

AI is part of how the work here gets made: finding sources, pulling figures out of long documents, drafting and restructuring. The arguments and the judgement are mine. Figures are checked against primary sources where a primary source exists, and the source is named next to the number so you can check it too.

That reduces errors and does not eliminate them. This is an exploratory surface and some of what is here will turn out to be wrong. If you find something, email **chris@lacoven.com** and I will correct it.

Read this paper with its figures at lacoven.com/writing/the-record-you-no-longer-write

LaCo Ventures LLC · 407 Lincoln Rd, Suite 6H PMB 1277, Miami Beach, FL 33139