LACOVEN

WORKING PAPER

Nobody wants the blank box

Chris LaColla

31 July 2026

ABSTRACT

Every major travel brand is now shipping a conversational planning interface, and the premise underneath all of them is that the traveler wants to type what they want into an empty field. That premise is being adopted without evidence, at a moment when the evidence available points the other way. Free text search has never converted particularly well in travel; the optimized booking flow, with its filtered attributes, structured choices and visual merchandising, won through two decades of testing against exactly this alternative. This paper argues that the blank box is the wrong interface for a considered travel purchase, that the context which would make it work accrues to the assistants while the operator starts from nothing, and that what an operator can build instead is a proposal.

The measurement now exists, and it says what the flow said

When I first made this argument the evidence was the accumulated design of the category. There is now a direct measurement, and it is worth stating before anything else because it cuts both ways. Adobe Digital Insights tracks traffic arriving at United States travel sites from assistant referrals. That traffic grew 194% year over year to May 2026, and 2,215% since Adobe began tracking it in October 2024. The volume argument is settled: this is a real channel and it is compounding.

The behavior of that traffic is where it gets interesting. Assistant-referred travel visitors are 21% more engaged than visitors from other sources, spend 70% longer per visit, and bounce 41% less. Adobe reads those patterns as more purposeful, higher-intent behavior, and that reading is reasonable. Someone who arrives after a conversation about their trip has already done the thinking.

They also convert 28% less.

Higher intent, longer sessions, deeper engagement, and a worse outcome. That combination is the whole argument in one line. The assistant is doing the specification work well. The surface it hands off to cannot use any of it, so a visitor who arrives knowing exactly what they want lands on a page that asks them where they want to go.

The honest counterpoint belongs here too. That conversion gap has narrowed by nearly 70% since October 2024. Something is closing it, whether that is the assistants getting better at qualifying, operators getting better at receiving the traffic, or the audience broadening past early adopters. If the gap closes entirely, the argument in this paper weakens considerably. I would watch that number more closely than any other in travel right now.

The context sits on the other side of the table

The value a model delivers in trip planning is roughly proportional to the context it holds on the traveler. This is the part of the design that determines everything downstream, and it is the part where the operator is structurally disadvantaged.

A general assistant accumulates behavioral history across every domain its user touches. It gets better at recommending because it has been watching for a long time and across a wide surface. A brand chat widget starts at zero on every session. The customer has to manually reconstruct their preferences, their travel patterns and their constraints before the conversation can produce anything past surface-level suggestion, and that reconstruction cost does not disappear because the interface looks modern. Hilton's AI Planner is a genuine signal that hotel brands are taking conversational AI seriously, and the commercial logic behind it is shakier than the engineering. A siloed planner on a brand domain cannot compete on context accumulation with the platforms whose entire business is context accumulation, so it will always feel less tailored than the thing the traveler already has open in another tab. If a demand-capture strategy depends on an interface that requires the traveler to do the context work, the leverage is on the wrong side of the equation.

What the optimized flow already learned

The booking funnel is one of the most heavily tested interfaces in consumer software. Filtered attributes, structured choices, visual browsing and curated result sets did not survive because designers preferred them. They survived because they beat the alternatives repeatedly, across operators, over many years, on conversion.

The mechanism is plain. A blank field hands the entire problem back to the user at the moment they are least equipped to specify it. Most travelers do not arrive knowing the vocabulary of the category, the tradeoffs that matter, or which of their preferences are even available to be expressed. A structured surface teaches the choice set while the user moves through it. A prompt assumes the choice set is already in their head.

This is why the same interface pattern keeps reappearing under new technology. Assistant products that begin as open conversation converge back toward cards, comparisons and constrained choices as soon as they are asked to close a transaction, because the transaction requires the user to select something specific and selection needs a visible set.

Logistics resolves faster than judgment

Generated itineraries expose the gap precisely. They resolve logistics well: sequence, timing, feasibility, the shape of a week. They produce plans that appear complete, organized and efficient.

What they rarely help with is the decision the traveler is actually stuck on, which is what kind of experience they are looking for. Hotels get slotted into a generic structure as interchangeable stops, and the consequence of choosing one over another never surfaces. The itinerary looks finished while the decision remains unresolved, and the gap shows up long before booking begins.

Tools that optimize for completion tend to avoid the harder work of interpretation, because completion is measurable and interpretation is not. That bias is worth naming, since it explains why so many planning products feel impressive in a demonstration and inert in use.

The tab problem, and the wrong conclusion drawn from it

The most cited evidence for a unified planning tool is that travelers keep many tabs open while researching. The observation is well supported. The conclusion drawn from it does not follow.

Cross-verification is the process. Travelers are suspicious of platform incentives, and comparing sources adds confidence to a purchase that is expensive and infrequent. A single source of truth asks the shopper to give up the one behavior that makes them comfortable committing, in exchange for convenience they did not ask for.

Planning is also idiosyncratic in a way that resists standardization. People process information, work through logistics and coordinate with the other people on the trip in vastly different ways. It is possible to aggregate some base of like-minded users. A universally appealing planning approach is a harder claim, and the infrequency of the purchase means there is little opportunity to build the habit that would make one stick.

Travel is important enough to people that they will tolerate scattered tabs and notes in three places. That tolerance is the ceiling on this category.

Recognition forms upstream of the comparison

There is a pricing observation that belongs here, because it says something about where the decision is really made.

In peak periods some hotels operate far above the market median without meaningfully testing elasticity. Revenue teams typically attribute this to amenities or service tier. The more likely driver is that travelers arrive already knowing what they want. Recognition of the property was established upstream, and by the time the shopper reaches a comparison surface the comparison is a formality.

If that is right, it relocates the commercial problem. The work that determines the outcome happens before the interface, in whatever formed the traveler's prior understanding. A blank box is downstream of the moment that mattered.

What the operator actually controls

Two things, and neither is a chat interface.

The first is internal consistency. Discovery mediated by models will not advantage the best product. It will advantage the most legible and most internally consistent one. Brands that let pricing, availability, positioning and merchandising drift independently lose visibility regardless of quality, because an assembled answer is only as coherent as the material it assembles.

The second is understanding the systems that already rank them. Expedia Group discloses that lodging ranking is driven by a weighted mix of price competitiveness, availability, reviews, booking history, cancellation behavior and compensation when offers are otherwise similar. Those inputs are coupled: raise prices and booking velocity drops, tighten cancellation terms and conversion drops, open availability and cancellation risk rises. Teams optimize one variable at a time while the platform evaluates the interaction, which is why visibility losses feel sudden and unexplained even when each decision looked reasonable alone.

What to build instead

The alternative to a blank box is a surface that arrives already specific.

An operator holds things a general assistant does not: live inventory, real availability, the actual fare, the photography, and a record of what this particular person engaged with. That is enough to render a proposal. The traveler's job becomes reacting to something concrete, which is the task the evidence says they perform well.

The personalization literature adds one constraint worth respecting. A meta-analysis of 53 experiments finds personalized advertising more effective than generic, with perceived relevance as the operative mechanism and personalization built on real user data considerably outperforming inferred personalization. The same body of work finds that covert personalization tends to outperform overt, and that high personalization levels in email can reduce purchase probability. So stated data should determine what the surface is, which offer leads and which alternates rank, and should not be announced back to the reader.

Limitations

The Adobe figures are channel-level observational data, not an experiment. Assistant-referred visitors differ from other visitors in ways beyond how they arrived, so the 28% conversion gap mixes the effect of the handoff with the composition of the audience, and neither Adobe nor anyone else has separated them. The narrowing of that gap has no published decomposition either, so my three candidate explanations are speculation. The broader claim that free text converts poorly in travel rests on the accumulated design of the category, and I am not aware of a published head-to-head experiment at scale. The Hilton assessment is a reading of the commercial logic from the outside; the product may hold context I cannot observe. The Expedia ranking inputs are disclosed as factors without published weights, so the coupling described is inference from the mechanics. The pricing observation comes from operator experience, and a controlled elasticity test is what would settle it. The personalization meta-analysis is retail-dominant, and its transfer to considered travel purchases is an assumption here that nobody has tested.

SOURCES

1. Hilton, AI Planner announcement, 2026
2. Expedia Group, disclosed ranking factors for lodging search
3. SiteMinder, Booking Lead Time: Hotel Trends https://www.siteminder.com/r/booking-lead-time/
4. Anderson & Han, The Billboard Effect: Still Alive and Well, Cornell SHA (2017)
5. Yeo, Chu & Li, How Persuasive Is Personalized Advertising? A Meta-Analytic Review, JAR 65(4) (2025)
6. Adobe Digital Insights, AI-sourced traffic to US travel sites, May 2026 https://business.adobe.com/blog/adobe-report-ai-traffic-travel-sites-surges-200-percent
7. 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
8. Skift, Adobe data on AI travel traffic across airlines, hotels and car rental, 17 June 2026 https://skift.com/2026/06/17/adobe-data-ai-travel-traffic-airlines-hotels-car/

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/nobody-wants-the-blank-box

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