



WORKING PAPER

The booking window is diverging, and both directions break the same model

Chris LaColla

31 July 2026



ABSTRACT

The interval between first interest and purchase is moving in opposite directions in different parts of travel at the same time. Hotel windows are compressing hard enough that a major operator now reports most of its bookings landing inside twelve days of arrival, while cruise windows are lengthening far enough that a majority of a year's sailings can be sold before the year begins. This paper argues that the two movements are the same problem seen from either end: both remove a different piece of the forecast that commercial teams price against, both get read as anomalies when they are the new baseline, and underneath both sits an identity horizon measured in days that neither side has priced.

Two directions at once

The averages conceal the divergence, so it is worth separating the verticals.

Hotels are compressing. Global average booking lead time on SiteMinder's platform is 36.4 days, inside a typical industry range of 20 to 60. That average is the last comfortable number in this paper. In March 2026 Accor's chief executive said 60% of the group's customers now book within twelve days of departure. A twelve-day modal window and a 36-day mean describe a distribution with a heavy front load, which is a different planning object from a distribution centered on five weeks.

Vacation rentals sit between, with severe seasonality. AirDNA's median United States lead time has stayed under 30 days since the pandemic. The median hides the shape: over half of July and August bookings are placed more than 30 days out, while half of January and February bookings land inside two weeks. The same portfolio is running two different businesses depending on the month.

Cruise is lengthening. Only 11% of cruise travelers book a year or more before sailing, but more than 40% of North American travel advisors report client windows stretching to nine or eighteen months. Carnival Corporation reported in March 2026 that 85% of its 2026 sailings were already booked, at record prices, largely against deposits collected months in advance.

Accor is the sharpest public statement of the compression, and it is consistent with the wider picture. STR's booking window indices show many markets running 10 to 20% shorter than their 2019 baseline, with urban and short-haul leisure compressing fastest. Industry reporting through 2026 puts last-minute reservations, meaning zero to six days out, at roughly 30 to 35% of all bookings against 15 to 20% before the pandemic. In major Chinese cities the last-minute share is reported as high as 65%. One industry, three distributions, and no shared planning assumption between them.

Compression removes the pacing signal

Revenue management is built on the premise that the shape of the booking curve to date predicts the shape of the rest of it. Pacing against a historical curve is what converts today's on-the-books position into a price for tomorrow.

Compression breaks that mechanism at the input. When the majority of demand arrives inside twelve days, the pricing curve for the eleven months preceding it is reading a period in which most of the eventual buyers had not yet entered the market. The curve stays accurate and stops being informative, and it goes quiet in a way that looks exactly like a soft market until the demand shows up.

My read is that travelers are responding to genuine uncertainty. They wait because waiting reduces the cost of being wrong, and nothing on the horizon makes the world feel more predictable. If that reading is right, the compression is durable and the models built on prior cycles will keep describing a demand environment that no longer exists.

The operational consequence is unglamorous. Yield management run as a periodic review function, with a weekly or biweekly cadence, is pricing demand it cannot see in time to act on. The teams that adapt to a twelve-day window are the ones that move to present-tense judgment, which costs more in attention than the models it replaces.

Commitment removes the ability to reprice

The cruise position looks like the opposite problem and is the same one.

A forward book sold 85% through at record prices is a genuinely strong commercial position, and it is also a position with almost no remaining optionality. Fare adjustments only reach inventory that is still for sale. When costs move materially, whether through fuel or through a change in the operating environment, an operator holding a committed book cannot reprice its way out, because the revenue is already contracted and the guest has already planned around it.

The risk that deserves watching in that configuration is cancellation behavior as sailing dates approach. Deposits collected a year in advance represent a commitment made under one set of household economic conditions and honored under another. A book that looks stable in March is a forecast about consumer confidence in October, and it is being carried on the balance sheet as though it were a fact. So compression removes the ability to forecast and commitment removes the ability to respond. Both are failures of the same underlying capability, which is the ability to act on demand inside the window that demand actually occupies.

The identity horizon underneath both

There is a third clock, and almost nobody puts it on the same axis as the other two.

WebKit's current tracking prevention policy blocks all third-party cookies without exception and deletes all script-writable storage, meaning JavaScript cookies, LocalStorage, IndexedDB and service-worker caches, after seven days of browser use without interaction with the site. When the arrival URL carries click identifiers, which is the shape of every tracked ad click, JavaScript-set cookies on that landing page are capped at 24 hours.

So the technical memory of a traveler runs somewhere between one day and seven. The decision runs between twelve days and eighteen months.

Two corrections are worth making because both circulate in stale form. The frequently repeated claim of a flat seven-day expiry describes the 2019 ITP 2.1 rule; the current rule keys off interaction: storage survives as long as the visitor keeps coming back. At multi-month windows that changes nothing, since a shopper researching in February and depositing in August will go seven days without a visit many times over. And Google's April 2025 decision not to deprecate third-party cookies is being read as a reprieve. It is not one. Third-party cookies never solved cross-device identity, and the display retargeting they enable decays over exactly the interval that matters: a large randomized experiment puts the four-week return-rate lift at 14.6%, with a third of the first week's effect arriving on day one and effectiveness falling from there.

The compression case and the commitment case both assume the operator can recognize a returning shopper. At twelve days that assumption survives only with repeat visits. At eighteen months it does not survive at all.

What this changes for planning

Three things follow, and none of them is a technology purchase.

Stop reporting a single lead time. A mean of 36 days across a distribution with a twelve-day mode and a long tail is a number that no decision should be made against. The planning object is the distribution, by season and by segment.

Price the optionality you are giving up. A deposit taken twelve months out is revenue and a forfeited call option on the fare. Operators book the first and rarely quantify the second, which is why a heavily committed book gets reported as unambiguously good news.

Treat the identity horizon as a commercial constraint. It gets filed as a tracking detail, which is why nobody prices it. The operator-held email address is the only identifier that survives both a multi-month window and a device change. Anything that depends on the browser remembering the shopper is, at these windows, a plan with no memory in it.

How I would know I am wrong

Long-window claims are where this class of argument most commonly overclaims, and the correction from experimental economics is severe enough that it should change how anyone in this category talks.

Blake, Nosko and Tadelis ran a randomized shutdown of eBay's brand-keyword search. Ninety-nine and a half percent of paid clicks were substituted by organic traffic. Conventional attribution implied a return above 4,100%. The experiment measured negative 63%. Last-click methods credit the channel with purchases by frequent buyers who were going to transact anyway, and the longer the window, the more of that substitution there is to hide.

So the falsification criteria for the argument above are these. If operators running present-tense pricing against compressed windows do not outperform those running historical pacing, measured on matched properties, the compression claim is decorative. If cancellation rates on long-dated deposits stay flat through a consumer downturn, the optionality argument is wrong and a committed book really

is unambiguously good news. And if return-visit rates against an operator-held identifier are indistinguishable from baseline direct traffic, then the identity horizon is not the binding constraint and something else explains the gap.

Limitations

The Accor figure is one operator's characterization of its own customer base in public remarks, not a published distribution, and Accor's mix skews toward markets and segments that may not generalize. The SiteMinder and AirDNA figures come from one platform's transactions each, so they carry that platform's mix. The last-minute share figures circulate through industry reporting as ranges, with no single published series behind them, which is why I have cited them as ranges; the underlying STR indices are subscription data I have not read directly. The advisor-reported cruise windows are self-reported and drawn from a survey population that skews toward advisors with long-lead clientele. The Carnival position is a single operator at a single point in a cycle. The retargeting experiment cited is retail-dominant, and its transfer to travel windows is an assumption of this paper that nobody has measured. Platform behavior changes without notice, and both the cookie-cap correction and the April 2025 Chrome reversal postdate widely circulated guidance.

SOURCES

1. SiteMinder, Booking Lead Time: Hotel Trends
<https://www.siteminder.com/r/booking-lead-time/>
2. AirDNA, Vacation Rental Metrics: Booking Lead Time
<https://www.airdna.co/blog/vacation-rental-metrics-booking-lead-time>
3. Sebastien Bazin, Accor, public remarks, March 2026
4. Carnival Corporation & plc, 2026 first-quarter results, March 2026
5. CLIA, State of the Cruise Industry Report 2025
6. STR, booking window indices versus 2019 baseline, reported 2026
7. WebKit, Tracking Prevention Policy; CNAME Cloaking and Bounce Tracking Defense (2020)
<https://webkit.org/tracking-prevention/>
8. Google Privacy Sandbox, Next steps for Privacy Sandbox and tracking protections in Chrome (22 April 2025)
<https://privacysandbox.com/news/privacy-sandbox-next-steps/>
9. Journal of Marketing Research 56(3), randomized field experiment on retargeting (2019)
10. Blake, Nosko & Tadelis, Consumer Heterogeneity and Paid Search Effectiveness, NBER WP 20171 (2014)
https://www.nber.org/system/files/working_papers/w20171/w20171.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/booking-windows-diverging

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