The World Cup Effect

How The 2026 Tournament Reshaped Short-Term Rental Income In U.S. Host Cities

CONTENTS

What's In This Report

01	The Hypothesis	3
02	How We Tested It	4
03	The Result: A Pricing Event, Not A Booking Boom	5
04	Income Per Listing, By City	6
05	Occupancy: Where The Hypothesis Broke Down	7
06	Confirming The Rate Result: Six Estimators, One Answer	9
07	Where It Mattered Most: Distance To The Venue	10
08	Stadium, Or Just Downtown?	11
09	Every City Has Its Own Halo	12
10	When It Mattered: Timing And The Match Calendar	13
11	Where And When Together: The Halo Breathes	14
12	What We Expected But Couldn't Confirm	16
13	Why Cities Differed: The Baseline Price Mechanism	17
14	Stress-Testing The Method: The Placebo Test	18
15	What This Data Does Not Support	19
—	Three Things To Carry Away	21
A	Appendix A · Methodology	23
B	Appendix B · Data Sources	24
C	Appendix C · Limitations, In Full	25
D	Appendix D · The Model, Written Out	26

HOW TO READ THIS REPORT

Trust The Rate Figure
Six estimators agree on it. Use the point estimate.

Occupancy Is Exploratory
Use it for the shape of the story, not the magnitude.

Not A League Table
City ordering reflects baseline price, not performance.

01

SECTION 01

The Hypothesis

Before we looked at a single row of data, we had a guess, and it was the same guess most people would make. A mega-event like the World Cup brings hundreds of thousands of visitors into a handful of cities over a few short weeks. The obvious hypothesis is that short-term rental hosts in those cities would see their calendars fill up. More visitors, more nights booked, more occupancy. That was the prediction we set out to test.

If that hypothesis were true, we'd expect to see three specific things in the data:

PREDICTED

Occupancy Climbs

Host-city occupancy rising well above what similar non-host cities saw in the same weeks.

PREDICTED

Rate Follows Demand

Rate increases as a secondary effect, because demand for nights outstripped supply.

PREDICTED

More Matches, More Lift

The effect scaling with how many matches a city hosted.

The hypothesis didn't survive contact with the data. What we found instead was close to the mirror image of what we expected, and the rest of this report is the evidence trail that got us there.

SECTION 02

How We Tested It

Here's the short version of the method, so you can trust what follows. The full model lives in Appendix A and Appendix D.

01

Ten Hosts, Seven Controls

The closest control city sits 330 km from any host metro — far enough that the tournament couldn't inflate it.

02

Three Channels, Modeled Apart

Rate, nights filled, and income per listing, each with no shared assumptions.

03

Every Listing Against Itself

Change measured within each listing, never one group of listings against another.

We compared ten U.S. host metros against seven similar cities that hosted no matches, using those seven as a control group. The closest control city sits 330 km from any host metro — far enough that the tournament in one city couldn't itself be inflating the numbers in a nearby "control."

We measured three things independently: the rate hosts charged, the share of nights they filled, and the resulting income per listing. Each was modeled separately, with no shared assumptions, so that agreement between them would mean something rather than being built into the math.

Most importantly, we measured change *within each individual listing* — comparing a listing to itself before the tournament, rather than comparing one group of listings to a different group. That turned out to matter a great deal. An earlier version of this analysis skipped that step and produced a false positive (Section 14).

SECTION 03

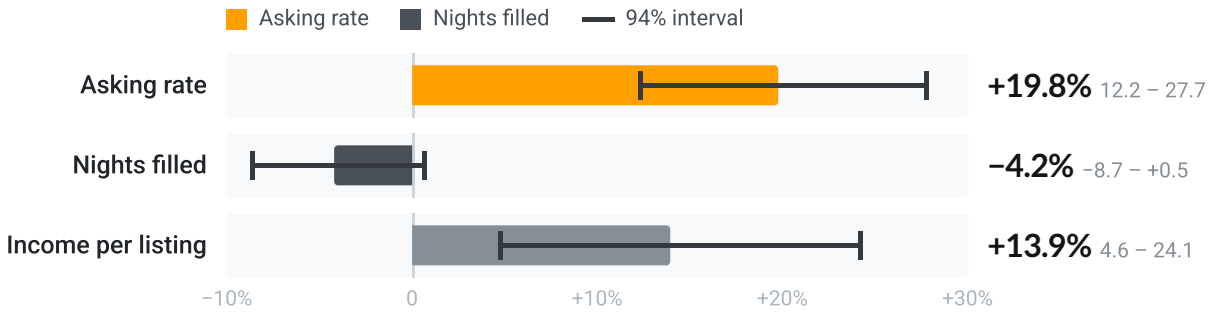
The Result: A Pricing Event, Not A Booking Boom

Here's what we actually found, against what the hypothesis predicted.

Figure 1 — RATE · OCCUPANCY · INCOME

A Pricing Event, Not A Booking Boom

Change during the tournament window vs matched control metros, with 94% credible intervals



+19.8% Rate is the solid estimate. Nights filled and income are exploratory — the nights interval touches zero. Use them for the shape of the story, not for precise magnitudes.

Three separately-modeled channels. Rate and occupancy reconcile with the independently-measured income change to within less than 1% in log terms.

- **Nights filled: down about 4%, not up.** The hypothesis predicted this would climb. It fell.
- **Asking rate: up about 20%.** The hypothesis treated this as a secondary effect. It turned out to be the entire story.
- **Income per listing: up about 14%**, driven almost entirely by rate, not occupancy.

This isn't a small miss on the original hypothesis; it's an inversion of it. Hosts didn't fill more nights — they charged more for the nights they were already going to fill. Three separate models point at the same conclusion.

The honest caveat: the occupancy and income figures are softer than the rate figure, since the underlying booked-nights data carries a quality issue in 2024 that touches every metro equally. The rate number is the one to trust down to the point estimate.

SECTION 04

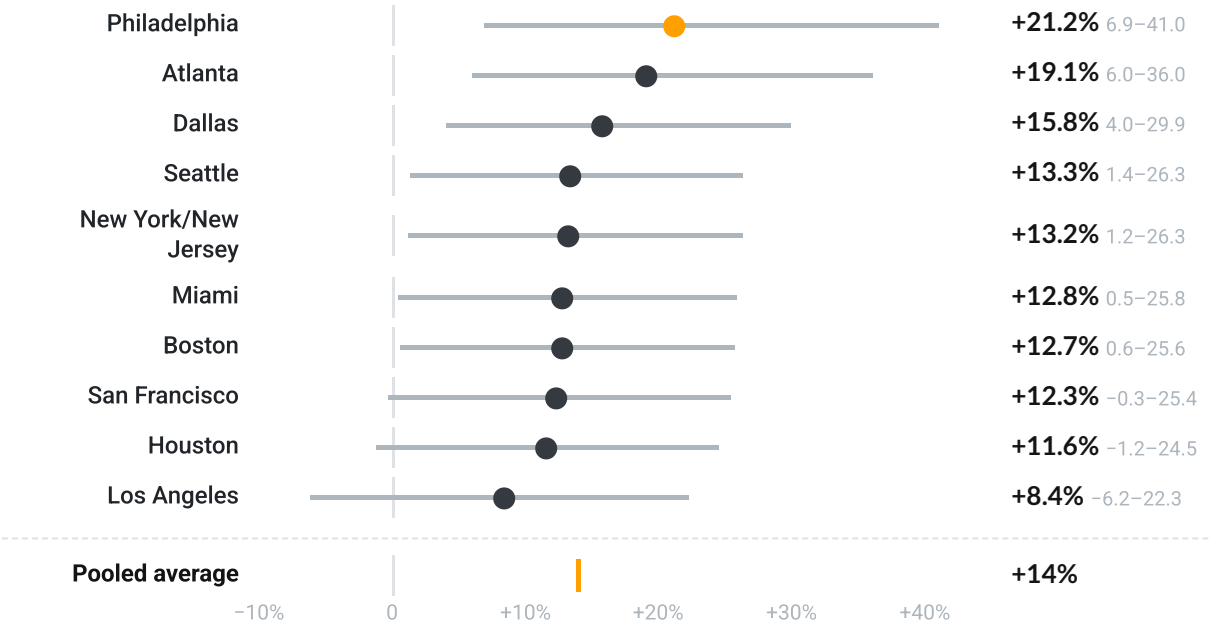
Income Per Listing, By City

Every one of the ten host cities showed a positive income lift. None showed a loss.

Figure 2 — INCOME BY METRO

Ten Cities, One Overlapping Range

Income lift per listing vs matched control metros, point estimate and 94% interval



Not a leaderboard: the intervals overlap almost completely and three of the ten still reach below zero. Section 13 explains what actually drives the spread.

It would be easy to read this as a ranking. It isn't one. The ordering here mostly reflects each city's baseline price, not how well any of them "performed."

SECTION 05

Occupancy: Where The Hypothesis Broke Down

This is the section where the original hypothesis fails outright. If the World Cup were a demand-volume event, this is the chart that should have lit up. It didn't.

Figure 3 — OCCUPANCY

The Cities Did Not Fill Up — Occupancy Fell Modestly

Market occupancy change in the event window, percentage points vs control metros



−4.7%	+1.5%	−6.2 pp
Host metros, mean change	Control metros, mean change	Host minus control

Hosts raised rates ~20% and lost ~4% of nights filled. That is a description, not an elasticity — no price/demand trade-off is quantified here.

Occupancy fell in nine of the ten host metros relative to control cities, by about 6 percentage points on average. San Francisco was the only exception. A second, independent measurement — scraped booked-nights rather than scraped availability — agrees on both direction and size.

> **This was never a story about cities filling up. It was a story about the price of the nights that were already filling.**

We're not going to tell you that means demand was "inelastic," because we can't isolate a real elasticity here: supply, listing mix, and the event were all shifting together. What we can say plainly is that this was never a story about cities filling up.

06

SECTION 06

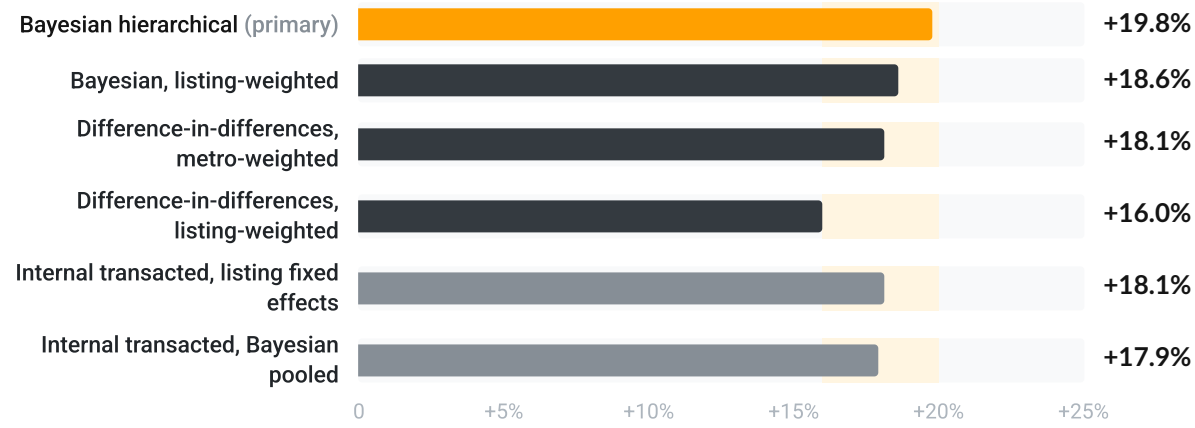
Confirming The Rate Result: Six Estimators, One Answer

Since the rate finding is now carrying the whole hypothesis, we stress-tested it as hard as we could before trusting it.

Figure 4 ROBUSTNESS

Six Estimators Land In The Same Narrow Band

Asking-rate lift by method — two statistical frameworks, two independent datasets



The tinted band marks +16% to +20% — every estimator falls inside it. This is why we call the pricing result a finding rather than a hunch.

SECTION 07

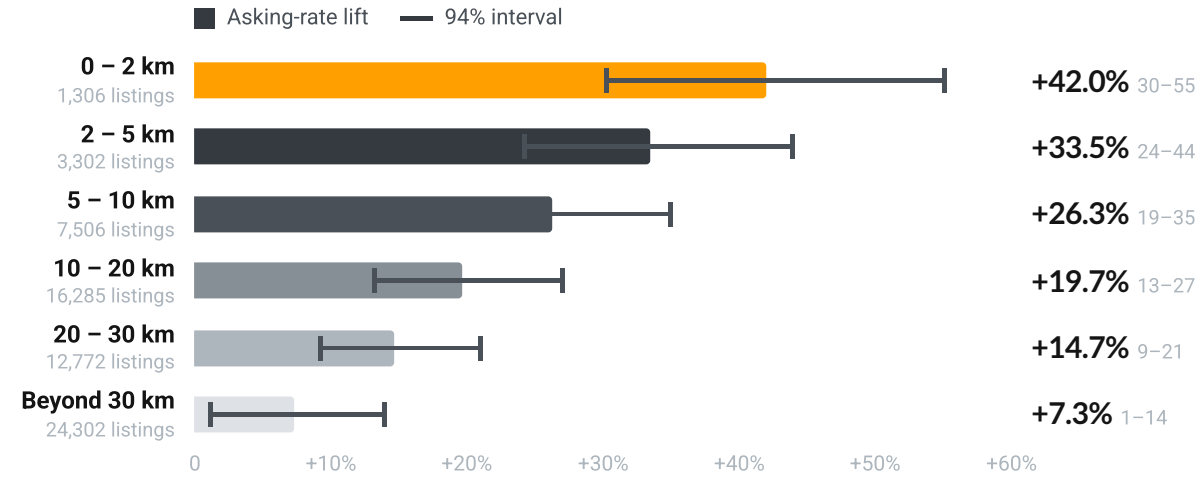
Where It Mattered Most: Distance To The Venue

Once we accepted that this was a pricing story, the next question was obvious: was the increase spread evenly across a host city, or concentrated somewhere specific? It was concentrated, and heavily.

Figure 5 — DISTANCE TO VENUE

Where A Listing Sat Mattered More Than Which Host City It Was In

Asking-rate lift vs control metros, by distance from the World Cup venue



Write the near-venue figure as a range: +42% (Bayesian, pooled), +35% (frequentist within-listing) and +30% (winsorised) are all defensible. Distance to downtown is flat over the same data — this is a venue effect, not a city-centre effect.

The closest ring is the most sensitive number in this report: a small number of outlier listings can swing it several points.

Which host city you were in mattered less than where in that city you were. The responsible way to state the near-venue effect is a range — roughly **+30% to +42%**. We'd rather give you the honest range than round up to the biggest number.

SECTION 08

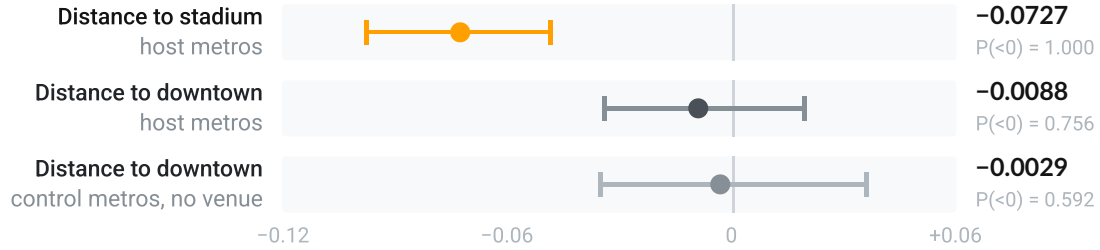
Stadium, Or Just Downtown?

Before accepting "proximity to the stadium" as the explanation, we had to rule out a more boring one: maybe this was never about the stadium at all, just about being downtown. Several venues sit well outside city centres — AT&T Stadium is in Arlington, Gillette in Foxborough, MetLife in East Rutherford — so the two are easy to confuse.

Figure 6 VENUE VS CITY CENTRE

Only The Stadium Slope Is Distinguishable From Zero

Slope per standard deviation of log-distance, both distances entered in the model at once



The stadium interval sits entirely left of zero; both downtown intervals straddle it. There is a 99.8% probability the stadium slope is genuinely steeper than either downtown comparison.

Distance to the stadium came back clearly negative. Distance to downtown didn't — in host cities or in the control cities that have no venue at all. The alternative explanation didn't hold up. This really is a venue effect.

09

SECTION 09

Every City Has Its Own Halo

The pooled curve in Section 7 averages across ten cities whose rental inventory sits in very different places relative to their venues. Four things are worth knowing before reading your own city's numbers.

Dallas & Houston — The Clean Case

Real stadium-adjacent inventory and the largest near-venue premiums. Closer really was better, and it was something a host could act on.

Philadelphia & New York/New Jersey

Steepest decay curves on paper, but near-venue bands are nearly empty: 10 listings inside 2 km in Philadelphia, none in New York/New Jersey.

Boston — A Comparison Among Far Listings

Gillette Stadium sits about 34 km out, so Boston has almost no listings inside 20 km. Its curve compares far listings to farther ones.

Atlanta — A Gradient, Not A Hard Line

Its 2–5 km band reads above its 0–2 km band, likely because those bands aren't cleanly separable at that sample size.

Treat the pooled curve in Section 7 as the number to trust. City-level variation is for understanding your own inventory, not for building a ranking — the ordering mostly reflects baseline price and where inventory happens to sit.

SECTION 10

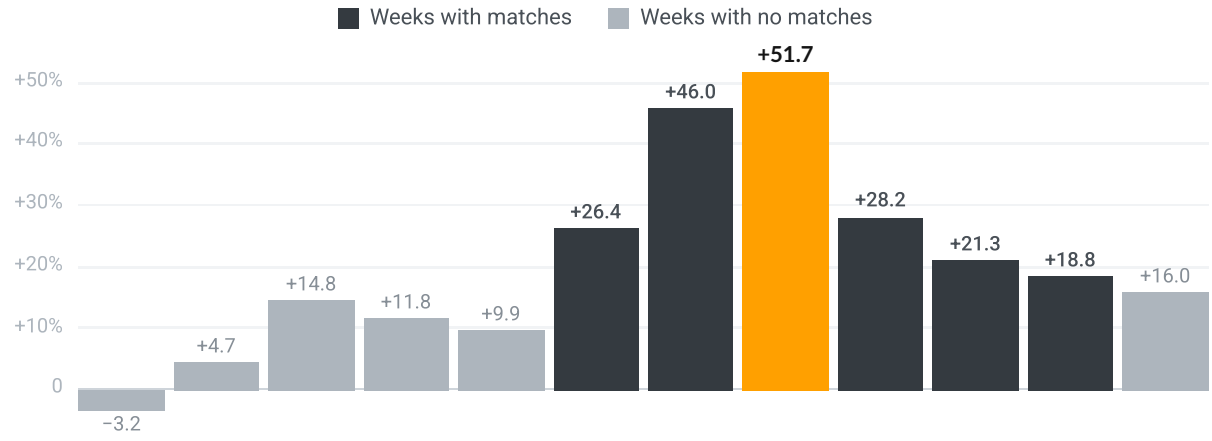
When It Mattered: Timing And The Match Calendar

If proximity was the "where," this is the "when." The premium wasn't constant — it rose and fell with the match schedule itself.

Figure 7 — EVENT WEEK

The Premium Arrived And Left With The Match Calendar

Asking-rate lift for stay-nights in each week, relative to matched control metros



Week	−6	−5	−4	−3	−2	−1	0	1	2	3	4	5
Matches	—	—	—	—	—	7	19	23	13	6	4	—

This is anticipation in pricing for future stay-dates, not evidence about when guests booked. Booking dates are not observable here.

Intervals (94%) span roughly ±13 points on each week; weeks −6 to −2 include zero, weeks −1 through 5 do not.

There's a measurable run-up before the first match: rates for future stay-dates were already climbing weeks in advance. That's anticipation baked into pricing — a sign this was priced in ahead of time, not a surprise that hit all at once.

SECTION 11

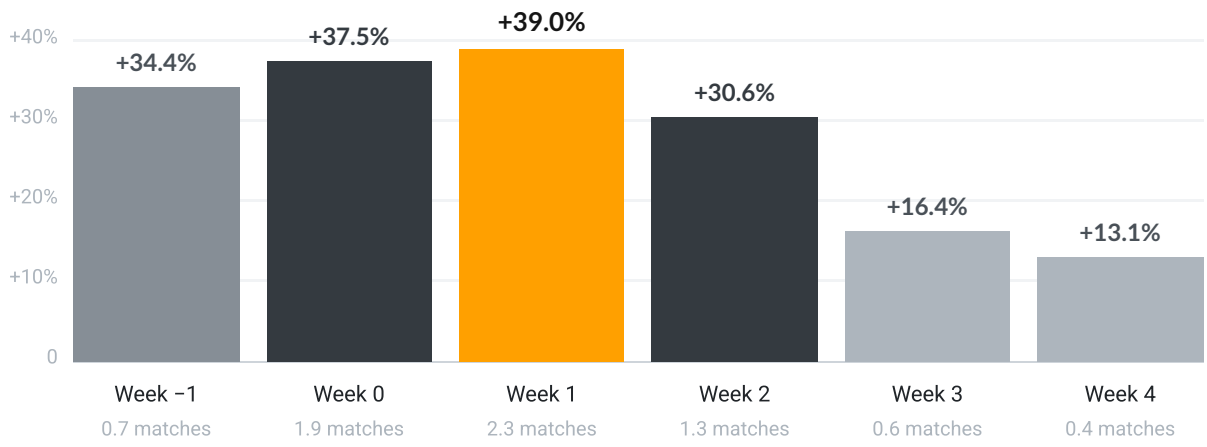
Where And When Together: The Halo Breathes

This is the most decision-relevant chart in the report, because it shows location and timing interacting rather than acting on their own.

Figure 8 — NEAR VS FAR, SAME METRO-WEEK

The Premium Was Rented, Not Owned

Near-venue premium (listings within 5 km vs beyond 30 km) in the same metro and week



+17 pp Going from zero to two matches in a metro-week widens the near-versus-far gap by roughly 17 percentage points. Proximity paid most in the weeks the venue actually hosted matches.

A within-metro spatial comparison, not a control difference-in-differences: it says how much more a listing at 2.9 km gained than one at 52.8 km in the same week, not the absolute lift.

Being close to a venue was worth roughly three times more in the heaviest match weeks than in the quiet ones. Proximity was rented, not owned.

It wasn't a fixed asset that paid off evenly for six weeks — the halo expanded and contracted with the match calendar.

One caveat that makes this conservative rather than inflated: because several venues sit outside their city centres, "beyond 30 km" is sometimes a city's best neighbourhood. In Boston, 82% of listings sit beyond 30 km from Gillette Stadium. There, the comparison

is really "near the stadium versus the best part of town," which understates the effect if anything.

SECTION 12

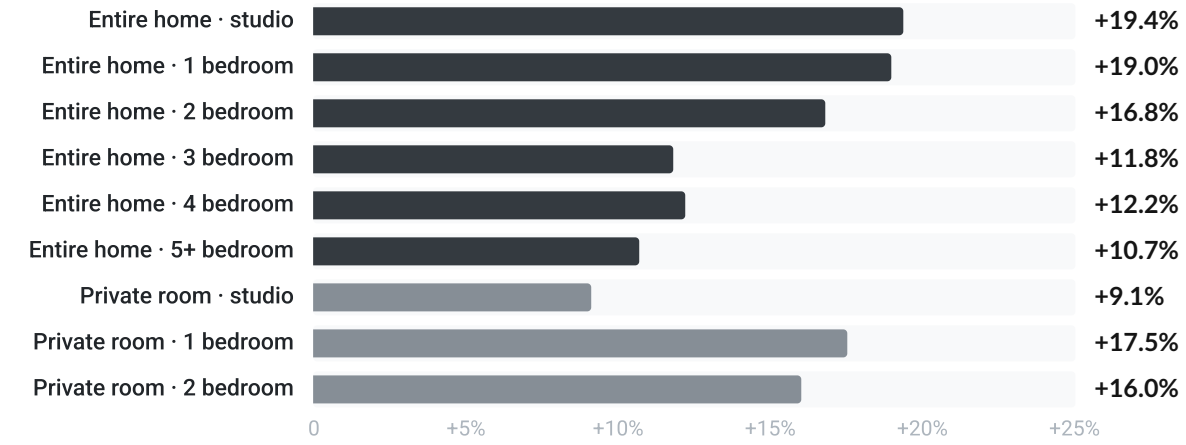
What We Expected But Couldn't Confirm

Two more predictions from the original hypothesis didn't hold up, and both are worth stating plainly.

Figure 9 PROPERTY TYPE

What You Listed Mattered Far Less Than Where It Was

Asking-rate lift by property type and size — a 10-point spread, against 35 points across distance bands



Studios, private rooms, and five-bedroom homes all moved within about 10 percentage points of each other.

We expected more matches to mean a bigger city-level lift. We couldn't detect that relationship at all. Every measure of match volume produced an interval spanning zero, and the sign flipped depending on the method. That means we couldn't detect an effect — not that there definitely isn't one. Within the tournament, heavier match weeks clearly saw bigger lifts (Section 11); across only ten cities there simply isn't enough variation in match count to detect a city-level relationship, even a modest real one.

SECTION 13

Why Cities Differed: The Baseline Price Mechanism

Match volume explained none of the spread between cities. Something else explained most of it.

Figure 10 — WHY CITIES DIFFERED

Expect A Dollar Premium, Not A Percentage

Median nightly asking price before the tournament vs the metro's asking-rate lift



+$55

The event added roughly a fixed dollar amount per night — about +$55, varying only ~$9 between cities. A fixed dollar premium is a big percentage in a cheap market and a small one in an expensive one. Baseline price explains about 80% of the between-city variance; match volume explains none of it.

Do not build a league table from this: the per-city intervals overlap almost completely, and the ordering reflects baseline price, not performance. Spearman(baseline, % lift) = −0.83.

Cheaper cities showed bigger percentage lifts. Los Angeles looks like a modest performer and Philadelphia like a strong one, but they made similar money on very different starting prices. This is the most useful mechanism in the report for a host: **expect a dollar premium, not a percentage.** What percentage you personally see depends entirely on what you already charge.

14

SECTION 14

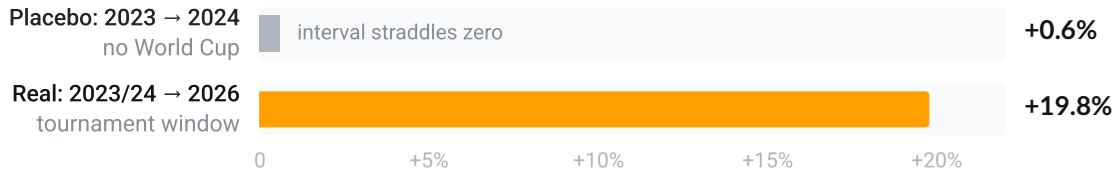
Stress-Testing The Method: The Placebo Test

Before you believe any of the above, the method itself has to survive a null test. We re-ran the entire analysis pretending 2024 was the "treatment year" against a 2023 baseline — two ordinary summers, no tournament. A sound design should return approximately zero.

Figure 11 PLACEBO TEST

A Year With No Tournament Returns Essentially Nothing

Estimated asking-rate effect, placebo window vs the real event window



An earlier version of this study, run on a smaller set of internally-managed listings, failed this same test — returning a false +16% "effect" in a year with no World Cup.

That earlier failure is why the analysis looks the way it does. A small portfolio of listings churns between years, so a simple metro-level average ends up comparing different properties to each other rather than the same properties over time. The fix was measuring change within each individual listing, across a full market-wide dataset. Every number in this report comes from that corrected, tested version.

SECTION 15

What This Data Does Not Support

THE MOST LIKELY MISREADING

Don't say: "The World Cup made guests willing to pay 42% more near stadiums."

Say instead: "Asking rates within 2 km of World Cup venues ran roughly 30–42% above matched control metros during the tournament."

We never observe willingness to pay in this data. We observe posted rates on nights that sold. That number bundles genuine guest demand together with however hosts and their pricing tools responded, and the two can't be separated here.

DON'T SAY	SAY INSTEAD
"Hosts earned 42% more near stadiums."	"Asking rates ran roughly 30–42% higher near venues. Income also depends on occupancy, which fell slightly."
"Occupancy held steady."	"Occupancy fell modestly, about 6 points relative to control metros."
"Demand was inelastic."	"Hosts raised rates about 20% and lost about 4% of nights filled." No elasticity is estimated here.
"Match volume doesn't affect a city's rates."	"We could not detect a relationship. The study is too small to rule out a modest one."
"Rates rose 20% in host cities," with no distance qualifier.	"Rates rose roughly 16–20% on average — from roughly 30–42% beside a venue to about 7% beyond 30 km."
"Travelers booked earlier."	"Rates for stay-dates before the tournament were already elevated." We don't observe when guests booked.
Anything about Kansas City.	Not covered at all. See Appendix C.
"Near MetLife, rates rose 42%."	New York/New Jersey is covered at the metro level (+22.7%), but zero listings sit within 2 km of MetLife itself.

The distance gradient, timing pattern, and metro-level lifts are treated as causal, conditional on three things holding: parallel pre-tournament trends (tested, Section 14), no spillover into control metros (checked — the nearest is 330 km away), and no other 2026-specific shock hitting host metros differently (not tested, see Appendix C).

IN CLOSING

Three Things To Carry Away

It Was A Pricing Event

Rates rose about 20%; nights filled fell about 4%. Income followed rate, not occupancy.

Proximity Was Rented

The near-venue premium tracked the match calendar — worth roughly three times more in heavy match weeks.

Expect Dollars, Not Percent

The event added about +$55 per night. What percentage that is depends on what you already charge.

Want This Analysis For Your Own Market?

iGMS Research runs rate, occupancy, and income benchmarks across 17 U.S. metros. Tell us which listings you manage and we'll show you where your pricing sat against the market during the tournament window.

Book a Demo

This report is the complete companion to our summary piece on what the 2026 World Cup did to short-term rental income. For questions about methodology or data access, reach out to the research team.

SUPPORTING MATERIAL

Appendices

Method, sources, limitations, and the statistical model. Everything from here on verifies the findings above rather than adding to them.

Appendix A · Methodology 23

Appendix B · Data Sources 24

Appendix C · Limitations, In Full 25

Appendix D · The Model, Written Out 26

APPENDIX A

Methodology

Design

A difference-in-differences study. Ten U.S. host metros (Atlanta, Boston, Dallas, Houston, Los Angeles, Miami, New York/New Jersey, Philadelphia, San Francisco, Seattle) compared against seven non-host donor metros (Charlotte, Denver, Indianapolis, Nashville, New Orleans, Orlando, Phoenix). The closest donor city sits 330 km from any host metro. The treated window runs June 12 to July 19, 2026, against the same calendar weeks in 2023 and 2024.

Model

A Bayesian hierarchical model with metro-level partial pooling, city-specific distance decay, and property-segment effects. The core design choice, and the one that made this trustworthy: every effect is measured by comparing each listing to itself before the tournament, never by comparing one group of listings to a different group.

Why The Intervals Are Wide, And Why That's Correct

Treatment was assigned to entire metros, not individual listings. The effective sample size behind city-level questions is 17 metros, not hundreds of thousands of listings. Ignoring that would make every interval in this report roughly 6 to 13 times too narrow. An independent metro-clustered bootstrap reproduces intervals of the same width.

APPENDIX B

Data Sources

Two independent sources, used deliberately so one could check the other.

1. **Market-wide scraped listing data.** Roughly 385,000 listings, 569,000 listing-years, and 10.9 million booked nights across 17 metros and three years. The basis for every headline figure.
2. **Internal transacted reservations.** A narrower dataset recording money actually paid, used throughout to corroborate direction and magnitude, never as the sole basis for a number.

APPENDIX C

Limitations, In Full

- **Asking price, not amount paid.** The gap between posted rate and amount paid likely shifted somewhat between baseline and event. Most of that cancels out in the comparison, but a residual likely remains — meaning these figures probably understate the true effect by around 10%.
- **Income and occupancy are exploratory.** The underlying booked-nights measure carries a data artifact in 2024, across every metro. The rate estimate is the solid one.
- **2025 is excluded entirely**, due to a coverage failure in the scraped record that year.
- **No elasticity estimate exists.** We cannot say how much of the gain was price versus fill in any structural sense.
- **The shape of the near-venue halo is unresolved.** It decays with distance, but whether smoothly or as a plateau-then-cliff isn't established.
- **No test exists for a competing 2026-specific shock.** The placebo test rules out a pre-existing divergence between host and control metros. It can't rule out something else entirely unrelated to the World Cup.
- **Kansas City is excluded, and no claim here covers it.** Every Kansas City price in the 2026 data was roughly 2.4× its 2023/24 level — across cheap, mid-range, and expensive listings alike, all multiplied by nearly the same factor. Real demand doesn't behave that way. Listings more than 30 km from the stadium also rose about 93%, which doesn't fit a match-tourism story. We couldn't identify what rescaled the series, so we removed the city rather than publish a number we couldn't explain.
- **New York/New Jersey is fully covered at the metro level; the Final's stadium surroundings are not.** Of its 4,544 listings, zero sit within 2 km of MetLife Stadium. The near-venue findings rest on metros that actually have near-venue supply: chiefly Atlanta, Seattle, Dallas, Houston, and Philadelphia.
- **One open, unverified question on New York/New Jersey.** Its listing count roughly halved between 2023 and 2024, coinciding with NYC's short-term-rental registration law (Local Law 18) entering enforcement. If that's the cause, it's a real regulatory shock rather than a data artifact, and it could be inflating that metro's measured 2026 lift. Inferred from timing alone, not verified.
- **Length of stay and booking lead time** are not covered at any meaningful scale, and **willingness to pay** is not measured — no rewording makes it measurable.

APPENDIX D

The Model, Written Out

For readers who want to see the statistical machinery underneath the findings. Skipping this loses none of the substantive findings.

D.1 · THE ESTIMAND

$$\log p_{i,t} = \alpha_{m(i)} + \delta_t + \gamma \cdot \text{host}_m \times 1[t = 2026] + \varepsilon_{i,t}$$

$\alpha_{m(i)}$ metro baseline price level

δ_t whatever happened to the whole U.S. market that year

γ the tournament effect — the quantity this report measures

$\varepsilon_{i,t}$ listing-level residual

D.2 · WHY THE INTERVALS ARE WIDE

$$\bar{y}_{mt} \sim \mathcal{N}(\mu_{mt}, \sqrt{s^2_{mt} + \sigma^2})$$

Two sources of noise. s reflects how precisely we know a single metro-year's average price, and shrinks with more listings. σ reflects a metro-year shock that doesn't shrink no matter how much data you add. This is why 385,000 listings don't buy a narrow interval on their own.

D.3 · PARTIAL POOLING

$$\gamma_m \sim \mathcal{N}(\mu_\gamma, \tau_\gamma), \quad m = 1 \ldots 10$$

Each host city gets its own effect, but all ten are treated as draws from one shared distribution — so cities inform each other's estimates instead of being analysed in isolation.

D.4 · WHAT SHRINKAGE DOES

$$w_m = \frac{\tau^2_\gamma}{\tau^2_\gamma + s^2_m}$$

How much a city's own data is trusted versus the pooled average. A noisier or more extreme city gets pulled harder toward the group mean — which is the lower-error estimate of that city's true effect, not just a cautious hedge.

D.5 · MEASURING CHANGE INSIDE ONE LISTING

$$\Delta_i = \log p_{i,2026} - \text{mean}(\log p_{i,base})$$

Every listing is compared to itself, never to a different set of listings. This is the specific fix for the failure that produced a false +16% effect in an earlier version of this study.

D.6 · STADIUM VERSUS DOWNTOWN, SIMULTANEOUSLY

$$\mathbb{E}[\bar{\Delta_c}] = \theta_m + \kappa_m z_c + \lambda_m w_c$$

Both distances — z to the stadium and w to downtown — enter the model at once, so each is measured holding the other fixed. Without this, a genuine venue effect and an ordinary downtown premium would be statistically indistinguishable.

D.7 · THE CAUSAL CONTRAST AT A GIVEN DISTANCE

$$DiD(z, w) = (\mu_\theta - \mu_U) + \mu_\kappa z + (\mu_\lambda - \mu_V) w$$

The host-metro change minus the counterfactual built from control metros, at any chosen distance. The v term prices in whatever a generic city core did on its own in 2026 — which is what makes this a genuine causal contrast.

CREDITS

Research and analysis by Sem Leontev, Data Scientist

© 2026 iGMS. All rights reserved.