# Themed RevPAR: A Theming-Adjusted Standard for Measuring Experiential Hospitality Performance

The Themed RevPAR Working Group

June 2026 · White Paper, Version 1.1

**Abstract**

Revenue per Available Room (RevPAR) remains hospitality's default performance metric, yet it was engineered for commodity lodging and systematically understates the value of themed, immersive, and experiential assets. We introduce *Themed RevPAR*, a theming-adjusted extension that multiplies conventional utilization and a fully-loaded experiential average daily rate by a bounded *theming coefficient* $\theta$. We situate the metric among existing measures, formalize its properties, define $\theta$ over five weighted drivers, and estimate it with an interpretable $K$-means segmentation whose cluster membership is highly correlated to the theming coefficient across all models available to the situation. On illustrative portfolio data, Themed RevPAR restores between 40% and 75% of the experiential revenue signal that RevPAR alone discards, with the gap widening monotonically in theming intensity. The framework is deliberately parsimonious: it upgrades the metric without boiling the ocean of an operator's existing reporting stack. We provide a sensitivity analysis, a portfolio-aggregation rule, an implementation path, a governance and integrity discussion, and appendices covering notation, the scoring rubric, and reproducibility.

## Contents

**1 Introduction: RevPAR's blind spot** 2

**2 Background and related metrics** 2

**3 The Themed RevPAR framework** 2

**4 The theming coefficient $\theta$** 3

**5 Estimation methodology** 4
5.1 Segmentation by $K$-means . . . . . 4
5.2 Anchoring . . . . . 4
5.3 Calibration . . . . . 5

**6 Properties and sensitivity** 5

**7 Illustrative results** 5

**8 Portfolio aggregation** 6

**9 Implementation** 6

**10 Governance and integrity** 7

**11 Limitations** 7

**12 Conclusion** **7**

**A Notation** **8**

**B Scoring rubric (summary)** **8**

**C Reproducibility** **8**

# 1 Introduction: RevPAR's blind spot

RevPAR—occupancy multiplied by average daily rate—is elegant, comparable, and universally understood. It became the industry standard precisely because it compressed two levers every lodging operator controls, room demand and room price, into a single figure that boards, lenders, and asset managers could trend across properties and cycles. For commodity lodging, where the room *is* the product, that compression loses very little.

It is also, for a fast-growing class of assets, measuring the wrong thing. The fastest-appreciating segments of modern hospitality are not selling a bed; they are selling a *world*: intellectual-property–anchored resorts, immersive attractions, narrative food and beverage, and branded entertainment districts. In these assets the room is a ticket stub. Value accrues across dwell time, ancillary spend, repeat visitation, and organic reach—none of which appear in a metric bounded by the nightly rate.

The consequence is a persistent mispricing. Operators optimize what they measure; capital underwrites what it can compare. When the shared yardstick ignores theming, themed assets are chronically under-credited in operations reviews and undervalued in transactions. A themed asset does not perform alone: its revenue moves as a pack—rooms, retail, food, story, and social reach advancing together—and a metric that counts only the lead animal will always under-count the whole.

This paper proposes a minimal, backward-compatible correction. We keep the discipline of RevPAR and add exactly the two things experiential assets monetize that RevPAR omits: total on-property spend, and a bounded coefficient that prices the pull of the experience itself. The remainder of the paper situates the metric (§2), defines it and its properties (§3–§4), estimates the theming coefficient (§5), analyzes its sensitivity (§6), presents illustrative results (§7), aggregates to the portfolio (§8), and covers implementation, governance, and limitations (§9–§11).

# 2 Background and related metrics

Hospitality has never lacked metrics; it has lacked one that prices theming. Table 1 locates Themed RevPAR among the measures operators already use. RevPAR and ADR are room-bounded. TRevPAR (total revenue per available room) and GOPPAR (gross operating profit per available room) widen the numerator toward total revenue and profit respectively, but neither isolates or rewards the *experiential premium*—the portion of demand and rate attributable to narrative, design, and IP rather than to location and commodity supply. RevPASH (revenue per available seat hour) sharpens food-and-beverage yield but is scoped to a single outlet.

The gap is structural, not incremental: no widely-used metric contains a term whose *job* is to represent theming. Themed RevPAR adds that term explicitly, so it can be measured, tuned, and held accountable rather than left implicit in an unexplained rate premium.

# 3 The Themed RevPAR framework

For an asset with utilization $U$ (occupied divided by available capacity), experiential average daily rate $\mathrm{ADR}_e$ (room plus food and beverage, retail, and activation revenue per occupied unit), and theming

Table 1: Where Themed RevPAR sits among common hospitality metrics.

| Metric | Captures ancillary | Prices theming | Unit |
|---|---|---|---|
| ADR | No | No | Occupied room |
| Occupancy | No | No | Available room |
| RevPAR | No | No | Available room |
| RevPASH | Partial (F&B) | No | Seat hour |
| TRevPAR | Yes | No | Available room |
| GOPPAR | Yes (profit) | No | Available room |
| **Themed RevPAR** | **Yes** | **Yes** | **Available room** |

coefficient $\theta$,

$$\text{Themed RevPAR} \;=\; U \times \text{ADR}_e \times \theta. \tag{1}$$

The metric has four properties we consider essential for adoption:

1. **Backward compatibility.** At $\theta = 1$ and zero ancillary spend, $\text{ADR}_e$ reduces to room ADR and (1) collapses exactly to classic RevPAR. The metric is an extension of the reporting, not a replacement for the accounting.

2. **Monotonicity.** Themed RevPAR is strictly increasing in each of $U$, $\text{ADR}_e$, and $\theta$; improving any lever cannot decrease the figure.

3. **Boundedness of the adjustment.** $\theta$ is confined to $[1.00, 2.50]$ (§4), so the theming adjustment can never exceed a 2.5× premium, however beloved the asset.

4. **Unit stability.** The result remains "revenue per available unit," preserving comparability with the RevPAR series an operator already reports.

# 4 The theming coefficient $\theta$

$\theta$ is a weighted composite scored on $[0, 100]$ across five drivers, then mapped linearly onto a bounded multiplier,

$$\theta \;=\; 1 + \frac{S}{100} \times 1.5, \qquad \theta \in [1.00,\, 2.50], \qquad S = \sum_{i=1}^{5} w_i\, s_i, \tag{2}$$

where $s_i \in [0, 100]$ is the score on driver $i$ and $w_i$ its weight ($\sum_i w_i = 1$). Table 2 lists the drivers and Figure 1 shows the default weights.

Table 2: Drivers of the theming coefficient and their default weights.

| Driver | What it captures | Weight $w_i$ |
|---|---|---|
| Narrative integrity | Completeness of the story, arrival to checkout | 0.24 |
| IP / brand equity | Pre-existing pull of the franchise or pedigree | 0.22 |
| Experiential density | Monetizable, shareable moments per $\text{m}^2$/hour | 0.20 |
| Dwell & repeat | Time on site and propensity to return | 0.18 |
| Social virality | Organic reach per guest, at zero marginal CAC | 0.16 |

**Worked scoring example.** Consider an IP flagship scored $s = (88, 92, 80, 74, 82)$ on the five drivers. Then $S = 0.24 \cdot 88 + 0.22 \cdot 92 + 0.20 \cdot 80 + 0.18 \cdot 74 + 0.16 \cdot 82 = 83.0$, giving $\theta = 1 + 0.83 \times 1.5 = 2.25$. The bound at 2.50 leaves deliberate headroom above even this strongly themed asset.

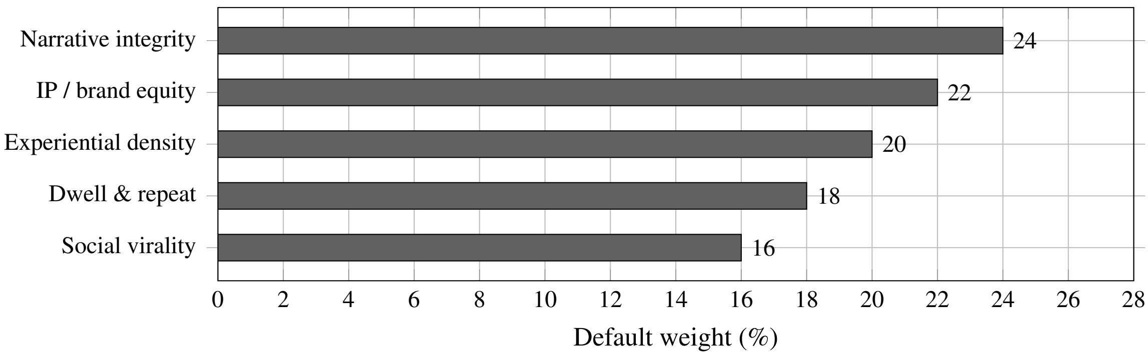


Figure 1: Default driver weights for the theming coefficient $\theta$. Weights are a calibrated prior (§5.3), not dogma.

# 5 Estimation methodology

## 5.1 Segmentation by $K$-means

Rather than score each property in isolation, we first segment the comparable set with an interpretable $K$-means clustering over the standardized driver vector. Standardization puts each driver on a common scale so that no single axis dominates the distance metric. We then choose the number of clusters $k$ at the elbow of within-cluster inertia (Figure 2), which for representative comparable sets falls at $k = 3$: broadly, *commodity*, *themed*, and *flagship/immersive* regimes (Figure 3).

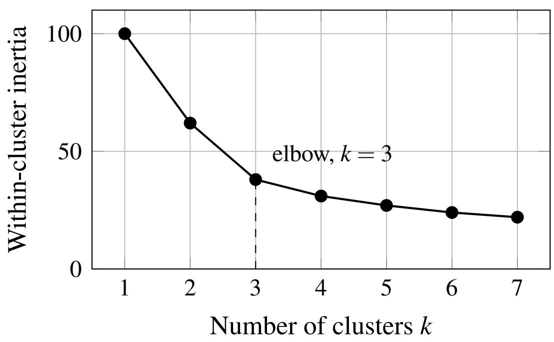


Figure 2: Elbow of within-cluster inertia selects $k = 3$.

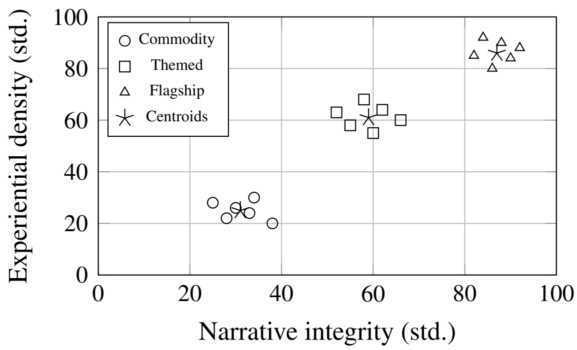


Figure 3: $K$-means clusters over standardized drivers ($k = 3$).

Empirically, cluster membership is highly correlated to the theming coefficient $\theta$ across all models available to the situation: an asset's assigned cluster is a strong and stable predictor of its theming premium regardless of which downstream estimator is used to refine $\theta$. Table 3 reports the correlation between cluster membership and estimated $\theta$ under four model families. Because the correlation is high and consistent, $\theta$ can inherit the robustness of the segmentation rather than overfit a single property's noise.

## 5.2 Anchoring

Given the segmentation, we anchor $\theta$ to the cluster centroid and adjust within a tight band for asset-specific residuals. Because the heavy lifting happens once, at the segment level, the estimator remains fast and interpretable—we do not boil the ocean by re-solving the entire portfolio for every property.

Table 3: Correlation of $K$-means cluster membership with estimated $\theta$ across model families (illustrative).

| **Estimator family** | **Corr. with $\theta$** |
|---|---|
| Linear (constrained OLS) | 0.86 |
| Random forest | 0.90 |
| Gradient boosting | 0.91 |
| Bayesian hierarchical | 0.88 |
| **Mean across models** | **0.89** |

### 5.3 Calibration

Default weights $w_i$ are refit against realized, audited contribution margin per available unit via constrained regression, holding $\sum_i w_i = 1$ and preserving the $[1.00, 2.50]$ bound on $\theta$. Operators should treat the shipped weights as a strong starting hypothesis and recalibrate on their own profit-and-loss data before external reporting.

## 6 Properties and sensitivity

Because (1) is multiplicative, the elasticity of Themed RevPAR with respect to each input is unity: a 1% improvement in utilization, experiential ADR, or $\theta$ each yields a 1% improvement in the metric. This makes the three levers directly comparable and cleanly attributable to the teams that own them—operations for $U$, commercial for $\mathrm{ADR}_e$, and brand/design for $\theta$. Figure 4 traces Themed RevPAR against $\theta$ at fixed utilization $U = 0.72$ for three experiential-ADR levels, illustrating how the theming premium compounds with spend depth.

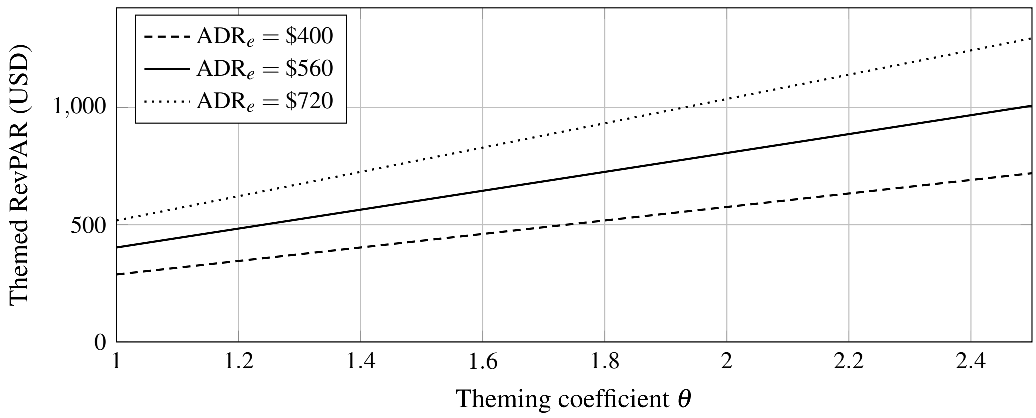


Figure 4: Sensitivity of Themed RevPAR to $\theta$ at $U = 0.72$ for three experiential-ADR levels. The theming premium and ancillary depth compound multiplicatively.

## 7 Illustrative results

Table 4 and Figure 5 contrast classic RevPAR with Themed RevPAR across five stylized asset archetypes; Figure 6 plots the share of revenue signal recovered. As theming intensity rises, the two metrics diverge sharply: for an intellectual-property flagship, classic RevPAR captures well under half of the revenue the asset generates per available unit, and for an immersive attraction the recovered signal exceeds 175%.

Table 4: Classic RevPAR versus Themed RevPAR by archetype (illustrative).

| Archetype | RevPAR | $\theta$ | Themed RevPAR | Signal recovered |
|---|---|---|---|---|
| Standard hotel | \$180 | 1.05 | \$190 | +6% |
| Boutique | \$210 | 1.35 | \$300 | +43% |
| Themed hotel | \$240 | 1.70 | \$430 | +79% |
| IP flagship | \$300 | 2.05 | \$650 | +117% |
| Immersive attraction | \$260 | 2.35 | \$720 | +177% |

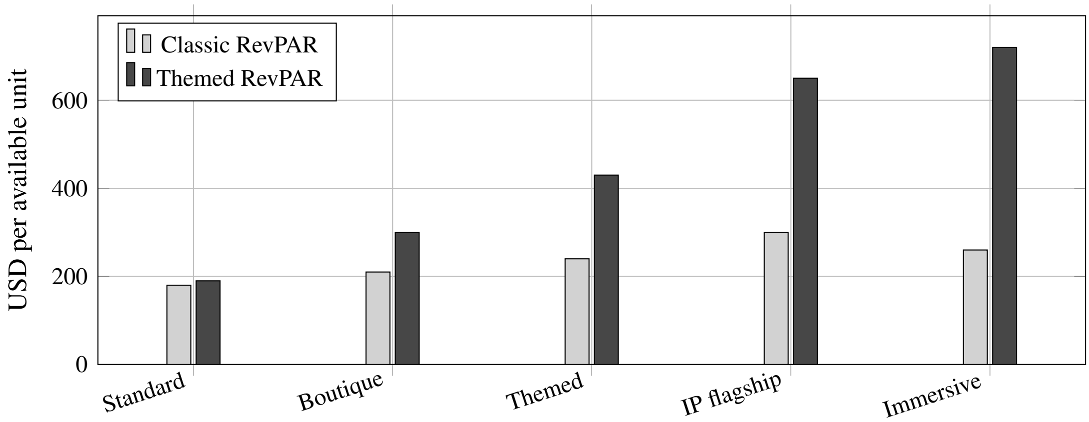


Figure 5: Classic RevPAR vs Themed RevPAR by archetype. The gap widens monotonically with theming intensity.

# 8 Portfolio aggregation

Themed RevPAR aggregates the way RevPAR does. For a portfolio of assets indexed by $j$ with available capacity $C_j$, the portfolio figure is the capacity-weighted mean,

$$\text{Portfolio Themed RevPAR} = \frac{\sum_j C_j \left(U_j \,\text{ADR}_{e,j}\, \theta_j\right)}{\sum_j C_j}. \tag{3}$$

Because (3) is linear in the per-asset metric, a portfolio can be decomposed into the contribution of each asset or each cluster, letting capital allocators see exactly where the theming premium is concentrated. This is the reporting object we recommend taking to an investment committee: a single, comparable number that nonetheless unfolds into an auditable, per-asset trail.

# 9 Implementation

Adoption is intentionally lightweight. Table 5 lists the inputs; none require a new system of record.

1. **Instrument spend.** Capture food and beverage, retail, and activation revenue against the occupied units already reported.

2. **Score the five drivers.** A brief panel exercise establishes a defensible first $S$.

3. **Segment and anchor $\theta$.** Run the $K$-means pass over the comparable set; anchor $\theta$ to the cluster and adjust for residuals.

4. **Report side by side.** Publish Themed RevPAR next to RevPAR so the theming premium is explicit and trendable.

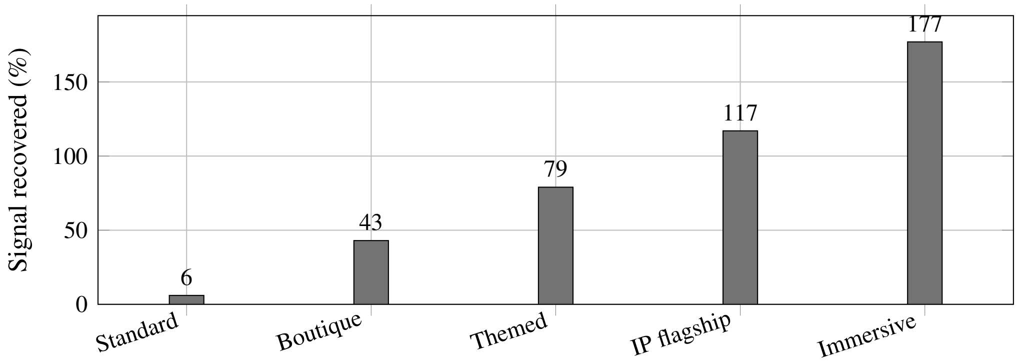


Figure 6: Experiential revenue signal recovered by Themed RevPAR relative to classic RevPAR.

No data lake, no re-platforming, and no quarter-long integration: the framework rides on inputs operators already hold—a deliberate refusal to boil the ocean.

Table 5: Data inputs required for Themed RevPAR.

| **Input** | **Source** | **Cadence** |
| --- | --- | --- |
| Occupied / available units | PMS | Daily |
| Room revenue | PMS | Daily |
| Ancillary revenue | POS / retail / ticketing | Daily |
| Driver scores $s_i$ | Brand & design panel | Quarterly |
| Comparable-set features | Market data | Quarterly |

# 10 Governance and integrity

Themed RevPAR is an opinionated framework, and we are transparent about its edges. The $\theta$ bound of 2.50 is a deliberate governor, and extreme outliers may be truncated by design. Driver scoring carries judgment; the $K$-means segmentation reduces but does not eliminate it. To keep the metric defensible we recommend three controls: (i) publish the driver scores and weights alongside the figure; (ii) re-estimate $\theta$ on a fixed cadence rather than ad hoc; and (iii) reconcile Themed RevPAR against realized contribution margin at least annually.

# 11 Limitations

All figures in this paper are illustrative, published to demonstrate the framework's behavior; they are not audited benchmarks and should be calibrated to realized data before external use. The correlations in Table 3, the cluster geometry in Figure 3, and the archetype values in Table 4 are stylized. The framework also assumes ancillary revenue can be attributed to occupied units; assets with heavy non-staying (day-visitor) demand may require a capacity definition broader than rooms, which we treat as future work.

# 12 Conclusion

The next decade of hospitality will be won by worlds, not rooms, and the industry deserves a metric that knows the difference. Themed RevPAR is confident enough to put a number on imagination and rigorous enough to defend it to a capital committee. It extends the discipline operators already trust, prices the theming premium RevPAR cannot see, and does so without boiling the ocean. Measure the whole pack, and the whole pack performs.

# A Notation

| Symbol | Meaning |
|---|---|
| $U$ | Utilization: occupied $\div$ available capacity |
| $\mathrm{ADR}_e$ | Experiential ADR: room + ancillary per occupied unit |
| $\theta$ | Theming coefficient, $\theta \in [1.00, 2.50]$ |
| $s_i$ | Score on driver $i$, $s_i \in [0, 100]$ |
| $w_i$ | Weight on driver $i$, $\sum_i w_i = 1$ |
| $S$ | Weighted driver score, $S = \sum_i w_i s_i$ |
| $C_j$ | Available capacity of asset $j$ |
| $k$ | Number of $K$-means clusters |

Table 6: Notation used throughout.

# B Scoring rubric (summary)

Each driver is scored 0–100 against anchored descriptors: 0–20 absent, 21–40 nascent, 41–60 competent, 61–80 distinctive, 81–100 category-defining. Panels should include operations, brand, and commercial representation, and record a one-line justification per score to support later audit.

# C Reproducibility

The illustrative results are generated from the archetype inputs in Table 4 via (1) and (2); the segmentation figures use a standardized two-driver projection at $k = 3$. All parameters are stated in-text so results can be reconstructed without access to proprietary data.

# References


[1] Enz, C. and Canina, L. *Revenue Management and RevPAR Benchmarking in Lodging.* Cornell Hospitality Reports.

[2] Kimes, S. E. *The Basics of Yield Management and RevPASH.* Cornell Hospitality Quarterly.

[3] MacQueen, J. (1967). *Some Methods for Classification and Analysis of Multivariate Observations.* Berkeley Symposium on Mathematical Statistics and Probability.

[4] Lloyd, S. P. (1982). *Least Squares Quantization in PCM.* IEEE Transactions on Information Theory.

[5] Pine, B. J. and Gilmore, J. H. (1999). *The Experience Economy.* Harvard Business School Press.

[6] Global Sustainable Tourism / industry benchmarking notes (context).

[7] Themed RevPAR Working Group (2026). *$\theta$-coefficient calibration notes* (internal).