



A to F. Same rating. Same city. Same year.

Eight Porto hotels a guest-rating filter lumps together at nine out of ten — on rates from €207 to €708 a night. The shared rating hides five letter grades of operator quality, and the rate doesn't track them. Every grade below is read from the guests' own words — and traces to a row you can open.

PREPARED BY

Rui Andrade · Hedonic Intelligence

WORKED EXAMPLE

8 Porto premium hotels · anonymised A-H

READ IN FULL

18,550 reviews · 4 platforms · 10 source languages ·
€207–€708 ADR

A category dashboard counts what guests mention. We read what they *mean* — every review, to the bottom.

A rating average and a category-sentiment tool tell you guests talked about breakfast, and tag it amber. They can't tell you which kind of thing it was — and that is the part that moves the asset.

01 An *isolated* complaint — or a *structural* pattern.

02 A guest who *will come back* — or one who *quietly won't*.

03 A recovery that *worked* — or one *buried* in a polite reply.

Those three distinctions are the lens. Each of the 18,550 reviews is read end to end — 100-plus structured fields per review, in its original language — then read again as a set. Every letter grade in the pages that follow is that read, anchored to the verbatim sentence behind it.

What the score *can't* tell you.

A single real review from one hotel in the set. The rating says "a bad stay"; a category tool adds three negative tags. Neither tells you it's a lost repeat guest — on an anniversary, after two failed recoveries — or what to do about it.

WHAT A CATEGORY-SENTIMENT TOOL SEES

"...the room they gave us had no natural light; they kindly swapped it, but the new room smelled strange, had hairs in the bed, and the water pressure cut out. It's a 5-star hotel, at the end of the day, and they didn't really act like it was."

Source: public review · upper-upscale hotel in the set · scored 4 / 10

Category tags returned:

room · negative

cleanliness · negative

F&B · negative

A low score and three tags. They tell you it went wrong — not what it costs, whether they're coming back, or what to fix.

WHAT THE READ SEES IN THE SAME WORDS

CHURN

Critical · will not return

OCCASION

Anniversary trip

REAL DRIVER

5-star promise broken — not isolated issues

LOYALTY

Repeat guest — worse than last stay

RECOVERY

Two attempts, both failed

NOT BLANKET

Several staff praised — it's product + one manager, not the team

"They kindly swapped the room... then offered another cocktail, which I didn't want."

One review, read this far, is a recoverable guest and a precise signal — repeat guest, occasion, brand promise on the line. **The same signal, repeated across the set, is the lever** a category tag can't reach.

Same rating band. *Five letter grades* apart — and the price doesn't track the grade.

Five operator-quality reads, graded A–F against the cohort, beside the rate each asset actually charges. The standard rating average flattens all five into one number — and hides the spread an underwriter is paid to see.

HOTEL	ADR	NLS	OPERATOR CULTURE	SERVICE RECOVERY	REPEAT GUESTS	LOYALTY TIER	OVERALL
Hotel B	€679	+50	A	A	B	—	A
Hotel A	€520	+45	A	A	A	—	A
Hotel G	€249	+33	B	C	C	—	B
Hotel C	€268	+37	C	C	C	C	C
Hotel D	€248	+25	C	D	D	F	D
Hotel F	€207	+35	D	C	C	F	D
Hotel E	€592	+41	C	D	D ↓	—	D
Hotel H	€708	+18	F	F	F	—	F

8 Porto premium hotels (upper-upscale to luxury) · 18,550 reviews · all clustered at ~9/10 guest rating · grades are cohort-relative. **ADR** = average offered nightly rate across collected room offers. **NLS** = net loyalty score (-100 to +100), strict-language read of repeat intent. **Loyalty tier** shown only for the 3 branded properties — 2 fail on elite-tier service; N/A for independents. **D ↓** = declining since an ownership change. Note the top and bottom of the table: **the two highest rates in the set sit at opposite ends of the grade.** Per-cell verbatim and row IDs on the engagement deliverable.

Hotel H — the highest rate in the set, the lowest loyalty.

Charging €708 a night — top of the cohort — on bottom-of-cohort retention. Three independent reads of operator quality converge on one grade, and they put the rate itself in question.

F

Operator culture, service recovery, repeat-guest retention — three separate reads, one letter.

Friction is resolved **3.3% of the time** (12 of 369 friction events; Wilson 95% CI 1.9–5.6%) against a **6.9% cohort baseline** (CI 6.2–7.6%) — roughly half the cohort rate. **65 guests state outright they will not return** — the most of any hotel in the set. This is a culture, not a building.

The decision: ADR €708 operator replacement, or a rate reset. The set's most expensive room is built on its weakest loyalty — the five-year question is whether that rate survives the retention, sized against the asset's NOI on the engagement.

THE VERBATIM BEHIND THE GRADE

EN · theft event · manager refused to call police

"The receptionist was smirking about the theft."

EN · 2025 · recent

"The so-called manager couldn't have been more unhelpful — doesn't care."

PT · repeat-guest cohort

"Sou cliente habitual... infelizmente não conhecem os clientes."

A regular guest — and they don't know their customers.

Each percentage carries a Wilson 95% CI and a row count; every quote is row-cited. Both inline on the engagement deliverable.

Hotel E — the decline a quarterly report hasn't caught yet.

Repeat guests name the operator transition themselves. The read catches the trajectory before the next quarter's numbers land.

—13_{pp}

The last six months sit 13 points of net loyalty below the property's prior baseline — strict NLS, 154 recent reviews — a slide a rating average hasn't surfaced yet.

Repeat guests at an acquisition target are the leading-edge signal — and they date the turn in their own words. **The read tracks prior-versus-current per aspect, before any operating report you would see in diligence.**

The decision: ADR €592 pre-acquisition diligence and management-contract renewal. The engagement sizes the operator-quality discount against the bid envelope — before the slide reaches the rating.

THE VERBATIM BEHIND THE READ

EN · 2024 · row-cited

"The new owner of just 9 months... lacks all nuance and love of a good hospitalitarian."

EN · repeat guest

"The breakfast has nothing more to do with what it was a year ago, before the new owners. Go back to your old concept."

JA · repeat guest · six-year comparator

"6年前に滞在した時から考えると朝ごはんがよくなかった"

Compared with my stay six years ago, the breakfast was not good.

Repeat guests identified by comparator language — "last time", "a year ago", "first stay" — then read prior-vs-current per aspect. Window and n labelled on every shift claim.

The brand's best guests, documenting the brand's failure.

At the cohort's two weakest branded hotels, named elite-tier members record brand-standard breach in writing — the exposure a franchisor can act on, and a bidder should price.

18+

Self-identified elite-tier members recording a negative experience at the two weakest branded hotels — on the public record. A deliberately conservative count, concentrated at the weaker of the two.

Franchise contracts carry brand-standard obligations. Documented failure on the highest-revenue guest cohort is flag risk. **The read lifts each elite-tier breach straight from public reviews** — the reviewer states their own status — with no operator data needed.

THE VERBATIM BEHIND THE RISK

EN · Hotel D · top-tier loyalty member

"In 36 hours in the hotel I experienced more than 20 issues."

EN · Hotel F · top-tier loyalty member

"Disorganization is the name of the game here."

Elite tier read from the reviewer's own stated status, cross-checked against the property's franchise brand. Brand names withheld from this artifact; named on the engagement deliverable.

The decision: price franchise-fee-loss exposure into the bid — elite-tier-denial rate against management-contract value — pre-acquisition, and ahead of the franchisor's own audit on held assets.

One worked example of thirty — then the rest.

The reads so far drilled four findings. Here is a fifth, worked, then nine more — each mapping to a different seat on the asset-side committee. The deliverable carries whichever the data warrants.

+44–60 pp

Reviews that name a staff member by name are 44 to 60 points more likely to carry explicit advocacy than reviews that don't — every hotel in the set.

When guests rave, they rave by name. Advocacy concentrates around a handful of named individuals — so the read shows where the goodwill actually sits, and where the concentration risk is before it walks out the door. A read no rating average can produce: it resolves every name out of the prose and ties it to loyalty. **The decision: know who the guests come back for — and protect them deliberately.**

<p>Is the operator delivering the experience the management agreement specifies?</p> <p>ASSET MANAGER · OWNER</p>	<p>Since the last renovation, has guest experience actually moved — and on which aspects?</p> <p>ASSET MANAGER · CAPEX</p>	<p>Which assets in the book are drifting — before the quarterly review flags them?</p> <p>PORTFOLIO · IC</p>
<p>Which aspect, fixed, would move the asset's standing the most per euro?</p> <p>ASSET MANAGER · CAPEX</p>	<p>Which guest segment is quietly leaving — and what is it worth in repeat revenue?</p> <p>ASSET MANAGER · REVENUE</p>	<p>Where is the asset leaving rate on the table — paid vs expected?</p> <p>REVENUE · ASSET MANAGER</p>
<p>How does the asset read against its named peer cohort — not its own history?</p> <p>INVESTMENT COMMITTEE</p>	<p>Is the asset delivering the position it is marketed and priced at?</p> <p>OWNER · ASSET STRATEGY</p>	<p>Which rival just changed hands — and is slipping fastest right now?</p> <p>ACQUISITIONS · IC</p>

Ten of more than thirty. Each carries its own sample-size floor. Alongside them, the read surfaces the patterns no fixed question knew to ask.

An honest perimeter.

WHERE HEDONIC WORKS

Periodic, by design. Continuous monitoring of your own portfolio is your operator platform's job. Hedonic is the read your platform can't run — on targets, on rivals, on operator-transition deltas.

No operator cooperation. No PMS feed, no NDA, no site visit. The read works on any property with public review volume.

The cohort is a worked example. The same read runs on your own book, at the cadence your committee uses.

WHAT REVIEWS DON'T CAPTURE — AND WHAT WE WON'T CLAIM

Capex cost modelling. We name the lever and the decision; the cost engineer sizes the build.

Whether your team executes. We read what guests noticed; execution is your call.

Small-number claims. A finding resting on a handful of reviews is flagged as directional, never sold as fact — every shift claim carries its window and its n.

01 · INPUT

Public *review prose*

Multilingual, across four platforms — read in the original language, no translation step.

02 · EXTRACTION

→ *100+ structured fields* per review

Schema-checked per review — aspects, friction root cause, named staff, recovery, journey stage, churn driver — each with an extraction-confidence score; low-confidence reads are flagged, not counted.

03 · READ

→ A cohort-relative *grade*

The reads combine into letter grades — each finding closing in a decision the right seat can act on.

The same read, *your portfolio*.

278,000+

reviews read across 119 hotels in the production pipeline.

100+

structured fields extracted per review, read in the native language.

Your acquisition pipeline, your held assets, your operator-transition deltas — the same read that graded this cohort runs on each, closing in the decisions you are already trying to make: operator selection and replacement, franchise-fee risk, repeat-guest diligence, rate-vs-retention, source-market reallocation.

Name the asset. The read returns as a single PDF — the cohort-relative grade, the verbatim behind it, the decision behind each finding, and a labelled confidence band on every claim.

Deliver-first · NET-30 · no contract. First engagements via two-asset pilot: €4 000 flat, 14 days — and you are the judge: if the read doesn't surface something your operator's reporting hasn't, you don't pay. Standard reads and recurring portfolio mandates available — engagement scope, DPA and sub-processor register on request.