Revenue Management 101 for Engaged Hotel Owners

By Vikram Aditya Singh · Oct 14, 2026

Before you start

This book is for the owner who reads his own P&L and wants to sit in the revenue meeting as an equal. You won't be asked to run a forecast. You will be asked to understand one well enough to know when it's wrong.

I've kept the maths in. Every chapter carries the formula your revenue manager or your RMS actually uses, because an owner who can't see the formula can only argue about the answer. Where the textbook and the course versions disagree, I say which one is right and why.

Each chapter follows the same order. It opens with the question an owner should be asking and a short answer. Then it shows how the method works, with the formula set out in full. Then it runs the method on the specimen hotel, using the numbers in the workbook. It closes with the questions to put to your team, the red flags I've learned to look for, and the workbook sheet that carries the numbers.

The specimen hotel

Most of the numbers come from one hotel: 500 keys, full service, in a city with a midweek conference peak and a winter event season. It runs at 74% occupancy and a $150 ADR, takes $20.3m in rooms revenue and $35.8m in total, and makes $7.9m of EBITDA. Its comp set is five hotels. Chapter 23 adds a 120-key resort.

The figures are illustrative but they tie to each other. Read $ as your own currency. Everything is net of taxes and service charges.

Using the workbook

The workbook is called Revenue Management 101 for Engaged Hotel Owners, The Workbook. The second sheet, RM101 Map, lists every chapter with a link to its sheet. Row 4 of each sheet names its chapter and links back to the map.

Gold cells with blue figures are inputs. Black cells are formulas. Green cells are linked from another sheet. Put your own PMS, STR and channel numbers into the gold cells and the whole chain recomputes. Don't make a decision on the specimen figures.

Six words you'll need from page one

Term	What it means
Occupancy	Rooms sold ÷ rooms available
ADR	Rooms revenue ÷ rooms sold
RevPAR	Rooms revenue ÷ rooms available, or occupancy × ADR
Net ADR	ADR after commission and every other cost of acquiring the booking
Displacement	The business you turn away when you say yes to something else
Bid price	What the last room on a night is worth; the lowest rate you should accept for it

Part 1 · Foundations

1 · Where revenue management works

Is my hotel the kind of business this works for, and where will the money come from?

Almost certainly yes for the rooms. Less so, for now, in the restaurants and the spa. Revenue management decides, for every room night, which guest to accept, at what price and through which channel, so that the owner makes the most profit from capacity that perishes at midnight.

Sheryl Kimes set out six conditions under which this pays. Score each from 1 (absent) to 5 (strong).

Condition	Specimen score	Why
Relatively fixed capacity	5	500 keys and 32,000 sq ft of event space
High fixed cost, low variable cost	4	Variable cost is about a sixth of ADR
Perishable inventory	5	An unsold night is gone
Reservations made in advance	4	Groups book months out; OTA guests inside 14 days
Variable demand	4	Midweek peaks, a winter season, a summer trough
Known price sensitivity by segment	3	Fences are thin on the OTAs

The specimen scores 25 out of 30. The weakest condition is the one to fix first, and here it's price sensitivity: the hotel doesn't yet know how each segment reacts to a price move. Chapter 11 measures it.

Kimes and Chase then gave us two levers: how much you vary price, and how well you control duration (how long a guest occupies the capacity). Put each revenue centre on that grid. Rooms already have many prices and good duration control, so they need refining: length-of-stay controls and more fenced rates. The ballroom and the spa control duration well but charge one price, so add day-part, day-of-week and lead-time prices. The restaurants and bars control neither; fix duration first, then price. The course puts the expected uplift at 2 to 5% of revenue.

Ask your team

- Which of the six conditions is weakest here, and what are we doing about it?
- Which quadrant is each outlet in, and which lever are we pulling this year?

Red flags

- Revenue management is talked about only for rooms in a hotel where rooms earn less than 60% of revenue.
- Nobody can say what a 1% rise in rate does to volume in each segment.

Workbook: 20 Conditions & Quadrants.

2 · The scoreboard

Which number tells me the hotel is doing well?

No single one. RevPAR tells you how well the rooms sell. Net RevPAR tells you what you keep after paying to win the booking. TRevPAR tells you what each room brings in across the whole hotel, and GOPPAR tells you what's left as profit. Read all four together.

$$\text{RevPAR} = \text{Occupancy} \times \text{ADR} = \frac{\text{Rooms revenue}}{\text{Rooms available}}$$

$$\text{Net RevPAR} = \frac{\text{Rooms revenue} - \text{commissions} - \text{other acquisition costs}}{\text{Rooms available}}$$

$$\text{TRevPAR} = \frac{\text{Total operating revenue}}{\text{Rooms available}} \qquad \text{GOPPAR} = \frac{\text{Gross operating profit}}{\text{Rooms available}}$$

Rooms available means every room, every night, whether you sold it or not. That's the point. A hotel can push occupancy up with cheap OTA business and RevPAR will look fine while Net RevPAR and GOPPAR fall.

On the specimen, 74% occupancy at a $150.02 ADR gives a RevPAR of $111.01. Acquisition costs take 6.2% of rooms revenue, so Net RevPAR is $104.14. Total revenue of $35.76m over 182,500 available room nights gives a TRevPAR of $195.95, which is 1.77 times RevPAR. F&B and banquets are 36% of the hotel's revenue. EBITDA per available room is $43.29, a 22.1% margin. Where only EBITDA is reported, as here, it stands in for GOPPAR; use GOPPAR from the USALI P&L when you have it.

The number I'd look at hardest is flow-through: how much of each extra unit of revenue reached EBITDA.

$$\text{Flow-through} = \frac{\Delta\,\text{EBITDA}}{\Delta\,\text{Revenue}}$$

Revenue rose $1.86m on last year and EBITDA rose $0.5m, so flow-through was 26.9%. A full-service hotel should be keeping 35 to 45% of new revenue. At 27%, new revenue is being bought with cost. Fix the cost base before you buy more volume.

Ask your team

- What's our Net RevPAR, and how has it moved against RevPAR this year?
- What was flow-through on the year's growth, and which department absorbed it?

Red flags

- The monthly pack leads with occupancy.
- RevPAR is up and GOPPAR is flat or down.

Workbook: Inputs, 01 KPIs & Indices, 13 Total Revenue.

3 · Fair share and the comp set

Are we beating the market, or just riding it?

You find out by comparing the hotel with the four to seven hotels a guest would genuinely shop against it. STR reports three indices, and 100 means you're taking exactly your share.

```
\text{MPI} = \frac{\text{Occ}_{\text{hotel}}}{\text{Occ}_{\text{comp set}}}
\times 100 \qquad \text{ARI} = \frac{\text{ADR}_{\text{hotel}}}
{\text{ADR}_{\text{comp set}}} \times 100 \qquad \text{RGI} =
\frac{\text{RevPAR}_{\text{hotel}}}{\text{RevPAR}_{\text{comp set}}} \times
100
```

RGI is MPI × ARI ÷ 100, so you can always see whether a gap comes from volume or from price.

The specimen runs 74% against the comp set's 76%, an MPI of 97.4. Its ADR is $150.02 against $173.75, an ARI of 86.3. RGI is 84.1. So the hotel fills close to its share of rooms and sells them 14% cheaper. Its problem is price. At an RGI of 100 the hotel would take $24.10m in rooms revenue; it takes $20.26m. That's a $3.84m gap a year.

The index is only as honest as the comp set. The rules STR and Demand360 apply are a good floor: four to seven hotels, none with more than 40% of the set's rooms, and at least two owners. I score candidates on product match and demand overlap (35% weight each), location and review score (15% each), and drop anything below about 3. Review the set every year. Never choose it to flatter the index.

STR looks back. To look forward, compare your share of the comp set's rooms on the books with your share of its supply. The specimen has 16.1% of the comp set's rooms, so any arrival week where it holds less than 16.1% of the set's bookings is a week to act on now.

Ask your team

- Is our RGI gap a volume problem or a price problem? Show me MPI and ARI separately.
- When was the comp set last scored, and who would a guest really compare us with?

- Which weeks ahead are we pacing below our forward fair share?

Red flags

- The comp set hasn't changed in five years, or it includes a hotel nobody on the team thinks of as a competitor.
- RGI is reported without MPI and ARI beside it.
- An RGI below 100 two weeks running with no action taken.

Workbook: 01 KPIs & Indices, 29 Demand & Comp Reports, 30 Rate Shopping (section A).

Part 2 · Demand

4 · Segmentation

Who fills my hotel, and which of them is worth the most?

You need three answers for every room night, and they mustn't be mixed. The segment is who booked and how many rooms. The rate category is the terms they booked on. The channel is the road the booking came by. A corporate guest at a negotiated rate who books through the GDS is one guest with three labels. Put the channel into the segment code and you lose the ability to say what the guest is worth apart from what it cost to get him.

USALI 12 gives the standard.

Group	Segments	Rule
Transient	Retail, Discount, Negotiated, Qualified, Wholesale	Fewer than 10 rooms a night
Group	Corporate, Association, Government, Tour and wholesale, SMERF	10 or more rooms a night under one contract
Contract	Contract	A fixed block for more than 30 days, such as airline crew

Three rules move real money. Record revenue gross or net consistently, and know which: a wholesale rate is net, an OTA booking is often loaded gross. When two discounts stack, the deeper offer decides the segment, so a loyalty member on a flash sale is Discount. And revenue lands on the dates the guest stays, not the date the contract was signed.

Rank segments on what they leave you per room night, rooms and F&B together.

```
\text{Net ADR} = \text{ADR} \times (1 - \text{acquisition cost \%})
```

```
\text{Contribution per RN} = \text{Net ADR} \times m_{\text{rooms}} +
\text{F\&B per RN} \times m_{\text{F\&B}}
```

The specimen uses 80% as the rooms margin and 35% for F&B.

Segment (PMS code)	Share of RNs	Net ADR ($)	Contribution per RN ($)	Rank
Wedding and social groups	8%	165.38	228.55	1
MICE groups	14%	149.45	161.12	2
Brand.com and voice (BAR)	12%	177.00	154.73	3
Corporate negotiated (LRA)	24%	150.35	131.66	4
Online travel agencies	18%	135.30	117.87	5
Long stay and embassy	3%	131.25	111.56	6
Wholesale and tour	7%	117.50	101.88	7
Loyalty redemption	4%	112.50	100.94	8
Airline crew	10%	77.50	65.94	9

The wedding guest is worth $229 a night to this hotel. The BAR guest is worth $155, even though he pays the highest room rate. That's the banquet spend. A system that protects rooms on ADR alone will give the wedding's rooms away.

The specimen's PMS codes mix segment and channel, as most hotels' do. Column L of sheet 02 maps each one to USALI 12.

Ask your team

- Are segment, rate category and channel three separate fields in our PMS?
- Rank our segments on contribution per room night, F&B included. Does our protection follow that order?

Red flags

- A segment called “OTA” or “Website”.
- Group revenue booked to the month the contract was signed.
- More than 5% of room nights in “Other”.

Workbook: 02 Segments.

5 · The booking curve, pickup and pace

Are we ahead or behind for next Tuesday, and how do we know?

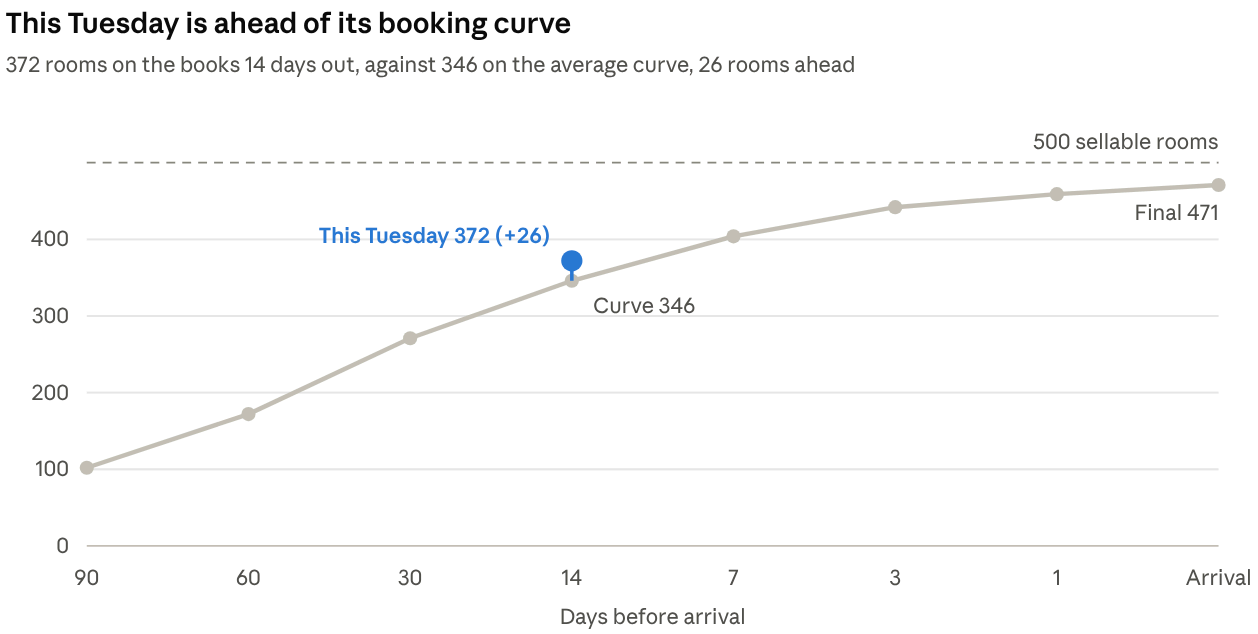
By comparing what's on the books today with what was on the books at the same point before past Tuesdays. Each night of the week fills along its own curve. You build it from history: rooms on the books at 90, 60, 30, 14, 7, 3 and 1 days out, for the same night of the week in the same season.

On the specimen's eight past Tuesdays, 14 days out the hotel had 74% of its final rooms on the books, and it picked up another 124 rooms on average to arrival. Pickup turns that into a forecast in two ways.

```
F_{\text{additive}} = \text{OTB}_d + \overline{\text{pickup}}_d \qquad
F_{\text{multiplicative}} = \text{OTB}_d \times
\overline{\left(\frac{\text{final}}{\text{OTB}_d}\right)}
```

The additive method assumes the same number of rooms will still come; the multiplicative one assumes the same proportion. Additive is steadier when the hotel is near capacity. Multiplicative is right when the base varies a lot, which is why group pickup (chapter 17) always uses it.

This Tuesday, 14 days out, the hotel holds 372 rooms. The curve says it should hold 346, so pace against the curve is +7.4%. Same time last year it held 338, so pace against last year is +10.1%. The additive forecast is 496 and the multiplicative 506; blended, 501 rooms against 500 sellable. The hotel is ahead. Lift BAR a tier, close the discounts and put a two-night minimum on the shoulder nights.



Workbook sheet 03 Forecast · rooms on the books, average of eight past Tuesdays in season, against this Tuesday

Ask your team

- What's on the books for each of the next 90 nights against the curve and against last year?
- Which forecast method do we use at each horizon, and why?

Red flags

- Pace is reported only against last year. A night that was badly managed last year is a poor benchmark.
- One curve for every night of the week.

Workbook: 03 Forecast.

6 · Unconstraining

If we sold out, how many more could we have sold?

A sold-out night records the cap, not the demand. If you forecast from what you sold, the hotel will underprice its best nights for ever, because the history never shows how many guests were turned away.

Start with the denial log: every enquiry the reservations team, the website and the GDS turned down because the room or the rate was closed. Only some of those guests would really have booked, so apply a conversion share.

$$\text{Unconstrained demand} = \text{Rooms sold} + s \times \text{Denials} \qquad \text{Spill} = \text{Unconstrained} - \text{Rooms sold}$$

The specimen assumes 55% of denials would have booked. Over eight peak nights it sold 499 rooms on average; unconstrained demand was 532. That's 6.6% more demand than the history shows, a spill of 33 rooms a night, with a standard deviation of 21.5 rooms. Those 33 rooms are the room the hotel has for a price rise on those nights.

Projection detruncation and the EM algorithm refine this (Weatherford and Pölt, 2002), and a good RMS runs them. Neither works without denials recorded in the first place.

Ask your team

- Do we log denials and regrets on every channel, including the booking engine?
- What share of denials do we assume would have booked, and where does that number come from?

Red flags

- No denial data, or denials only from the phone.
- The forecast for a night that sold out last year equals last year's rooms sold.

Workbook: 04 Unconstraining.

7 · Forecast accuracy

How good is our forecast?

Ask for four measures by horizon, not one. Each answers something different.

$$\text{MAD} = \frac{1}{n}\sum |A_t - F_t| \qquad \text{MAPE} = \frac{1}{n}\sum \frac{|A_t - F_t|}{A_t} \qquad \text{WAPE} = \frac{\sum |A_t - F_t|}{\sum A_t}$$

$$\text{Bias} = \frac{1}{n}\sum (A_t - F_t) \qquad \text{Tracking signal} = \frac{\sum (A_t - F_t)}{\text{MAD}}$$

A is actual and F is forecast. MAD is the average miss in rooms. MAPE weights a 262-room night as heavily as a 462-room night; WAPE weights by volume, so it's the better headline. Bias tells you which way the forecast leans. The tracking signal tells you whether that lean is luck or method: beyond ±4, the method is biased.

Horizon	MAD (rooms)	WAPE	Bias (rooms)	Tracking signal
30 days out	35.3	8.9%	−15.5	−4.4
14 days out	20.1	5.1%	−8.7	−4.3
7 days out	9.7	2.4%	−4.1	−4.2
1 day out	2.9	0.7%	−1.1	−3.8

A WAPE of 5% at 14 days looks good. But every horizon is negative and the tracking signal is beyond −4 at 30, 14 and 7 days. The forecast is consistently too high. On the Tuesday test in sheet 03 the MAPE is only 2.7%, and the tracking signal is −6.0. A hotel with that forecast closes discounts too early on soft nights and holds rooms for guests who never come.

Ask your team

- Show me MAD, WAPE and bias at 30, 14, 7 and 1 day out, for the last quarter.
- What's the tracking signal, and if it's outside ±4, what have we changed?

Red flags

- Accuracy reported as one MAPE with no bias.
- Forecast accuracy measured only on the day before arrival.

Workbook: 03 Forecast (rows 32 to 43), 21 Demand Control (rows 29 to 46).

8 · Hot, warm and cold nights

What should we be doing differently on each night in the next fortnight?

Sort the nights into zones by forecast occupancy and give each zone a set of instructions. That's the demand-control chart, and it's the one page I'd put in front of a GM every morning.

```
\text{Forecast occupancy}_n = \frac{\text{Unconstrained room-night forecast}_n}{\text{Sellable rooms}}
```

The specimen calls a night hot at 100% or more, warm from 80% and cold below. Each zone carries a minimum BAR tier from the ladder in chapter 10.

Zone	Trigger	Minimum BAR (Deluxe King)	Instruction	Nights in the next 14
Hot	⩾ 100%	BAR 7 ($228)	Close below BAR 7; two-night minimum if the shoulder is soft	3
Warm	80–99%	BAR 4 ($181)	Close below BAR 4; hold groups to displacement	5
Cold	< 80%	BAR 1 ($144)	All rates open; push fenced offers	6

Use the room-night forecast, arrivals plus stayovers. The course divides arrivals by capacity, which understates any night with a lot of stayovers. On the specimen, next Tuesday forecasts at 104% and is hot. The Sunday before it forecasts at 58%.

Ask your team

- Which nights in the next fortnight are hot, warm and cold, and what's open on each?
- When did we last move the triggers, and on what evidence?

Red flags

- The same rates open on a hot Tuesday and a cold Sunday.
- Zones set on arrivals rather than room nights.

Workbook: 21 Demand Control.

Part 3 · Price

9 · Price response and elasticity

What's the right price for tonight?

The one that earns the most contribution from the rooms you have left. That depends on how many guests you lose for each dollar you add, and on whether you have more demand than rooms.

Start with a straight-line demand curve for tonight's transient business. D is rooms demanded, P is BAR, c is the variable cost of an occupied room and C is the transient rooms left to sell.

$$D = a - bP \qquad \pi = (P - c)(a - bP)$$

Set the derivative of contribution to zero and you get the best price when rooms aren't the limit. If that price would sell more rooms than you have, the cap sets the price instead.

$$P^{*} = \frac{1}{2}\left(\frac{a}{b} + c\right) \qquad P_{\text{cap}} = \frac{a - C}{b} \qquad \varepsilon = -\frac{bP}{D}$$

On the specimen, a is 520 rooms, b is 1.76 rooms lost per dollar, c is $23.75 and 180 transient rooms are left. P* is $159.60, but at that price 239 guests would want a room. The cap binds, so the price to post is $193.18 and it sells exactly 180 rooms. Elasticity at that price is −1.89.

Two things follow. Revenue peaks where elasticity is −1; contribution peaks at a higher price, because each extra guest costs $23.75 to serve. And when demand exceeds rooms, the price comes from the cap. A revenue manager who prices a sell-out night from the elasticity alone leaves money on the table.

Ask your team

- On a night we expect to sell out, how did we set the price? From demand, or from the comp set?
- What's our variable cost per occupied room, and is it in the pricing?

Red flags

- The same BAR posted on a night with 33 rooms of spill as on a night with none.
- Prices set to maximise revenue with no cost in the calculation.

Workbook: 05 Price Response.

10 · The BAR ladder and attributes

What are BAR 1, BAR 2 and BAR 3, and who decides which one we sell?

BAR is the best available rate: the lowest unrestricted rate a guest can book tonight. A ladder of BAR tiers lets the hotel move price in planned steps, triggered by forecast occupancy, rather than inventing a new number every day. Every other rate category is set off the tier in force.

$$\text{BAR}_n = \text{BAR}_1 \times (1 + s)^{\,n-1}$$

The specimen has eight tiers with an 8% step, starting at $143.75 for a Deluxe King.

Tier	Opens at forecast occupancy	Deluxe King ($)	Club ($)	Junior Suite ($)	Suite ($)
BAR 1	0%	144	187	266	374
BAR 2	55%	155	202	287	403
BAR 3	62%	168	218	311	437
BAR 4	68%	181	235	335	471
BAR 5	74%	196	255	363	510
BAR 6	80%	211	274	390	549
BAR 7	86%	228	296	422	593
BAR 8	92%	246	320	455	640

Room types ride on the tier: Club at +30%, Junior Suite at +85%, Suite at +160%. Fenced rates ride on it too. Advance purchase at 15% off, a non-refundable rate, a two-night package: each is a discount off the tier in force, with a fence that keeps the BAR guest from taking it. A fence has to be something the BAR guest won't accept, such as paying 21 days ahead.

The hotel can also sell what the guest values beyond the room type. At the specimen's assumed take rates, five attributes (high floor with a view, guaranteed early check-in, late check-out, lounge access, a Junior Suite upgrade at check-in) bring $544,000 a year.

$$\text{Attribute revenue} = \text{price} \times \text{take rate} \times \text{occupied room nights}$$

Ask your team

- Show me the ladder, the triggers, and which tier is loaded for each of the next 30 nights.
- Which fences do we use, and how often does a BAR guest get round them?

Red flags

- Twenty rate codes and no ladder.
- The suite category sells out first, every peak. It's underpriced.

Workbook: 06 BAR Ladder.

11 · Should we move the price?

The revenue manager wants to cut BAR by 10% to fill the hotel. Should I agree?

Only if volume will rise by more than the break-even amount. Work that out before anyone looks at elasticity. CM is contribution margin per room night at today's price; ΔP is the price change and ΔVC any change in variable cost.

```
\%\Delta Q_{\text{break-even}} = \frac{-(\Delta P - \Delta VC)}{CM + \Delta P - \Delta VC}
```

Then measure elasticity from your own STR data, so you know what volume change to expect. Regress the log of your fair-share index on the log of your ADR index, week by week.

```
\ln(\text{FSI}_t) = \alpha + E \ln(\text{ARI}_t) \qquad \text{FSI}_t = \frac{\text{Rooms sold}_t}{\text{Fair share}_t} \qquad \frac{Q_1}{Q_0} = \left(\frac{P_1}{P_0}\right)^{E}
```

On the specimen's brand.com BAR of $184.38, CM is $160.63. Twenty weeks of data give E = −0.72 with an R^2 of 0.66.

Test	Break-even volume change	Elasticity predicts	Verdict	Contribution a night at 180 rooms
Cut 10%	+14.5%	+7.8%	Cut destroys contribution	−$1,676
Rise 8%	−8.4%	−5.4%	Rise pays	+$961

The cut needs 14.5% more rooms and elasticity says it'll get 7.8%. The rise can afford to lose 8.4% of rooms and is expected to lose 5.4%. When the comp set raises its BAR by $8.75 and you hold, you need 5.4% more volume to justify holding.

Inelastic demand gives you room to test a rise. The break-even line, which carries your margin, decides whether you keep it. If the elasticity estimate comes out positive, or the R^2 is under about 0.3, don't use it; the data are telling you something else moved.

Ask your team

- What's the break-even volume change for this price move?
- What's our measured elasticity, over how many weeks, and what's the R^2?

Red flags

- A price cut justified by "we need the occupancy".
- An elasticity number nobody can trace to data.

Workbook: 22 Price Test.

12 · Rate shopping and parity

The comp set dropped its rate. Should we follow?

Not automatically. A rate shopper tells you where you stand. Your demand tells you what to charge. Read your position as an index and set it beside the demand zone.

```
\text{Price index}_n = \frac{\text{Own BAR}_n}{\text{Comp-set median BAR}_n}
\times 100
```

Night	Own BAR ($)	Comp median ($)	Index	Zone	Read
Wed +1	181	174	104	Warm	Consistent
Mon +6	181	198	92	Warm	Consistent
Tue +7	228	245	93	Hot	Underpriced on a hot night: lift
Sun +5	144	130	111	Cold	Above market on a cold night: fence an offer

The shop is only useful if it compares like with like. Same room category, two adults, the refundable BAR and the lowest non-refundable rate, taxes normalised, public and member rates kept apart, desktop and mobile both. The specimen's specification (seven hotels, six channels, two lengths of stay, 90 nights daily and the rest of the year weekly) runs to 326,000 shop calls a month. Know what you're paying for.

Parity is the other use. When an OTA shows a lower rate than brand.com, some direct guests switch and you pay commission on a booking you'd have had anyway. On the specimen, 18% of shops find an OTA undercutting; of 16,400 direct-intent room nights a year, that exposes about 2,950, and 30% of those switch. At $41.81 more acquisition cost a night, the leak is $37,000 a year. Look for the source before you fight the OTA. It's usually a wholesale or bed-bank rate being resold.

Ask your team

- What's our price index tonight and for the next hot night, and what zone is each in?
- Where did the last parity breach come from?

Red flags

- The rate follows the comp set down on a night the hotel forecasts to sell out.
- The shop compares our refundable rate with their non-refundable one.

Workbook: 30 Rate Shopping.

13 · Reputation as pricing power

Is the review score a marketing number or a revenue number?

It's a revenue number. Chris Anderson's Cornell study (2012) found that a one-point rise on a five-point review scale let a hotel raise its price by 11.2% and keep its occupancy.

$$\text{Price headroom} = \Delta\,\text{score}_{5\text{-point}} \times 11.2\% \qquad \text{Revenue at stake} = \text{Rooms revenue} \times \text{headroom}$$

The specimen scores 7.9 out of 10 on its leading OTA and targets 8.7. That's 0.4 on a five-point scale, 4.5% of price headroom, and $908,000 of rooms revenue a year. Chapter 25 counts only half of it, because the score has to move first and the rate follows.

The order matters. Value for money and the age of the product drag scores down. Fix the product and how the hotel answers reviews, then raise the rate.

Ask your team

- What's our score on each major site, and what are guests marking us down for?
- When we last raised rates, did the value-for-money score move?

Red flags

- Rates raised while the score is falling.
- Nobody owns review responses.

Workbook: 17 Reputation.

Part 4 · Inventory

14 · Nested protection

How many rooms should we hold back for guests who'll pay more later?

Hold a room for the higher rate as long as the chance of selling it at that rate, times the rate, is worth more than the lower rate you could take now. That's Littlewood's rule (1972), and every RMS still starts from it.

For two classes, with r_1 the higher net rate and D_1 its demand, protect y_1 rooms where:

```
r_1 \, P(D_1 > y_1) = r_2 \quad\Rightarrow\quad y_1 = \mu_1 + \sigma_1 \,
\Phi^{-1}\!\left(1 - \frac{r_2}{r_1}\right)
```

Φ^{-1} is the inverse of the normal distribution. With more than two classes, Belobaba's EMSR-b (1992) treats all the classes above each line as one: add their means, add their variances, and weight their rates by demand.

```
\bar{\mu}_j = \sum_{i \le j} \mu_i \qquad \bar{\sigma}_j = \sqrt{\sum_{i \le
j} \sigma_i^2} \qquad \bar{r}_j = \frac{\sum_{i \le j} r_i \mu_i}
{\bar{\mu}_j} \qquad y_j = \bar{\mu}_j + \bar{\sigma}_j \, \Phi^{-1}\!\left(1
- \frac{r_{j+1}}{\bar{r}_j}\right)
```

The booking limit for the next class down is the rooms available less y_j. Limits nest: a higher class can always take a lower class's room, never the other way round.

On a peak Tuesday the specimen has 500 rooms, 243 of them already committed to LRA corporate, crew and group blocks. That leaves 257 for the dynamic classes.

Class	Net rate ($)	Mean demand	Std dev	Protect for this class and above	Limit for the next class
BAR flexible, brand.com	177.00	70	20	49	208

Class	Net rate ($)	Mean demand	Std dev	Protect for this class and above	Limit for the next class
Advance purchase (−15%)	150.50	60	18	105	152
OTA, net of 18%	135.25	110	30	209	48
Wholesale, net	117.50	80	25		

Wholesale can have at most 48 rooms on that Tuesday, whatever it asks for. Run this for each room type, because each has its own capacity and its own ladder.

The spreadsheet that ships with the EHL course gets this wrong. It sets $r_1 P(D_1 > y)$ equal to $r_2 P(D_2 > y)$. Littlewood compares $r_1 P(D_1 > y)$ with r_2 itself, because the lower-rate booking is in hand. On the course's own three-class example, the error lets the $100 class book 71 rooms when it should have 52. Over a 400,000-night simulation that costs $149 a night.

Ask your team

- How many rooms are protected for BAR on each of the next five peak nights, and by room type?
- Are our limits nested?

Red flags

- Allocations by channel (“the OTAs get 60 rooms”) in place of nested limits by net rate.
- Protection set on gross rates, so the OTA looks as valuable as brand.com.

Workbook: 07 Protection, 27 EHL Cases (section B).

15 · Bid prices and length of stay

Why did we turn down a one-night booking on a Tuesday when we had rooms?

Because the room was worth more to a guest staying three nights through Tuesday. A bid price is the value of the last room on a night. A stay is worth taking if what it pays covers the bid prices of every night it uses.

```
\text{BP}_n \approx \bar{r} \times P(D_n > \text{rooms left}_n) \qquad
\text{Accept if} \quad \sum_{n \in \text{stay}} r \ge \sum_{n \in
\text{stay}} \text{BP}_n
```

On the specimen, with $155 as the weighted net rate of the demand it could displace, the bid prices for the week are $10 on Monday, $154 on Tuesday, $151 on Wednesday and $4 on Thursday.

Request	Stay revenue ($)	Sum of bid prices ($)	Decision
1 night Tuesday, OTA	135.25	154	Decline or reprice
3 nights Mon–Wed, OTA	405.75	315	Accept
2 nights Tue–Wed, advance purchase	301.00	305	Decline or reprice
1 night Monday, OTA	135.25	10	Accept

The heuristic is good enough to explain the idea. A proper RMS solves the week as one linear programme, so that a stay straddling several nights is valued against all of them at once. x_j is the number of requests of product j accepted, D_j its demand and C_n the rooms left on night n.

```
\max \sum_j r_j \, \text{LOS}_j \, x_j \quad \text{subject to} \quad \sum_{j \ni n} x_j \le C_n \;\; \forall n, \qquad 0 \le x_j \le D_j
```

The bid price of a night is the shadow price of its capacity constraint: what the hotel would lose if it had one room fewer. On the specimen's week the LP gives $93.50 on Monday, $177 on Tuesday and Wednesday, $150.50 on Thursday, $135.25 on Friday and nothing at the weekend.

That produces controls no single rule could. Tuesday and Wednesday arrivals sell BAR only. A Thursday arrival on an OTA rate is closed for one or two nights (worth $135.25 and $270.50 against bid prices of $150.50 and $285.75) and open for three ($405.75 against $285.75), because the third night reaches the empty weekend. A blanket two-night minimum gets that wrong.

Ask your team

- What's the bid price for each of the next 14 nights?
- Which length-of-stay restrictions are loaded, and did they come from bid prices or from habit?

Red flags

- A two-night minimum that sits on every Tuesday all year.
- Single-night requests accepted on a peak night while multi-night requests through it are turned away.

Workbook: 08 Bid Price, 26 SDM Bid Price LP.

16 · Overbooking

How many rooms should we oversell, and what happens when we get it wrong?

Oversell until the next booking's expected gain from filling a no-show equals its expected cost of walking a guest. Each reservation shows up with probability 1 − p, where p is the no-show and same-day cancellation rate, so arrivals follow a binomial distribution. Take n reservations for C rooms, with m the contribution of a filled room and W the cost of a walk:

```
X \sim \text{Bin}(n, 1-p) \qquad \text{EV}(n) = m \, E[\min(X, C)] - W \, E[(X - C)^{+}]
```

The specimen has a 5.5% no-show rate, a $123.57 contribution on a filled room and a $275 cost of walking a guest. The best number to accept is 526: 26 rooms over, with 0.9 walks expected a night. Against no overbooking that's worth $2,675 a night. The hotel currently stops at 512. Moving to 526 on its 140 sell-out nights is worth $178,000 a year.

The same decision, room by room, is the critical fractile. Overbook the k-th room as long as:

```
P(\text{no-shows} \ge k) \ge \frac{W}{W + E}
```

where E is the contribution lost on an empty room. With W at $275 and E at $123.57 the ratio is 0.69. On the specimen's 40-night no-show table, that says overbook by 6 on a peak night. Build one table for each day of the week and season.

Sometimes the walk can't be priced, or a contract says nobody gets walked. Then cap how often a walk can happen. Toh and DeKay (2002) give the authorised level X for a working inventory C and a service level z:

```
(1-p)^2 X^2 - \left[2C(1-p) + z^2 p(1-p)\right] X + C^2 = 0
```

Take the smaller root. With 454 standard rooms of working inventory and 90% of peak nights walk-free, the house cap is 514. It differs from 526 because it answers a different question. Use the expected-value answer for policy and the service-level cap where a no-walk clause applies.

Groups are handled at cut-off. A wedding block of 150 whose comparable blocks used 85% would get 128 rooms by rule of thumb. The economic answer allocates the fewest rooms that keep the chance of the group needing more below R_low ÷ (R_low + R_high), here 117.50 ÷ (117.50 + 177) = 0.40. That's 129 rooms. Release the other 21 to transient.

When you do walk, walk well. Never walk a top-tier loyalty member, a returning guest, or anyone arriving after 22:00 without a car waiting. Walk the earliest-departing, lowest-value reservation.

Ask your team

- What's our overbooking level on a sell-out Tuesday, and how was it set?
- How many guests did we walk last year, and what did each walk cost us?
- Who decides whom to walk, and against what list?

Red flags

- "We never overbook." That's $178,000 a year on the specimen.
- One overbooking number for every night of the week.
- No record of walks or their cost.

Workbook: 09 Overbooking, 23 Overbooking II, 27 EHL Cases (section A).

Part 5 · Business mix

17 · Groups: displacement and the floor

A group wants 150 rooms for three nights at $131. Do we take it?

Take it if what the group leaves you, rooms and F&B, after the rooms it won't pick up, is more than the transient business it pushes out. Work it night by night, because a group is cheap on a soft night and expensive on a peak one.

For each night, the rooms free for the group are the sellable rooms less what's committed less the transient forecast. Anything the group picks up beyond that displaces a transient guest.

$$\text{Displaced}_n = \max(0,\ \text{Picked up}_n - \text{Free}_n) \qquad \Delta\Pi_n = \text{Picked up}_n \times c_g - \text{Displaced}_n \times c_t$$

c_g is the group's contribution per room night (net rate times the rooms margin, plus F&B times the F&B margin), and c_t is the same for the transient guest displaced. Picked up is the block less wash.

Night	Block	Picked up (12% wash)	Free for the group	Displaced	Net change ($)
Monday	150	132	120	12	+17,445
Tuesday	150	132	10	122	+2,553
Wednesday	150	132	30	102	+5,261
Total					+25,259

Accept. But Tuesday barely pays. Quote Sunday to Tuesday instead, or price Tuesday on its own.

The floor is the lowest group rate that leaves the hotel no worse off. Chris Anderson's form, as taught in the Cornell course, is:

$$G = VC - \frac{AC}{N} + \frac{1}{N}\sum_{n}\sum_{k} L_{nk}\,(R_k - VC)$$

N is group room nights, VC the variable cost per room, AC the group's ancillary contribution, L the rooms displaced from class k on night n and R that class's net rate. For a 400-room-night block Sunday to Wednesday that displaces 320 transient rooms (90% from the OTA class), the course formula gives $74.73. Add the F&B the displaced transient guests would have spent and the floor is $83.83. On rooms alone it would be $116.29. That's the floor. Never open with it.

The opening quote comes from the win curve. Fit a logistic curve to your own won and lost log, then quote where expected contribution peaks.

$$P(\text{win} \mid q) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 q)}} \qquad E[\Pi] = P(\text{win} \mid q) \times (q - G) \times N$$

On the specimen, expected contribution at the floor is $68. At $150 it's $17,211, the peak. Quote to the right of the peak when other groups are coming, to the left when the period is soft.

The LP in chapter 15 catches something the night-by-night method misses. A 60-room group on Tuesday and Wednesday takes 120 room nights and displaces 140, because three-night stays through Tuesday are lost whole. On the LP the break-even rate is $194.49 on rooms and $152.92 with the group's F&B. The course's Solver tool adds variable cost to lost revenue rather than to lost contribution, which is wrong.

Forecast group pickup multiplicatively. A 120-room block with 38 rooms picked up, at the point where comparable groups had 36% of their final count, will end at about 106. Release 14 rooms to transient at cut-off.

Ask your team

- Show me displacement for this group night by night, with F&B on both sides.
- What's the floor, what are we quoting, and where does the win curve come from?
- What wash does this account's own history show?

Red flags

- Groups priced on the average across the dates.
- A group accepted on a peak Tuesday because "it fills the shoulder".
- No lost-business log, so no win curve.

Workbook: 10 Group Displacement, 25 Group Break-even, 26 SDM Bid Price LP (rows 52 to 62).

18 · The long-term segment mix

How much of the hotel should we sell to groups, as a strategy?

Give a room to groups when the expected value of the next transient room falls below the group's value. The x-th transient room is only worth its rate times the chance that transient demand reaches x.

```
E[\text{x-th transient room}] = r_t \times P(D_t \ge x) \qquad y^{*} = \mu_t
+ \sigma_t \, \Phi^{-1}\!\left(1 - \frac{r_g}{r_t}\right)
```

On a shoulder-season weekday the specimen expects 340 transient rooms with a standard deviation of 50. Transient pays $155 net and the group's room element $112.50.

Basis	Group ÷ transient	Protect for transient	Rooms for groups	Mix
Room rate only	0.73	310	190	62% : 38%
Contribution, F&B included	0.97	245	255	49% : 51%

Count the banquet spend and 65 rooms move to groups. That's a decision about how big the sales team should be and how the function space is sold, so it belongs in the budget.

Ask your team

- What mix of transient and group are we aiming for on each day type, and on what basis?
- Is our group sales effort sized to that mix?

Red flags

- Group targets set as last year plus 5%.
- The mix argued on room rate alone.

Workbook: 24 Segment Mix.

19 · Corporate accounts and LRA

Should we give this account last-room availability?

Only if its rate clears the bid price on the nights it'll use. LRA means the hotel must sell the negotiated rate on its best nights. NLRA lets the hotel close it.

```
\text{Net value} = N(1 - s)\,(r\,m_{\text{rooms}} + f\,m_{\text{F\&B}}) + N s\,(r - \text{BP}_{\text{peak}})\,m_{\text{rooms}}
```

N is the account's room nights, s the share on peak nights, r its rate, f F&B per room night and BP the peak bid price ($154 here).

Account	Room nights	Rate ($)	On peak nights	Net value ($)	Decision
Pharma major	1,200	136.25	30%	96,003	Accept
Embassy (long stay)	900	122.50	10%	86,326	Accept

Account	Room nights	Rate ($)	On peak nights	Net value ($)	Decision
Global consulting firm	1,800	147.50	65%	75,422	NLRA or a higher rate
IT services	2,400	127.50	70%	46,014	NLRA or a higher rate

The IT services account brings the most room nights and is worth the least, because 70% of them land on nights the hotel could sell for more. The pharma account pays less than the consulting firm and is worth more, because most of its nights are soft ones.

Ask your team

- For each top-20 account, what share of its room nights fall on peak nights?
- Which accounts have LRA, and does each rate clear the peak bid price?

Red flags

- Accounts ranked on volume.
- LRA granted as a sweetener in the RFP.

Workbook: 16 Corporate Accounts.

20 · Distribution and digital demand

Which channels are actually making us money?

Rank them on what they leave you after every cost of the booking. Net ADR takes off the percentage cost and any fixed cost per booking, spread over the length of stay.

```
\text{Net ADR}_c = \text{ADR}_c\,(1 - k_c) - \frac{F_c}{\text{LOS}} \qquad \text{Acquisition cost \%} = \frac{\sum \text{channel costs}}{\text{Rooms revenue}}
```

Channel	Share of RNs	ADR ($)	Cost	Net ADR ($)
Brand.com and app	12%	184.38	3% + $3.13	177.11
Voice and direct	6%	172.50	2% + $5.00	166.27
Group and MICE sales	22%	158.13	2%	154.96
GDS (corporate)	20%	155.00	4% + $6.88	144.98
OTAs	18%	165.00	18%	135.30

Channel	Share of RNs	ADR ($)	Cost	Net ADR ($)
Wholesale	7%	117.50	net	117.50

The hotel spends $1.25m a year to acquire its bookings, 6.2% of rooms revenue. Moving 3% of room nights from the OTAs to brand.com at the same rate saves $93,000. Don't count on the billboard effect to do it for you. About 7% of OTA shoppers go on to book direct (Kalibri Labs and HSMAI, 2016), well below the 7.5 to 26% uplift often quoted from Anderson (2009). Treat the OTAs as paid acquisition and judge them on net rate.

Paid search needs the same discipline. On the specimen's month of search advertising, $55,860 bought 735 bookings.

```
\text{CPA} = \frac{\text{Spend}}{\text{Bookings}} \qquad \text{ROAS} =
\frac{\text{Revenue}}{\text{Spend}} \qquad \text{Value per click} =
\text{Contribution per booking} \times \text{Conversion}
```

CPA is $76 and ROAS 4.4, so acquisition costs 22.9% of revenue. That's dearer than the OTA's 18%. Bid on the marginal click. Each booking leaves $265.50, conversion is 2.5%, so a click is worth $6.64. Raise the bid while the extra cost of the extra clicks stays below that; on the specimen that stops between $2.75 and $3.25. The average CPC at that point is about $1.80, which is why averages flatter a bid.

The hardest question is how many of those bookings would have come anyway. A pause test on brand search found only 35% truly incremental. On the whole spend the return looks like 2.5 times; on the true increment it's 0.48.

Ask your team

- What's Net ADR by channel, and what's our total acquisition cost as a share of rooms revenue?
- What's our paid-search CPA against OTA commission, and when did we last run an incrementality test?

Red flags

- The channel report shows volume and gross ADR only.
- Brand-keyword bidding judged on ROAS with no pause test.

Workbook: 11 Distribution, 31 Digital Demand.

21 · Function space

Two enquiries want the ballroom on the same Saturday. Which one wins?

The one that leaves the most total contribution per square foot, rooms included. Measure the space the way you measure rooms, per unit of capacity per period.

```
\text{RevPAST} = \frac{\text{Function space revenue}}{\text{sq ft} \times \text{sessions}} \qquad \text{ConPAST} = \frac{\text{Function space contribution}}{\text{sq ft} \times \text{sessions}}
```

The specimen has 32,000 sq ft and three sessions a day, 1,095 a year. Its $7.15m of banquet revenue is a RevPAST of $0.20 and a ConPAST of $0.08.

Enquiry	Banquet ($)	Room nights	Total contribution ($)	Per sq ft ($)
Pharma conference, 300, 180 rooms × 3	40,000	540	76,585	2.39
Wedding, 450 guests, 120 rooms × 2	51,250	240	47,427	1.48
Wedding, 600 guests, no rooms	67,500	0	27,000	0.84

The biggest banquet cheque is worth the least. Check the conference's rooms against displacement (chapter 17) before it wins.

Ask your team

- What's our RevPAST and ConPAST by day part?
- When two enquiries compete, who decides, and on what number?

Red flags

- Catering takes space bookings without a rooms check.
- Space held on a first-come basis.

Workbook: 12 Function Space.

22 · Guest value

Which guests should we fight to keep?

The ones whose lifetime value is high against what it cost to win them. Price the guest as well as the night.

$$\text{LTV} = a \left(1 + \frac{\rho}{1 + d - \rho}\right) \qquad \frac{\text{LTV}}{\text{CAC}}$$

a is annual contribution, ρ the retention rate and d the discount rate (12% here). CAC is what it took to win the first stay: commission, media and sales time.

Guest	Contribution a stay ($)	Stays a year	Retention	LTV ($)	CAC ($)	LTV ÷ CAC
Corporate negotiated	275	4	70%	2,933	43.75	67
MICE planner	8,125	2	60%	35,000	750	47
Brand.com leisure	325	1	30%	444	35	13
OTA leisure	262.50	1	15%	303	61.25	5
Wedding host family	4,750	0.2	25%	1,223	562.50	2

The MICE planner is the most valuable person who walks into the hotel. The OTA guest is worth a direct-only benefit on the next stay, because moving him to brand.com raises his value more than anything else you can do for him.

Use recency, frequency and monetary value (RFM, each scored 1 to 5) to decide who's protected on a peak night. A returning CEO who scores 555 is never walked. Know the data law where your guests live and where the data sits: GDPR in the EU and UK, CCPA and CPRA in California, PDPA in Singapore and Thailand, the DPDP Act 2023 in India. Collect only what a stated purpose needs, record consent, and take counsel.

Ask your team

- Who are our 50 highest-value guests and planners, and what happens when they book?
- Do we know LTV and CAC by segment?

Red flags

- The CRM holds names and no value.
- Marketing consent isn't recorded.

Workbook: 32 CRM & Guest Value.

23 · Resorts

Does any of this change for a resort?

The maths is the same, but a resort sells stays rather than nights, and most of its money is spent after check-in. The specimen resort has 120 keys and four seasons.

Length of stay is the first lever. A minimum stay pays when the room nights it fills on the shoulder days are worth more than the short-stay arrivals it turns away.

Season	Nights	BAR ($)	Forecast occupancy	MinLOS	Arrivals lost	Room nights gained
Peak (Christmas to New Year)	12	525	97%	4	30	210
High (Oct to Mar)	150	325	78%	2	120	600
Shoulder	90	225	60%	1	0	0
Monsoon low	113	150	38%	1	0	0

Meal plans are the second. Price the plan on contribution. At a $325 room-only rate, two guests and a 15% package discount on menu value:

Plan	Package price ($)	Room rate inside it ($)	Contribution a night ($)
EP, room only	325	325	260
CP, breakfast	363	318	281
MAP, breakfast and one meal	444	304	324
AP, all meals	508	293	357

All-inclusive leaves $97 a night more than room only, even though its room rate is $32 lower. Report the room rate to STR net of the meal component, or your ARI will flatter you.

The third is the room-type ladder. The private-pool villas sell out 95 nights a year at a 180% premium. The category that sells out first is underpriced; widen its differential until its sell-outs match the rooms below it, and don't upgrade into villas for free on peak nights. With F&B at $77.50 and spa and activities at $30 per guest a day, total revenue per occupied room is $540, and 40% of it is outside rooms. Resorts live on TRevPAR and GOPPAR.

Ask your team

- What MinLOS do we run each season, and what did it cost in turned-away arrivals?
- Which plan earns us the most per night?

Red flags

- The villas sell out first every peak.
- A meal-plan room rate reported gross to STR.

Workbook: 33 Resorts.

Part 6 · Governance

24 · Budget and variance

We beat the rooms budget by $43,000. Is that good news?

Not necessarily. Split the variance into price, volume and the joint effect, and the story is often different from the headline. The three parts add up exactly to the total.

$$\text{Price} = (\text{ADR}_a - \text{ADR}_b) \times \text{RN}_b \qquad \text{Volume} = (\text{RN}_a - \text{RN}_b) \times \text{ADR}_b \qquad \text{Joint} = \Delta\text{ADR} \times \Delta\text{RN}$$

The specimen budgeted 140,640 room nights at $143.75 and delivered 135,050 at $150.02.

Part	Amount ($)
Price variance	+881,603
Volume variance	−803,563
Joint variance	−35,041
Total	+43,000

The hotel lost 5,590 room nights and made it back on rate. The negative joint term says the rooms it lost would have carried the higher rate. In plain terms, rate was bought with volume. Read that beside the RGI of 84 in chapter 3 and you have a hotel that priced up into a market it was already behind.

The budget itself should be built month by month: this year's base, times a seasonal index for each month, plus the strategy.

$$\text{Occ}_m = \text{Occ}_{\text{base}} \times I^{\text{occ}}_m + u_{\text{occ}} \qquad \text{ADR}_m = \text{ADR}_{\text{base}} \times I^{\text{ADR}}_m \times (1 + u_{\text{ADR}})$$

With a two-point occupancy uplift and 6% on rate, the specimen's next-year rooms budget is $22.10m, 9.1% up. The uplifts are the programme in this book. Every point should trace to a chapter. Use your own three-year monthly history and fiscal calendar for the indices.

Ask your team

- Split last month's rooms variance into price, volume and joint.
- Which chapters of the programme carry the uplifts in next year's budget?

Red flags

- "Ahead of budget" with RGI falling.
- A budget built as last year plus a percentage.

Workbook: 14 Variance, 15 Budget.

25 · Running it: rhythm, systems and the prize

What should I expect my revenue team to do, what should I buy, and what's it all worth?

Revenue management is a calendar. Here's the one I'd hold a team to.

Cadence	Routine	Owner
Daily	Pickup and pace for the next 90 nights; BAR tier, fences and MinLOS; tonight's overbooking limit and walk plan	Revenue manager
Weekly	Revenue meeting: forecast by segment, group decisions on displacement, channel mix, rate shop	DOR with GM, DOSM and F&B
Monthly	STR indices, Net RevPAR, forecast accuracy, three-part variance, flow-through	DOR with Finance

Cadence	Routine	Owner
Quarterly	Segment contribution ranking, price tests, channel contracts	DOR
Yearly	Corporate RFP on displacement, budget by segment, BAR ladder reset, comp-set review	DOR, DOSM and GM

Ask for five numbers every month: RGI, Net RevPAR, WAPE at 14 days with its bias, flow-through and review score. On the specimen they read 84.1, $104.14, 5.1% with a bias of −8.7 rooms, 26.9% and 7.9.

Buy systems in this order: clean PMS data first, then a channel manager, a rate shopper, an RMS, and only then BI and CRM. An RMS fed bad segment codes, no denials and a comp set picked by habit will produce confident wrong answers, AI or not. On the specimen, a 3% uplift in rooms revenue from an RMS is worth $486,000 of contribution against a $75,000 licence and integration cost. That's 6.5 times the cost, paid back in under two months, and it breaks even at a 0.5% uplift. Licences are on quote, so get your own number.

What is the whole programme worth? Two lines that can't overlap. Closing half the gap to fair share is worth $1.54m of rooms contribution a year. Moving 3% of room nights from the OTAs to brand.com saves $93,000. Together that's $1.63m a year, $3,258 per key and 20.6% of this year's EBITDA.

The tools are how the first line gets closed. They sit inside it, so don't add them on top:

Inside the fair-share line	A year ($)	Chapter
Overbooking from 512 to 526 on sell-out nights	178,352	16
Attribute selling, at the rooms margin	434,861	10
Reputation from 7.9 to 8.7, half realised	363,059	13
BAR discipline, fences, mix and LOS controls	559,378	8, 10, 14, 15

Group and corporate repricing (chapters 17 and 19) are left out until the account data are loaded.

Ask your team

- Show me the calendar. Which routines happened last month and which didn't?

- What does our RMS decide, what do people override, and how often are the overrides right?
- What's our prize, on our numbers?

Red flags

- No weekly revenue meeting, or one the GM doesn't attend.
- An RMS bought before the PMS data were clean.
- Overrides that nobody logs.

Workbook: 19 Rhythm & KPIs, 28 Tech Stack, 18 Prize, Dashboard.

Appendices

A · The formula sheet

Measure	Formula	Chapter	Sheet
RevPAR	Occupancy × ADR, or rooms revenue ÷ rooms available	2	01
Net RevPAR	(Rooms revenue − acquisition costs) ÷ rooms available	2	01
TRevPAR, GOPPAR	Total revenue, or GOP, ÷ rooms available	2	01, 13
Flow-through	Δ EBITDA ÷ Δ revenue	2	13
MPI, ARI, RGI	Hotel ÷ comp set × 100, for occupancy, ADR and RevPAR	3	01
Contribution per RN	Net ADR × rooms margin + F&B per RN × F&B margin	4	02
Pickup forecast	OTB + average pickup, or OTB × average (final ÷ OTB)	5	03
Unconstrained demand	Sold + share converting × denials	6	04
WAPE, bias, tracking signal	Σ|A − F| ÷ ΣA; mean (A − F); Σ(A − F) ÷ MAD	7	03, 21
Optimal price	P* = (a ÷ b + c) ÷ 2; capped at (a − C) ÷ b	9	05

Measure	Formula	Chapter	Sheet
BAR tier	$BAR_1 \times (1 + step)^{(n - 1)}$	10	06
Break-even volume	$-(\Delta P - \Delta VC) \div (CM + \Delta P - \Delta VC)$	11	22
Elasticity	Slope of ln FSI on ln ARI	11	22
Price index	Own BAR ÷ comp median × 100	12	30
Littlewood	$y = \mu_1 + \sigma_1 \Phi^{-1}(1 - r_2 \div r_1)$	14	07
Bid price test	Accept if Σ rate ⩾ Σ bid prices over the stay	15	08, 26
Overbooking, critical fractile	Overbook the k-th room while P(no-shows ⩾ k) ⩾ W ÷ (W + E)	16	09, 23
Group floor	VC − AC ÷ N + ΣΣ L(R − VC) ÷ N	17	25
Transient protection	$\mu + \sigma \Phi^{-1}(1 - \text{group value} \div \text{transient value})$	18	24
Net ADR by channel	ADR × (1 − cost %) − fixed cost ÷ LOS	20	11
Value per click	Contribution per booking × conversion	20	31
RevPAST	Space revenue ÷ (sq ft × sessions)	21	12
LTV	Annual contribution × (1 + retention ÷ (1 + discount rate − retention))	22	32
Three-part variance	ΔADR × RN_b; ΔRN × ADR_b; ΔADR × ΔRN	24	14

B · Glossary

Term	Meaning
Acquisition cost	Everything paid to win a booking: commission, transaction fees, media
BAR	Best available rate; the lowest unrestricted rate on sale tonight
Bid price	The value of the last room on a night; the lowest net rate worth accepting for it

Term	Meaning
Booking curve	Rooms on the books by days before arrival, for one night type
Comp set	The four to seven hotels a guest would genuinely shop against yours
CTA, CTD	Closed to arrival, closed to departure on a given night
Denial	An enquiry turned away because the room or rate was closed
Displacement	The business turned away by saying yes to something else
EMSR-b	Belobaba's method for protection levels across several nested rate classes
Fence	A condition that stops a guest who would pay more from taking a discount
FSI	Fair-share index: rooms sold ÷ the hotel's fair share of the market's rooms sold
LRA, NLRA	Last-room availability, or not: whether a negotiated rate must be sold on the hotel's best nights
MinLOS	A minimum length of stay for arrivals on, or stays through, a night
Nesting	A higher rate class may always take a room left for a lower one
OTB	On the books: reservations held today for a future night
Pace	OTB against the curve or against the same time last year
Pickup	Rooms booked between a point in time and arrival
RMS	Revenue management system: forecasts, prices and controls
Shadow price	In an LP, what one more unit of a constraint is worth

Term	Meaning
SMERF	Social, military, educational, religious and fraternal groups
Spill	Demand turned away on a constrained night
STLY	Same time last year, same day of the week
Unconstrained demand	What demand would have been with no rooms or rates closed
USALI	The Uniform System of Accounts for the Lodging Industry, 12th revised edition
Walk	Relocating a guest with a reservation to another hotel
Wash	Group rooms blocked but not picked up

C · Sources

- Anderson, C.K. (2009), *The billboard effect: online travel agent impact on non-OTA reservation volume*, Cornell CHR Report 9(16).
- Anderson, C.K. (2012), *The impact of social media on lodging performance*, Cornell CHR Report 12(15).
- Belobaba, P. (1989), Application of a probabilistic decision model to airline seat inventory control, *Operations Research* 37(2).
- Belobaba, P. (1992), Optimal vs. heuristic methods for nested seat allocation, AGIFORS Reservations and Yield Management Study Group.
- Cornell eCornell, Revenue Management 360 (SHA531 to 535, 541 to 543, 774) and SHA775, SHA871, SHA872.
- Cross, R.G., Higbie, J.A. and Cross, Z.N. (2009), Revenue management's renaissance, *Cornell Hospitality Quarterly 50(1)*.
- EHL, Revenue Management for Market Leaders (Heo, 2020).
- Gupta, S. and Lehmann, D. (2005), *Managing Customers as Investments*.
- HFTP and AHLA, *Uniform System of Accounts for the Lodging Industry*, 12th revised edition.
- Kalibri Labs and HSMAI Foundation (2016), *Demystifying the Digital Marketplace*, Part 1.

- Kimes, S.E. (1989), Yield management: a tool for capacity-constrained service firms, *Journal of Operations Management* 8(4).
- Kimes, S.E. and Chase, R.B. (1998), The strategic levers of yield management, *Journal of Service Research* 1(2).
- Littlewood, K. (1972), Forecasting and control of passenger bookings, AGIFORS.
- Talluri, K. and van Ryzin, G. (2004), *The Theory and Practice of Revenue Management*, Springer.
- Toh, R.S. and DeKay, F. (2002), Hotel room-inventory management: an overbooking model, *Cornell HRA Quarterly* 43(4).
- Weatherford, L. and Pölt, S. (2002), Better unconstraining of airline demand data in revenue management systems, *Journal of Revenue and Pricing Management* 1(3).

Non multa, sed multum.