Excelsior Restaurant Revenue Management

Standard operating procedures · Edition II · Oct 3, 2026 · @Vikram

How to use this manual

This manual turns Excelsior's restaurant revenue doctrine into twenty-six procedures that an outlet team runs every day, week, month and quarter. It is the how-to. The Practice Bible is the why, and *The Mathematics* workbook holds every calculation. Each SOP names its workbook sheet; change the gold cells and every number below recomputes.

Governing thought. A restaurant sells seat-hours, not meals; a banquet department sells square-foot sessions. Profit comes from pricing every plate on purpose and selling every seat-hour and every session at the right price, for the right duration, to the right guest.

The chain. Run the procedures in order. Each feeds the next, and a later fix will not hold while an earlier link is broken.

#	Procedure	Answers	Sheet
1	Costing	What does each plate really cost?	01, 02
2	Pricing	What should it sell for?	03
3	Menu engineering and design	Which dishes earn their place, and where do they sit?	04
4	RevPASH	What does each seat earn each hour?	05
5	Hot and cold periods	Which tool in which hour?	06
6	Table mix and combinable tables	Do the tables fit the parties?	07, 17 D
7	Meal duration	Where do the minutes go, and what are they worth?	08
8	The booking curve	What will tonight and next week bring?	09
9	Regression forecasting	What drives demand, and how many covers on Thursday?	17
10	Protection and displacement	How many seats to hold for better guests?	10

#	Procedure	Answers	Sheet
11	No-shows and overbooking	How many bookings to take?	11
12	The peak cover charge	How do we price the risk of the empty seat?	15
13	The queue test	How long will guests wait?	12
14	Labour to standard	How many servers each hour?	13
15	Banquets: RevPASF, RevPAST, ConPAST	What does each square foot earn?	16
16	The outlet P&L	What did it all earn?	14

Edition II. Ten further procedures extend the chain beyond the seat, on workbook sheets 18 to 26: 17 conditions and levers (18), 18 reservations and pacing (19), 19 channels (20), 20 beverage (21), 21 guest value (22), 22 the kitchen (23), 23 the hotel link (24), 24 variability and slack (25), 25 competitor benchmark (26) and 26 the technology layer. The programme in SOP 16 is corrected: overbooking is now valued after the card guarantee, not before it.

Roles. The F&B Director owns the programme. The Outlet Manager runs the daily and weekly routines. The Executive Chef owns costing and recipes. The Cost Controller owns yields, inventory and the actual-over-theoretical gap. The Revenue Manager owns the forecasts, protection, overbooking and the cover charge. The Director of Catering owns function space. The Excelsior asset manager reviews monthly on the owner's behalf.

The specimen. Worked examples use a fictitious 100-seat all-day dining room in an Indian five-star hotel: 31 tables, 16 trading hours, about 643 covers a day, ₹2.18 crore of revenue a month. Its menu is deliberately international: Butter Chicken, Lamb Biryani, Margherita Pizza, Spaghetti Carbonara, Club Sandwich, Grilled Norwegian Salmon, New Zealand Lamb Chops and Vegetable Lasagne. The hotel also runs 19,000 square feet of function space. The numbers are realistic, not real. Replace every gold input with the outlet's own data before any decision.

Conventions. Rupees, net of GST. Metric units, except function space, which the Indian trade quotes in square feet. Percentages of net revenue unless stated.

Data an outlet must hold. Twelve weeks of hourly POS (covers, revenue, party size, seated and closed times). Two years of daily covers for regression. Standard recipe cards and current invoices. The floor plan with every table size. The roster and payroll. The

reservations export with booking dates, party sizes and no-shows. The hotel's occupancy and events diary. Banquet event orders by venue and session.

What the research allows us to claim. Four restraints are built into every number in this manual. Shorter meals return less than a quarter of their arithmetic (Thompson 2009), so the workbook keeps 40% of the capped gain. Combinable tables win only in smaller rooms where small parties dominate (Thompson 2002). There is no "sweet spot" on a menu page (Yang 2012). Banquet F&B earns far less contribution than room hire (Kimes and McGuire 2001), so space is ranked on contribution, not revenue.

SOP 1 · Costing the plate

Every plate is costed from one ingredient master at edible-portion cost, and the kitchen's actual food cost is reconciled to that theory every month. Sheets 01 Ingredients and 02 Recipes.

Owner Executive Chef, with the Cost Controller. **Frequency** Yield tests quarterly and on any price move above 10%. Recipe costs monthly. Actual-over-theoretical gap monthly.

1.1 Run a yield test

1. Take a typical delivery. Weigh it as purchased (AP), to the gram.
2. Process it exactly as the recipe card says: bone, skin, pin-bone, trim, peel.
3. Weigh the usable part. Log the trim separately, with anything sold or reused (bones to stock).
4. Yield % = usable weight ÷ AP weight. Repeat on three deliveries and average.
5. Enter AP price and yield in the master. The sheet computes EP cost.

```
\text{EP cost per kg} = \frac{\text{AP price per kg}}{\text{yield}}
```

Specimen: Norwegian salmon at ₹2,200 a kg with a 78% yield costs ₹2,821 a kg on the plate, 28% above the invoice. Lettuce at a 65% yield costs 54% above it.

1.2 Keep one ingredient master

- One row per purchased item: code (I01, I02…), unit, AP price, yield, EP cost.
- House preparations (makhani gravy, carbonara base, pizza dough, stocks) get their own row, costed by the batch and divided by batch yield.

- Update AP prices from the latest invoice every month. Never type a cost into a recipe; recipes look up EP cost by code.

1.3 Cost the standard recipe

1. Write the recipe for one portion: ingredient code and quantity in the master's unit.
2. Line cost = quantity × EP cost. Raw cost = the sum of the lines.
3. Add the Q-factor for uncosted items (oil, garnish, condiments, salt). House setting 8%; audit it quarterly.
4. Plate cost = raw cost × (1 + Q). Food cost % = plate cost ÷ net price. Contribution margin = price − plate cost.

```
\text{Plate cost} = \Big(\sum_i q_i \times EP_i\Big)(1 + Q)
```

Specimen: Butter Chicken costs ₹200 a plate, sells at ₹945, runs 21% food cost and leaves ₹745. New Zealand Lamb Chops cost ₹924, run 42% and leave ₹1,271. The higher food cost earns more rupees.

1.4 Reconcile actual to theoretical

1. Month-end actual food cost = (opening stock + purchases − closing stock − staff meals and transfers out) ÷ food revenue.
2. Theoretical food cost = Σ(units sold × plate cost) ÷ food revenue. Sheet 04 gives it: 27.5% for the specimen mains.
3. Gap = actual − theoretical, in points. Standard: 3 points or less. The specimen P&L carries 4 points, which is the problem to fix.
4. Above 3 points: test portions at the pass for a week, check waste logs, recount high-value stock (lamb, salmon, cheese, saffron, butter), and check that every comp and staff meal was recorded.

Common errors. Costing at invoice instead of EP. Recipe cards that no longer match what the cooks make. A Q-factor set once and never audited. Judging dishes by food cost % instead of rupees of margin.

SOP 2 · Pricing the plate

Price every main by blending the factor and contribution-margin methods, test it against the market, and end it the house way. Sheet 03 Pricing.

Owner F&B Director, with the Executive Chef. **Frequency** At every menu revision, and for any new dish before it is printed.

2.1 Procedure

1. Take the plate cost from SOP 1.
2. Factor price = plate cost ÷ target food cost %. House target for mains 30%; for accompaniments (breads, rice) 10%.
3. Contribution-margin price = plate cost + target margin. House target for mains ₹700.
4. Blended price = the mean of the two. Accompaniments use the factor price alone.
5. Competitive band: what three comparable rooms charge for the same dish. A blended price above the band needs a reason the guest can see: an ingredient, a portion, a provenance.
6. House ending: round up to the next ₹50, take off ₹5. Prices end in 45 or 95.
7. Do not move the printed price yet. Hand the house price to SOP 3, which decides whether the dish can bear it.

```
P_{factor} = \frac{C}{FC^*} \qquad P_{CM} = C + CM^* \qquad P = \Big\lceil
\tfrac{P_{factor}+P_{CM}}{2} \Big\rceil_{50} - 5
```

2.2 Why two methods

The factor method alone overprices costly proteins and underprices cheap dishes. Lamb Chops, with a ₹924 plate, would sell at ₹3,081; the Club Sandwich, with a ₹125 plate, at ₹416. The contribution method asks every main to clear the same rupees towards the room, the brigade and the profit. The blend tempers both.

Dish	Plate ₹	Factor ₹	Cost + ₹700	House ₹	Menu ₹
New Zealand Lamb Chops	924	3,081	1,624	2,395	2,195
Lamb Biryani	463	1,543	1,163	1,395	1,345
Butter Chicken	200	666	900	795	945
Club Sandwich with fries	125	416	825	645	695
Garlic Naan, basket of three	29	295	—	295	245

Reading: Lamb Chops and Biryani sit below the house logic. Butter Chicken sits ₹150 above it and still sells 980 plates a month, so the market has already ruled. The house price is a guide, not a verdict.

Common errors. One markup for every dish. Pricing from the competitor's menu alone. Endings in 99, which read as discount. Moving prices without the SOP 3 test.

SOP 3 · Menu engineering and menu design

Classify every dish by popularity and rupee margin within its category, act on each class, test every price move against guest response, and lay out the page by how guests actually read. Sheet 04 Menu Eng.

Owner F&B Director, Executive Chef, Outlet Manager. **Frequency** Quarterly, on twelve weeks of POS data; before every reprint.

3.1 Classify (Kasavana and Smith)

1. Pull units sold per dish for twelve weeks. Work one category at a time: starters, mains, accompaniments, desserts. Never mix categories.
2. Mix % = units of the dish ÷ units in the category.
3. Popularity line = (1 ÷ items in the category) × 70%. With eight mains, 8.75%.
4. Item margin = price − plate cost. Menu margin = units × item margin.
5. Profitability line = total menu margin ÷ total units, the weighted average margin. Specimen: ₹734.
6. Popular and above the line = Star; popular and below = Plowhorse; unpopular and above = Puzzle; unpopular and below = Dog.

3.2 Act on each class

Class	Specimen dishes	Action
Star	Butter Chicken, Lamb Biryani, Spaghetti Carbonara	Protect. Hold price, placement, recipe and portion. Never discount.
Plowhorse	Margherita Pizza, Club Sandwich	Reprice in small steps, or re-engineer the portion or garnish to lower plate cost.
Puzzle	Grilled Norwegian Salmon, New Zealand Lamb Chops	Promote: lead the section, a better description, servers trained to recommend, a half portion at lunch.

Class	Specimen dishes	Action
Dog	Vegetable Lasagne	Remove, or rebuild as a Puzzle. Keep one only if it serves a guest need, such as the one vegetarian main of substance.

3.3 Test a price move

1. Propose the new price, guided by SOP 2.
2. Apply elasticity: new units = units × (1 + e × % price change). Use e = −0.8 for premium dining until the outlet's own history gives a better number.
3. Compare monthly margin before and after. Move only where margin rises.
4. After the reprint, watch the mix for eight weeks. The real response becomes the outlet's elasticity.

Specimen: Margherita ₹795 → ₹845 and Club Sandwich ₹695 → ₹745 lose about 80 plates a month between them and gain ₹22,767 of margin.

3.4 Design the page

1. Design the mains page first. Guests choose the main and build the meal around it.
2. Guests read a menu like a book, top-left to bottom-right, in five or six fixations. The old "upper-right sweet spot" did not survive eye-tracking (Yang 2012). Lead each section with the dish you most want sold.
3. Prices as plain numbers after the description: no ₹ sign, no dot leaders, no price column. A Cornell field test found spend about 8% higher, though the overall result was not statistically significant; treat it as directional (Yang, Kimes and Sessarego 2009).
4. One quiet box per section, at most. Five to seven items a section. Names that place the dish; one-line descriptions.
5. Plowhorses go last in a section, where they still sell unaided. Dogs go, or sit mid-section until removed.

Common errors. Engineering the whole menu as one list. Removing a Star to "freshen" the menu. Discounting Stars. Designing for a sweet spot that does not exist. Repricing without measuring what happened next.

SOP 4 · Measuring RevPASH

Measure revenue per available seat-hour by hour, and split it into occupancy and spend so you know which lever to pull. Sheet 05 RevPASH.

Owner Outlet Manager, with the Revenue Manager. **Frequency** Weekly by hour; monthly by day part.

4.1 Procedure

1. Export the POS by hour: covers seated, net revenue, and guests and tables in house at the half-hour.
2. Seat-hours available = seats × trading hours. Specimen: 100 × 16 = 1,600 a day.
3. RevPASH for the hour = revenue in the hour ÷ seats.
4. Seat occupancy = guests in house ÷ seats. Revenue per occupied seat-hour = revenue ÷ guests in house.
5. Check the identity on every row: RevPASH = seat occupancy × revenue per occupied seat-hour.
6. Seat fit = guests ÷ seats at the occupied tables. Table occupancy = tables in use ÷ tables.
7. Read each hour: low occupancy is a demand problem; low spend per occupied seat-hour is a menu, upsell or duration problem.

```
\text{RevPASH} = \frac{\text{revenue}}{\text{seats} \times \text{hours}} = \text{seat occupancy} \times \frac{\text{revenue}}{\text{occupied seat-hours}}
```

Measure	Specimen
RevPASH, whole day	₹454
RevPASH, 20:00	₹1,147 (88% occupied × ₹1,303)
RevPASH, 16:00	₹52 (10% occupied × ₹520)
Average seat occupancy	44%
Share of revenue from the three hot hours	43%

Common errors. Reporting covers and average check separately and missing the hour that hides the problem. Using a daily total, which averages a full dinner with an empty afternoon. Counting seats that are not set.

SOP 5 · Hot and cold periods

Label every hour hot, warm or cold, and apply only the tools that fit it. Sheet 06 Hot-Cold.

Owner Outlet Manager. **Frequency** Set quarterly; checked weekly against the forecast.

5.1 Procedure

1. Take seat occupancy and average check by hour from SOP 4.
2. Hot: occupancy at or above 70%. Cold: below 40%. Warm: between. House thresholds; adjust per outlet.
3. Split hot and cold by check: above or below the day's average.
4. Assign the tool set. Write it into the shift brief for that hour. Never run a tool across the whole day.

Period	Tools
Hot, high check · protect	Duration control, protection levels, cover charge on sell-out nights, no discounts, set menus for groups
Hot, low check · reprice	Raise price or upsell, premium specials, cut slow courses, timed slots
Warm · nudge	Suggestive selling, bundled courses, targeted release of tables to offers
Cold, low check · stimulate	Off-peak menus, residents' offers, events, early-bird pricing
Cold, high check · reposition	Afternoon tea, private dining, chef's table, takeaway

Specimen: three hot hours (13:00, 20:00, 21:00), six warm, seven cold. The afternoon from 15:00 to 18:00 is an empty room with a payroll, and the first place to reposition.

Common errors. A dining offer that runs at 20:00. A duration push at 16:00. A cover charge on a cold night. Treating breakfast, lunch and dinner as one business.

SOP 6 · Table mix and combinable tables

Match the tables to the parties who actually come, using Little's Law by party size, and build the room on twos that join where the evidence says it pays. Sheet 07 Table Mix, and sheet 17 section D for lunch.

Owner F&B Director, Outlet Manager. **Frequency** Twice a year, and before any refurbishment.

6.1 Size the mix for the peak

1. From twelve weeks of POS take the peak-hour party arrivals (λ), the share of parties by size, and the average seated time by size.
2. Assign each size to its table: 1–2 to a two-top, 3–4 to a four-top, 5–6 to a six-top, 7–8 to an eight.
3. Tables in use for each size = λ × share × minutes ÷ 60 (Little's Law).
4. Sum by table type, add a 10–15% variability buffer and round up: the required count.
5. Utilisation now = tables in use ÷ tables of that type. Above 90% is where parties queue, even in a half-empty room.
6. Seat fit = guests ÷ seats at the assigned tables. 75–85% is typical; above 90% is good.
7. If required seats exceed the room, the answer is duration (SOP 7) or price (SOP 2), not furniture.

```
L = \lambda W \qquad \text{tables in use}_k = \lambda \cdot s_k \cdot \frac{d_k}{60}
```

Table	Today	In use at 20:00	Required (+10%)	Utilisation now
Two-top	14	12.1	14	86%
Four-top	15	9.4	11	63%
Six-top	2	4.4	5	218%

Reading: large parties queue while four-tops run more than a third empty. Required seats total 102 against 100, so the Saturday peak needs the duration programme as well.

6.2 Why a room built on twos

Two is the common unit of almost every party. A floor of two-tops that push together seats a couple at two chairs, a four at two tables and a six at three, so seat fit rises without a single new chair.

The evidence, and its condition:

- Chevys Arrowhead found 52–56% of parties were ones and twos, yet the room was mostly four-tops. Moving to 40 two-tops, 24 fours, 5 sixes and 3 eights lifted revenue about 5 points against control stores, with payback inside a year (Kimes and Thompson 2004; Kimes 2004).
- Combinable tables win in rooms of about 100 seats where small parties dominate. At 200 seats, or where large parties are common, dedicated tables of the right sizes win, because a large party waits for adjacent tables to clear and the held tables earn nothing meanwhile (Thompson 2002).

The rule: dedicated two-tops as the base, a block of joinable twos for fours and sixes, one round eight. Re-set twice a day: twos for lunch, joins for dinner.

6.3 The Thursday lunch plan

Feed the regression forecast (SOP 9) into the plan.

1. Peak-hour covers = forecast × peak share. Specimen: 163 × 45% = 73 covers.
2. Parties in house = peak covers ÷ average party × duration ÷ 60, split by party size.
3. Seats blocked, fixed layout = parties × the seats of the table they must take. A couple at a four-top blocks four.
4. Seats blocked, flexible layout = parties × the seats they need, in twos.
5. Seats freed = fixed − flexible. Extra covers = seats freed × flexible seat fit. Value = extra covers × lunch check.

Lunch peak	Today's floor	Twos that join
Seats blocked	81	72
Seat fit	76%	85%
Seats freed	—	9
Extra covers a peak hour	—	9

Lunch peak	Today's floor	Twos that join
Value a Thursday, if the demand exists	—	₹13,567

At lunch 70% of parties are one or two guests. That is the room Thompson's condition describes.

Common errors. Designing the floor around the largest party. Seating couples at four-tops at the peak. Making every table combinable in a large room. Ignoring that larger parties stay longer.

SOP 7 · Meal duration

Measure the meal by stage, cut only the controllable waits, run the programme only on nights that turn guests away, and value it at what the research says actually arrives. Sheet 08 Duration.

Owner Outlet Manager, with the Executive Chef. **Frequency** A stopwatch study twice a year; the programme on every spill night.

7.1 Run the study

1. Choose ten dinner services across the week. Time at least 100 parties.
2. Record each timestamp: arrived, seated, order taken, first course served, plates cleared, check requested, paid, departed.
3. Mean and standard deviation of each stage. Whole-meal spread = √(sum of stage variances).
4. Mark each stage controllable or not. Eating and the farewell are the guest's; the waits between are the house's.

Dinner stage	Mean min	Controllable	Target cut	How
Greet to seated	4	Yes	1	Table ready before arrival; host map
Seated to order	9	Yes	3	Water and menus within 60 seconds; specials told once

Dinner stage	Mean min	Controllable	Target cut	How
Order to first course	14	Yes	3	Fire on order; mise en place by forecast
Eating, all courses	46	No	0	Never rush it
Cleared to check requested	6	Yes	2	Clear within two minutes; offer dessert once
Check requested to paid	5	Yes	2	Handheld payment at the table
Paid to departed	3	No	0	Never rush it
Total	**87**		**11**	New mean 76, σ 14.8

Guests accept meals about 20% shorter without losing satisfaction, provided the minutes come out of ordering and paying, not eating (Kimes, Wirtz and Noone 2002; Noone et al. 2007). Eleven minutes is 13%.

7.2 Value the programme honestly

1. Turns in the peak window = window minutes ÷ duration. Specimen: 180 ÷ 87 = 2.07 now; 180 ÷ 76 = 2.37 after.
2. Capacity freed = change in turns × seats in play. Specimen: 27.3 covers a service.
3. Covers spilled = parties turned away × average party. Specimen: 6 × 2.9 = 17.4.
4. Arithmetic gain = the lower of the two: 17.4 covers.
5. Realised gain = arithmetic × 40%. Parties do not arrive the moment a table frees; across 1,200 simulated scenarios the real gain from shorter meals was under a quarter of the naive figure (Thompson 2009). Specimen: 7.0 covers.
6. Revenue = realised covers × dinner check. Profit = revenue × contribution ratio (SOP 16).

Specimen: 7.0 covers × ₹1,744 = ₹12,159 a sell-out night; ₹31.6 lakh a year across 260 nights; ₹1.50 lakh of profit a month. The naive figure would have claimed ₹3.8 lakh. ₹1.50 lakh is the number to defend in front of an owner.

7.3 Run it on the night

- The forecast (SOP 8) flags spill nights a week ahead.

- On those nights: tables pre-set, menus and water within 60 seconds, fire on order, the check printed when dessert is cleared, payment at the table.
- On every other night, let the meal breathe. A faster table on an empty night earns nothing.

Common errors. Rushing the guest instead of the process. Running the programme every night. Valuing it at turn-time arithmetic. Not measuring before and after.

SOP 8 · The booking curve

Forecast every reserved service from the booking curve and pickup, score it every month, and spread the week with a day-of-week index. Sheet 09 Forecast.

Owner Revenue Manager, with the Reservations Manager. **Frequency** Daily pace check; weekly refresh; monthly accuracy review.

8.1 Build the curve

1. Export reservations with the date each booking was made.
2. For each past service of one type (Saturday dinner, weekday lunch), count covers on the books at 30, 21, 14, 7, 3 and 1 days out, and the final count.
3. Use at least eight comparable services; twenty is better. Drop holidays and one-offs.
4. Share of final = column average ÷ final average.
5. Additive pickup = average of (final − on books). Multiplicative ratio = average of (final ÷ on books).

Specimen Saturday: 22% of final at 30 days, 51% at 14, 69% at 7, 94% at 1.

8.2 Forecast a future service

1. Read today's covers on the books and the days left.
2. Pace = on books ÷ the curve's average at that point − 1.
3. Additive forecast = on books + average pickup. Multiplicative = on books × average ratio. Blended = the mean.
4. Spill = forecast − reservable capacity, if positive.
5. Act. Ahead of pace or spilling: protect prime slots (SOP 10), switch on the cover charge (SOP 12), run the duration programme (SOP 7). Behind pace: release slots to offers and walk-ins now.

$$\hat{F}_{add} = B_t + \overline{(F - B_t)} \qquad \hat{F}_{mult} = B_t \times \overline{F / B_t} \qquad \text{pace} = \frac{B_t}{\bar{B}_t} - 1$$

Specimen: seven days out, 128 on the books against a curve of 144, pace −11%. Additive 193, multiplicative 186, blended 189 against 200 seats. Release slots now.

8.3 Score the forecast

- MAPE = average of |actual − forecast| ÷ actual. Standard: under 10% for a reservations-led room, under 15% for walk-in-led.
- Bias = average of (actual − forecast). Bias that persists two months means the method is wrong.
- Poor forecasts can inflate staffing by up to half again (Thompson's labour-scheduling work), which is why accuracy is a P&L line, not a statistic.

Specimen: MAPE 2.3% on eight recent Saturdays the curve did not see; bias −1.0 covers.

8.4 The week ahead

1. Smoothed weekly covers = α × this week + (1 − α) × last smoothed, with α of 0.2 to 0.4.
2. Day-of-week index = average covers that day ÷ average daily covers.
3. Next week by day = smoothed total ÷ 7 × index. This drives the roster (SOP 14) and prep.

Specimen: next week 4,876 covers; Monday index 0.84, Saturday 1.19.

Common errors. Forecasting from memory. Never scoring the forecast. One curve for every day of the week.

SOP 9 · Regression forecasting

Where demand is walk-in-led, as at a hotel lunch, there is no booking curve to read. Forecast from the drivers instead, with multiple regression, and turn the forecast into a table plan. Sheet 17 Regression.

Owner Revenue Manager. **Frequency** Forecast weekly; model refitted quarterly.

9.1 Build the model

1. Pull 104 weeks of the day part (here, Thursday lunch) from the POS: actual covers.

2. Beside each week, the drivers: a week counter for trend, hotel occupancy from the PMS, conference delegates in-house from the events diary, a public-holiday flag, a rain flag from the weather history, and an offer flag.
3. Fit with LINEST, or Data › Data Analysis › Regression, on the first 92 weeks. Hold the last 12 back.
4. Keep a driver only if its t-statistic exceeds 2 and its sign makes sense to an operator.
5. Score on the hold-out: MAPE on weeks the model never saw. If hold-out MAPE is far above training MAPE, the model is memorising, not forecasting.

```
\text{Covers} = b_0 + b_1\,\text{trend} + b_2\,\text{occupancy} +
b_3\,\text{delegates} + b_4\,\text{holiday} + b_5\,\text{rain} +
b_6\,\text{offer}
```

Driver	Coefficient	Reads as
Base	37.1	Covers before any driver
Trend	+0.15 a week	Growth
Hotel occupancy	+99.9 per 100%	Ten points of occupancy add 10 covers
Delegates	+0.11 each	One delegate in nine lunches here
Public holiday	−29.6	Offices closed
Rain	+11.6	Guests stay in
Business-lunch offer	+12.7	Covers the offer brings

Fit: R^2 88%; standard error 6.4 covers; MAPE 3.7% in training and 4.7% on the hold-out. Regression for restaurant demand has a long record (Forst 1992); Lasek, Cercone and Saunders (2016) survey the methods.

9.2 Forecast next Thursday

1. Enter next week's drivers. Specimen: week 105, occupancy 82%, 140 delegates, no holiday, no rain, offer on.
2. Forecast = the fitted equation. Specimen: 163 covers.
3. Range = forecast ± 1.28 × standard error, for 80% confidence. Specimen: 155 to 171.
4. Plan prep and staff to forecast + half a standard error. Specimen: 166.
5. Feed the forecast to the table plan (SOP 6.3) and the roster (SOP 14).

Common errors. Fitting without a hold-out. Keeping drivers because they raise R^2. Forecasting a walk-in day part from a booking curve.

SOP 10 · Nested protection levels and displacement

Hold seats at the peak seating for higher-value guests with nested booking limits, fence the lower classes, and price every group against what it displaces. Sheet 10 Protection.

Owner Revenue Manager. **Frequency** Limits weekly for each peak seating; every group priced on receipt.

10.1 Define the classes

List the ways the same seat is sold, highest spend per cover first, with mean demand and standard deviation from the booking history. Specimen at 20:00: chef's tasting ₹4,500; à la carte ₹1,750; dining offer (25% off) ₹1,300; group set menu ₹950.

10.2 Two classes: Littlewood's rule

Accept a lower-value booking only while the chance that higher-value demand exceeds the seats left is below the price ratio (Littlewood 1972).

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

Specimen: à la carte (₹1,750; mean 55, σ 14) against the offer (₹1,300). Ratio 74%. Protect 46 seats for à la carte.

10.3 Many classes: EMSR-b

1. For each line j, aggregate classes 1 to j: mean = sum of means; σ = √(sum of variances); weighted spend = Σ(spend × mean) ÷ Σ mean.
2. Protection for classes 1 to j = aggregate mean + aggregate σ × NORM.S.INV(1 − next class's spend ÷ weighted spend).
3. Booking limit for class j + 1 = capacity − that protection.
4. Limits nest: a higher class may always take a lower class's seat; never the reverse (Belobaba 1989).

Class	Spend ₹	Demand μ, σ	Protect for this and above	Booking limit
Chef's tasting	4,500	10, 4	11	100
À la carte	1,750	55, 14	61	89
Dining offer	1,300	35, 10	100	39
Group set menu	950	30, 10	—	0

10.4 Fence the classes

Limits only hold if a guest cannot buy the cheaper class at will. Offers only at 19:00 and 22:00. Groups only with a deposit and a set menu. The tasting menu bookable only online with prepayment. Aggregator discounts closed on forecast spill nights.

10.5 Price a group

1. Covers displaced = the lower of group size and (forecast non-group demand + group − seats).
2. Group contribution = covers × (food + beverage per cover) × contribution ratio.
3. Displaced contribution = covers displaced × weighted spend of the displaced classes × contribution ratio.
4. Accept if group contribution ⩾ displaced. Minimum group price = displaced spend × displaced covers ÷ group covers − group beverage.
5. If declined, offer a price, an earlier seating, or a cold night, where the group displaces nobody.

Specimen: 30 covers at ₹950 food plus ₹350 beverage on Saturday at 20:00 bring ₹22,249 of contribution and displace ₹31,962. Net −₹9,713. Quote at least ₹1,545, or the same menu at 19:00 or on a Tuesday.

Common errors. First come, first served at the peak. Offers with no fence. Accepting groups on revenue instead of displaced contribution.

SOP 11 · No-shows and overbooking

Cut no-shows first, then accept bookings up to the limit where one more booking's expected margin equals its expected walk cost. Sheet 11 Overbooking.

Owner Revenue Manager, with the Reservations Manager. **Frequency** No-show rate monthly; the limit weekly for each peak seating.

11.1 Know the rate

Industry data put no-shows near 3.5% and cancellations near 11% of bookings, with late cancellations the costly part; deposits have taken single restaurants from 15% to about 1% (SevenRooms 2025). UK diners run 4–5% (OpenTable). Prepaid ticketing at the top end runs near 1.5%. The specimen uses 12% for no-shows plus cancellations inside 24 hours at the peak: an unguarded book.

11.2 Cut no-shows

1. Measure the rate by seating and party size from twelve weeks of reservations.
2. Card guarantee on every peak-seating booking, with the charge stated. Guests accept a card guarantee far more readily than a non-refundable deposit (Kimes 2009).
3. Confirmation the day before; a call to large parties.
4. Release unconfirmed tables two hours before the seating.
5. Target: from 12% to 5% or less. Specimen: 1.7 tables recovered each peak seating; ₹1.05 lakh of margin a month.

11.3 Set the limit (newsvendor)

1. Cost of an empty table, Cu = average party × dinner check × contribution ratio. Specimen: 2.9 × ₹1,744 × 57% = ₹2,890.
2. Cost of a walk, Co = comped drinks, a voucher and the goodwill lost. Specimen: ₹2,500.
3. Critical ratio = Cu ÷ (Cu + Co). Specimen: 54%.
4. For each booking level n, shows follow a binomial with probability 1 − no-show rate; the workbook sums the exact binomial.
5. Expected value = expected seated × Cu − expected walks × Co. Choose the n with the highest value.

```
n^* = \arg\max_n \Big( \mathbb{E}[\min(S_n, c)]\, C_u - \mathbb{E}[(S_n - c)^+]\, C_o \Big), \quad S_n \sim \text{Bin}(n, 1-p)
```

Specimen: 24 reservable tables. At today's 12% no-shows, book 27: expected walks 0.54 a seating, ₹4,744 more a seating than no overbooking. Once the card guarantee cuts no-shows to 5%, the best book is 25, adding ₹1,251 a seating, ₹27,096 a month. Only that second figure enters the P&L: the guarantee has already recovered most of the tables overbooking would cover.

11.4 House rules

- Overbook only after no-shows have been cut, and re-run the limit when the rate falls.
- Never overbook six-tops or larger. A walked large party is a lost account.
- Keep walk-in tables as the buffer.
- A walked guest is seated within 15 minutes with a drink on the house, or rebooked with a gesture the manager signs off.

Common errors. Overbooking by instinct. Overbooking before taking guarantees. Setting the walk cost at the price of a drink.

SOP 12 · The peak cover charge

On forecast sell-out nights, take a prepaid cover charge, fully redeemable against food and beverage and forfeited on a no-show. It joins SOP 10 and SOP 11 in one policy: it attacks the no-show at source, and it sets a spend floor just below what a peak seat is worth. Sheet 15 Cover Charge.

Owner Revenue Manager, with the Outlet Manager. **Frequency** Switched on or off daily from the forecast; reviewed monthly.

12.1 How the pieces link

- **Overbooking (SOP 11)** manages the risk that a guest does not come, and pays for it in walks.
- **Displacement (SOP 10)** says what a peak seat is worth: the weighted spend of the guests it would turn away. Specimen: ₹1,868 a cover.
- **The cover charge** sits at or below that floor. A guest who dines as planned pays nothing extra. Only two guests feel it: the no-show, who forfeits, and the guest who would have spent less than the seat is worth, who tops up.

12.2 Procedure

1. Trigger: the forecast (SOP 8) shows spill for the seating, or pace is 10% or more ahead of the curve at 14 days.
2. Set the charge at or below the displacement floor, rounded to ₹500. Specimen: ₹1,500 a guest against a ₹1,868 floor.
3. Re-run the overbooking limit at the lower no-show rate. Specimen: no-shows fall to 3% and the optimal book tightens from 27 to 25.
4. Value the policy per seating, in four lines:
    a. Tighter, more certain book = best expected value with the charge − best today.
    b. Forfeits = bookings × no-show rate × party × charge. Pure margin.
    c. Top-ups = guests seated × share below the charge × average top-up × contribution ratio.
    d. Lost bookings = guests who decline to prepay, net of the waitlist that refills them, at Cu a table.
5. Switch on only if the net is positive. It is positive only when the room is turning guests away.

Per sell-out seating	₹
Tighter, safer book	1,785
Forfeited charges	3,268
Spend-floor top-ups	1,748
Lost bookings, refilled from the 17 covers spilled	0
Net	**6,801**

Specimen: ₹17.7 lakh a year across 260 sell-out nights, ₹1.47 lakh a month. Over the card guarantee and the overbooking it replaces, the cover charge adds ₹15,061 a month. That increment, not the gross, is the number in the P&L.

12.3 House rules

- The two prime seatings only. Never 19:00 or 22:00, never breakfast or lunch, never a cold night: on a cold night a charge is a reason not to come.
- Shown at booking and on the confirmation. Redeemable in full. Forfeited on a no-show or a cancellation inside 24 hours.

- Groups at peak pay a deposit on a set menu priced above the displacement floor, not the cover charge.
- In-house guests and the top loyalty tier are exempt.
- Review monthly: no-show rate, bookings lost, waitlist conversion, guest comments.

Common errors. Running it every night. Setting the charge above the displacement floor, which turns a guarantee into a price rise. Booking the gross value instead of the increment over the card guarantee.

SOP 13 · The queue test

Test the peak hour with Erlang C, and plan a luxury room's peak at 80–85% table utilisation. Sheet 12 Queue.

Owner Revenue Manager, Outlet Manager. **Frequency** Quarterly, and whenever peak demand or duration changes.

13.1 Procedure

1. Tables, c: all tables in play at the peak.
2. Arrivals, λ: parties seated plus parties turned away in the busiest hour.
3. Service rate per table, μ = 60 ÷ dinner duration in minutes.
4. Offered load a = λ ÷ μ. Utilisation ρ = a ÷ c. Above 100% the queue grows without limit.
5. Probability a party waits and the average wait, from Erlang C.
6. Standard: average wait under 5 minutes in a luxury room. Above 15, guests leave.
7. If the wait is too long: shorten duration (SOP 7), move demand to warm hours with price (SOP 5), or protect fewer seats for low-value classes (SOP 10).

```
P_{wait} = \frac{\frac{a^c}{c!}\frac{c}{c-a}}{\sum_{k=0}^{c-1}\frac{a^k}{k!}
+ \frac{a^c}{c!}\frac{c}{c-a}} \qquad W_q = \frac{P_{wait}}{c\mu - \lambda}
```

Parties an hour	Utilisation	P(wait)	Average wait
14	65%	2%	0.1 min
16	75%	9%	1.0 min
18	84%	27%	4.7 min
20	94%	63%	27.2 min

Parties an hour	Utilisation	P(wait)	Average wait
22	103%	100%	unbounded

Reading: two more parties an hour take the wait from five minutes to 27. Pooling tables as identical servers is a first view; SOP 6 decides whether the right-sized table is free. Seating walk-ins optimally, rather than first come first served, lifted peak revenue 7.7–13.1% in simulation without longer waits (Bertsimas and Shioda 2003).

Common errors. Reading a half-full room as spare capacity. Planning the peak at 95%. Treating the table count, not the right-sized table, as the constraint.

SOP 14 · Labour to standard

Roster servers hour by hour to the forecast, holding a 20% quality reserve on purpose. Sheet 13 Labour.

Owner Outlet Manager. **Frequency** Roster weekly from the forecast, in 15-minute intervals where the system allows; reviewed daily against the book.

14.1 Procedure

1. Guests seated by hour from the forecasts (SOP 8 and SOP 9) and the RevPASH profile (SOP 4).
2. Productivity p: guests one server holds to standard at once. House setting 12; measure it in a service study.
3. Quality reserve q: the share of each server's capacity withheld. House setting 20% for luxury.
4. Servers required = guests ÷ (p × (1 − q)), rounded up, never below one.
5. Roster to the required column. Put breaks in the cold hours (SOP 5).
6. Check each hour: reserve held = 1 − guests ÷ (scheduled × p). Below q is under-staffed.

```
N_h = \left\lceil \frac{G_h}{p\,(1-q)} \right\rceil
```

Specimen: 82 server-hours required, 77 scheduled. Five hours under-staffed: 09:00, 10:00, 12:00, 13:00 and 14:00. Dinner is staffed to standard; breakfast and lunch are not.

14.2 The principle

A server at 100% load serves covers. A server at 80% load serves guests. The reserve raises labour cost per cover, and that is exactly what the premium price pays for.

Common errors. Rostering by habit instead of forecast. Cutting the reserve to hit a payroll percentage. Breaks in the hot hours.

SOP 15 · Banquets: RevPASF, RevPAST and ConPAST

For function space the scarce unit is a square foot for a session. Measure it three ways, rank it on contribution, read every venue by session, and price every enquiry against what it displaces. Sheet 16 Banquets.

Owner Director of Catering, with the Revenue Manager. **Frequency** Monthly by venue and session; every enquiry priced on receipt; the hot-and-cold calendar set by season.

15.1 The three measures

1. RevPASF = total function revenue (F&B, room hire and other) ÷ square feet, for the month.
2. RevPAST = revenue ÷ (square feet × sessions available). Three sessions a day: morning, afternoon, evening.
3. ConPAST = contribution ÷ (square feet × sessions available), the measure Kimes and McGuire (2001) introduced because margins differ so much by stream.
4. Check the identity: RevPAST = session occupancy × revenue per occupied square-foot session. It is RevPASH in another room.

$$\text{RevPASF} = \frac{\text{revenue}}{\text{sq ft}} \qquad \text{RevPAST} = \frac{\text{revenue}}{\text{sq ft} \times \text{sessions}} \qquad \text{ConPAST} = \frac{\text{contribution}}{\text{sq ft} \times \text{sessions}}$$

Contribution ratios: banquet F&B 40% in the specimen (published ranges run 30–55%); room hire and other 85–95%. A rupee of hire is worth more than two of catering.

Venue	Sq ft	Revenue ₹ lakh	RevPASF ₹	RevPAST ₹	ConPAST ₹	Sessions sold
Grand Ballroom	6,000	255.6	4,260	45.8	21.6	44%
Junior Ballroom	2,400	79.2	3,300	35.5	16.6	45%

Venue	Sq ft	Revenue ₹ lakh	RevPASF ₹	RevPAST ₹	ConPAST ₹	Sessions sold
Garden Lawn	10,000	180.6	1,806	19.4	9.5	15%
Boardroom	600	15.9	2,650	28.5	17.5	54%
All space	**19,000**	**531.3**	**2,796**	**30.1**	**14.4**	**40%**

Reading: the lawn earns the most per event and the least per square foot, because it sells 15% of its sessions. Ranked on ConPAST, the Boardroom beats the Junior Ballroom.

15.2 Space fit

1. Space fit = average guests per event ÷ the venue's banquet capacity. It is to function space what seat fit is to tables.
2. Below 60%: move the event to a smaller room, or divide the room with an air wall and sell the other half.

Specimen: Grand Ballroom 48% (239 guests in a 500-guest room); Junior Ballroom 50%; Garden Lawn 50%; Boardroom 35%.

15.3 Hot and cold by session

1. Session occupancy by venue and session.
2. Hot at 80% or above, cold below 20% (Kimes and McGuire 2001). Set the calendar by month as well: February to May carries 43–45% of India's wedding demand (HVS 2023).
3. Hot sessions: contracted minimum spend per square foot, a forfeitable deposit, a dated release clause. Cold: day packages, MICE, social lunches.

Grand Ballroom	Sessions sold	Occupancy	Period
Morning	6	19%	Cold
Afternoon	9	29%	Warm
Evening	26	84%	Hot

15.4 Price an enquiry against displacement

1. Contribution if it happens = guests × F&B per head × F&B ratio + hire × hire ratio.
2. Expected contribution = contribution × probability it materialises.

3. Compare enquiries for the same session on expected contribution, and per square foot.
4. Minimum spend = the F&B per head at which the weaker enquiry matches the stronger.

Saturday evening	Conference dinner	Wedding sangeet
Guests	250	450
F&B per head	₹1,800	₹2,800
Room hire	₹1.5 lakh	₹3 lakh
Probability	95%	70%
Expected contribution	₹2.92 lakh	₹5.31 lakh

Decision: hold the date for the wedding on a dated option; offer the conference the Junior Ballroom or a Thursday. To win the date, the conference would need ₹4,318 a head in F&B. The banquet version of the peak cover charge is that minimum spend, contracted with a forfeitable deposit.

Common errors. Ranking space on revenue. Quoting a hot Saturday like a cold Tuesday. A 200-guest wedding in a 500-guest ballroom. Confirming the first enquiry without weighing the second.

SOP 16 · The outlet P&L

Close every month on the USALI departmental P&L, read prime cost first, and price every improvement in rupees of departmental profit. Sheet 14 P&L.

Owner Financial Controller, with the F&B Director; reviewed by the Excelsior asset manager. **Frequency** Monthly.

16.1 Procedure

1. Revenue, split food and beverage.
2. Cost of food = food revenue × (theoretical food cost from SOP 3 + the actual-over-theoretical gap from SOP 1). Cost of beverage = beverage revenue × beverage cost %.
3. Payroll = fixed salaried brigade + flexible hours from the roster (SOP 14).
4. Prime cost = cost of sales + payroll. Read it first.

5. Other expenses: variable (china, glass, linen, laundry, supplies, card commissions) and fixed (contracts, licences, base energy).
6. Departmental profit = revenue − prime cost − other expenses.
7. Break-even covers = fixed costs ÷ contribution per cover. Margin of safety = 1 − break-even ÷ actual covers.
8. Flow-through on a new cover at constant price = the contribution ratio. Price every move in SOPs 1–15 at that rate.

USALI's 12th Revised Edition is mandatory for periods from 1 January 2026; close on it.

$$\text{Break-even covers} = \frac{F}{\bar{c} - v} \qquad \text{CM ratio} = \frac{\bar{c} - v}{\bar{c}}$$

Per month	₹ lakh	% of revenue
Outlet revenue	217.8	100.0
Cost of sales	65.0	29.8
Payroll	42.0	19.3
Prime cost	107.0	49.1
Other expenses	28.3	13.0
Departmental profit	82.5	37.9

Specimen: break-even 216 covers a day against 643 served; margin of safety 66%; contribution ratio 57%.

16.2 The programme, priced

Move	Profit a month
Duration programme on sell-out nights, realised at 40% (SOP 7)	₹1.50 lakh
Card guarantee: no-shows 12% → 5% (SOP 11)	₹1.05 lakh
Overbooking at the peak, after the guarantee (SOP 11)	₹0.27 lakh
Two Plowhorses up ₹50 (SOP 3)	₹0.23 lakh

Move	Profit a month
Total	**₹3.05 lakh, +3.7%**
Option: peak cover charge over the guarantee and overbooking (SOP 12)	+₹0.15 lakh

Edition I showed ₹4.14 lakh: it added overbooking valued at 12% no-shows to a card guarantee that cuts them to 5%. Edition II values overbooking after the guarantee. About ₹37 lakh a year, before the table-mix change, the labour fix or the banquet programme, and without a new seat.

Benchmarks (replace with property data): food cost 24–32% actual; beverage 20–26%; payroll 18–28%; prime cost 45–60% in a busy all-day room, 55–70% in fine dining; departmental profit 30–40% all-day, 15–30% signature fine dining. US hotel F&B departments averaged a 29.1% margin, with labour 59.4% of F&B expense (CBRE 2025). In India, F&B and banquets carry 42.3% of hotel revenue (FHRAI and Horwath HTL 2017–18): the department is not an amenity.

Common errors. Reading food cost % without prime cost. Pricing improvements in covers instead of rupees. Ignoring the actual-over-theoretical gap. Booking gross values where only the increment is real.

SOP 17 · Conditions and levers

Before any procedure, score whether revenue management fits the outlet and name the levers it can pull. Sheet 18 Conditions & Levers.

Owner F&B Director, with the Excelsior asset manager. **Frequency** At engagement start and once a year.

17.1 Score the six conditions

1. Score each of Kimes's six conditions from 1 (absent) to 5 (strong): fixed capacity, perishable inventory, time-variable demand, segmentable demand, advance sale, low marginal sales cost.
2. Readiness = average score ÷ 5.
3. Above 80%: implement the whole chain. 60–80%: start with duration and pricing. Below 60%: fix the product first.

Specimen: 87%. The weakest condition is advance sale (3), because lunch is walk-in; that is why lunch is forecast by regression (SOP 9).

17.2 Place the outlet on the price × duration grid

Restaurants sit in Kimes and Chase's third quadrant: fixed price, unpredictable duration. Move toward the second (hotels and airlines) in that order: control duration (SOP 7), then vary price by hour and guest (SOPs 10, 12).

17.3 Name the five levers

Lever	Procedures
Capacity	6, 15, 22
Duration	7, 18
Price	2, 3, 10, 12, 20
Information	8, 9, 21, 25
Sales management	11, 18, 19

The 2025 systematic review in the *British Food Journal* adds information and sales management to the three classic levers.

17.4 Run Kimes's five steps

Baseline, drivers, strategy, implementation, monitoring. Track each step's status on sheet 18; never start strategy without the baseline.

Common errors. Skipping the readiness test. Varying price before duration is under control.

SOP 18 · Reservations and pacing

Sell the book in quarter-hours the kitchen can fire, book each table for its party's real duration, quote waits honestly, account for every way a booking fails to arrive, and climb the deposit ladder only as far as the night needs. Sheet 19 Reservations.

Owner Reservations Manager, with the Executive Chef. **Frequency** Pacing daily; slot lengths and ladder quarterly.

18.1 Pace to the pass

1. Kitchen capacity = covers the line can fire at standard in 15 minutes, from a pass study. Specimen: 24.
2. Load per slot = booked covers + expected walk-ins.
3. Any slot above capacity: move bookings 15 minutes either side before confirming more.

Specimen Saturday: 32 covers at 20:00 against 24. Four slots run over the line.

18.2 Book the table, not the hour

Slot = seated minutes for the party size (SOP 6) + 10 minutes reset, rounded up to 15. Specimen: 90 minutes for twos, 120 for fours, 135 for fives and sixes, 150 for larger. State the slot at booking.

18.3 Quote the wait

Wait = position × minutes between releases of that table size + 5. Quote a little long and seat early. Unexplained, uncertain and unoccupied waits feel longer (Maister 1985): explain, update, and give the waiting guest a drink.

18.4 Count every loss

Covers arriving = booked × (1 − no-shows) × (1 − late cancellations) × (1 − party shrinkage). Specimen: 100 × 95% × 96% × 97% = 88.5. Counting clean no-shows alone misses half the gap.

18.5 The deposit ladder

Rung	No-show rate	When
Nothing	12%	Never at the peak
Reminders and confirmation	8%	Every booking
Card hold, charge stated	5%	Peak seatings
Refundable deposit	5%	Large parties, holidays
Peak cover charge (SOP 12)	3%	Sell-out nights only
Prepaid ticket	1.5%	Chef's table, tasting menu

OpenTable reports deposits cut no-shows by 60% where used, and that UK diners accept a refundable hold of about £35 a head. Rates in the table are house assumptions.

Common errors. Pacing by seats instead of by the pass. One slot length for every party. Quoting short to keep the guest.

SOP 19 · Channels

Measure every booking channel by net contribution a cover, and fence discounted channels out of the hot hours. Sheet 20 Channels.

Owner Revenue Manager. **Frequency** Monthly; fences set weekly by hour.

1. For each channel: share of covers, discount, commission, fixed fee a cover.
2. Net a cover = check × (1 − discount) × (1 − commission) − fee.
3. Contribution = net − variable cost. Compare with a direct cover.
4. At hot hours, sell a channel only if its contribution is at least 90% of direct. Below that, warm and cold hours only.

Channel	Net a cover	vs direct	Hot hours
Direct and in-house	₹1,744	100%	Sell
Reservation platform, ₹60 fee	₹1,684	94%	Sell
Corporate, 10% off	₹1,570	90%	Sell

Channel	Net a cover	vs direct	Hot hours
Bank offer, 15% off	₹1,482	85%	Warm and cold only
Discovery app, 20% off + 6%	₹1,312	72%	Warm and cold only

Specimen leakage: 5.2% of the check. Net RevPASH at 20:00 is ₹1,087 against a gross ₹1,147. Fees are illustrative; load each contract.

Common errors. Judging a channel by covers delivered. Letting an offer run at 20:00.

SOP 20 · Beverage: pour cost and the wine ladder

Cost every pour, build the wine list as a price ladder, and judge it by cash margin and movement up the ladder. Sheet 21 Beverage.

Owner Beverage Manager or Head Sommelier. **Frequency** Pour costs monthly; the ladder at every list revision.

1. Pour cost = bottle cost ÷ bottle ml × pour ml × 1.1 (mixers, garnish, spillage) ÷ price.
2. Set four price points first: accessible, core, premium, prestige. Then source wines to fill them.
3. Multiplier falls as cost rises. Specimen: 3.0×, 2.6×, 2.1×, 1.6×.
4. List price = landed cost × multiplier, rounded up to ₹500, less ₹5.
5. Review by cash margin a bottle, bottles a month and movement between tiers, never by cost percentage alone.

Tier	Landed ₹	List ₹	Cost %	Cash margin ₹
Accessible	2,200	6,995	31%	4,795
Core	3,800	9,995	38%	6,195
Premium	7,500	15,995	47%	8,495
Prestige	18,000	28,995	62%	10,995

Specimen: moving 20 bottles a month from Accessible to Core adds ₹28,000 of margin at a higher cost percentage. Bar seats earn ₹1,500 RevPASH against ₹454 in the dining room. Landed costs in India carry heavy state duties: use excise-inclusive invoices.

Common errors. One multiplier for every bottle. Rejecting a wine for its cost percentage. Pricing before deciding the tiers.

SOP 21 · Guest value

Protect tables by what a guest is worth over time, not by tonight's check. Sheet 22 Guest Value.

Owner Outlet Manager, with the CRM lead. **Frequency** Quarterly.

1. Segment the book: local regular, corporate host, in-house guest, occasion diner, offer seeker.
2. Margin a year = visits × check × contribution ratio (SOP 16).
3. Lifetime value = margin a year × r ÷ (1 + d − r) + margin a year, with retention r and discount rate d (12%).
4. Rank segments; tag guests in the reservation system; protect prime tables for the top two.

Segment	Lifetime value ₹	Rank
Local regular	45,935	1
Corporate host	28,277	2
Occasion diner	3,983	3
Offer seeker	2,922	4
In-house guest	1,665	5

A regular at ₹1,900 a visit is worth eleven times an occasion diner at ₹2,400. In-house guests rank low to the restaurant yet stay exempt from the cover charge, because the hotel owns that relationship.

Common errors. Protecting by price class alone. No guest tags in the book.

SOP 22 · The kitchen as capacity

When the line, not the room, is the constraint, rank orders by contribution per kitchen minute. Sheet 23 Kitchen.

Owner Executive Chef, with the Outlet Manager. **Frequency** Weekly for the peak; on any change to delivery terms.

1. Line capacity = plates an hour the pass can send at standard. Specimen: 96.
2. Utilisation = (dine-in + delivery plates) ÷ capacity. Specimen: 106% at the Saturday peak.
3. Contribution per kitchen minute = (price × (1 − commission) − packaging − plate cost) ÷ minutes. Specimen: dine-in ₹142, delivery ₹72.
4. Above 85%, pause delivery or extend its prep time through the peak. Below it, take delivery: it fills an idle line.

Common errors. Accepting every delivery order at 20:00. Pacing the book without the kitchen (SOP 18).

SOP 23 · The hotel link

Allocate package revenue and judge the outlet inside the hotel's total revenue and profit. Sheet 24 Hotel Link.

Owner Financial Controller, with the Director of Revenue. **Frequency** Monthly.

1. Package breakfasts = rooms × occupancy × share sold with breakfast × guests a room × take-up × days. Specimen: 3,975 a month.
2. Allocate each at stand-alone selling price (USALI). Specimen: ₹650, ₹25.8 lakh a month, 70% of morning revenue.
3. Report F&B (outlet, in-room dining, banquets) as a share of total hotel revenue, and TRevPAR. Specimen December: 60%; TRevPAR ₹19,297.
4. Price the programme per available room: ₹1,388 a month.

In India F&B and banquets carried 42.3% of hotel revenue across the year (FHRAI and Horwath HTL 2017–18); December runs higher on weddings.

Common errors. An unallocated package, which makes breakfast look empty and rooms look rich. Treating the restaurant as an amenity.

SOP 24 · Variability and slack

Plan the peak below the wait cliff, attack variability before adding headcount, pool queues, and price the slack the luxury guest is paying for. Sheet 25 Variability & Slack.

Owner Revenue Manager, with the Outlet Manager. **Frequency** Quarterly, with the queue test (SOP 13).

24.1 The wait cliff

Ten tables, a 60-minute meal: average waits of 1.5, 4.4, 12.3 and 40 minutes at 6, 7, 8 and 9 parties an hour (Erlang C). The engine is ρ ÷ (1 − ρ): 1.5 at 60%, 9 at 90%.

24.2 Variability before headcount

```
W_q \approx \frac{\rho}{1-\rho} \cdot \frac{c_a^2 + c_s^2}{2} \cdot \tau
```

Kingman's approximation: halving arrival variability (from 1.0 to 0.5) cuts the wait to 40% at the same load and payroll. Staggered booking slots (SOP 18) are the tool.

24.3 Pool the queue

Two rooms of five tables at four parties an hour each wait 33 minutes; one room of ten at eight an hour waits 12. One shared queue across cross-trained staff buys waiting time no salary pays for.

24.4 Price the slack

Fixed cost a cover = fixed cost of the service ÷ covers. Specimen: ₹1.2 lakh across 200 covers costs ₹667 a cover at 90% and ₹1,000 at 60%. The ₹333 is what the premium funds: the table that is not crowded and the server who is not stretched (SOP 14).

Common errors. Planning at the daily average. Hiring before smoothing arrivals. Siloed queues.

SOP 25 · Competitor benchmark

Shop the comp set's menu every quarter and read each dish's price index against its menu-engineering class. Sheet 26 Benchmark.

Owner F&B Director. **Frequency** Quarterly, and before every reprint.

1. Choose three to five rooms the same guest would consider for the same occasion.
2. Record their price for each comparable dish.
3. Price index = our price ÷ comp median.

4. Read it with the class (SOP 3): a Star under 95% has headroom; a Plowhorse above 110% is at risk; a Puzzle under 95% is promoted before it is repriced.

Specimen: menu index 96%. Both Plowhorse moves (₹845, ₹745) stay inside the band. The 2025 systematic review names competitor analysis the least-explored area of restaurant revenue management. Rival prices on sheet 26 are placeholders.

Common errors. Pricing from one rival. Shopping without reading the class.

SOP 26 · The technology layer

Buy tools that run this logic faster; reconcile every tool's numbers to its workbook sheet.

Job	What it automates	Sheets
Reservation platform	Booking curve, deposits, slot pacing, guest tags	09, 19, 22
Restaurant RMS and table optimiser	Protection, table assignment, duration	07, 10, 11
AI demand forecasting with rostering	Regression forecast to a 15-minute roster	13, 17
POS and recipe costing	Theoretical food cost, menu mix, check	01–05
Sales and catering system	Space, sessions, displacement	16

Test each vendor on two questions: what is its forecast error (SOP 8), and can it show its displacement logic (SOP 10)? A tool that cannot answer both is a black box, not a revenue system.

Operating rhythm

Revenue management is a calendar, not a project.

Cadence	Routine	Owner	SOP
Daily	Pace against the curve for the next 14 days	Revenue Manager	8
Daily	Protection and overbooking limits for tonight; cover charge on or off	Revenue Manager	10, 11, 12

Cadence	Routine	Owner	SOP
Daily	Roster against the book; tools set for each period	Outlet Manager	5, 14
Weekly	Curve refreshed; spill nights flagged; regression forecast for walk-in day parts	Revenue Manager	8, 9
Weekly	Hourly RevPASH review	Outlet Manager	4
Weekly	No-show list; groups and banquet enquiries priced on displacement	Reservations, Catering	10, 11, 15
Monthly	Actual against theoretical food cost	Cost Controller	1
Monthly	MAPE and bias	Revenue Manager	8, 9
Monthly	RevPASF, RevPAST, ConPAST by venue and session	Director of Catering	15
Monthly	Departmental P&L, prime cost, flow-through	Financial Controller	16
Quarterly	Yield tests; menu engineering; regression refit; queue test	Chef, Revenue Manager	1, 3, 9, 13
Twice a year	Duration study; table-mix review	Outlet Manager	6, 7

Re-measure after 8 to 12 weeks on six numbers: RevPASH, seat fit, dinner duration, forecast error, the food-cost gap and ConPAST.

Edition II routines. Daily: slot pacing against the kitchen (SOP 18) and delivery throttling at the peak (SOP 22). Weekly: channel fences by hour (SOP 19). Monthly: pour costs, the wine ladder and package allocation (SOPs 20, 23). Quarterly: guest value, the queue and variability test, and the competitor shop (SOPs 21, 24, 25). Yearly: the readiness score (SOP 17).

Glossary

- **AP / EP**: as-purchased and edible-portion weight or cost.
- **Q-factor**: a percentage added to plate cost for uncosted items.
- **Contribution margin (CM)**: price less plate cost; at outlet level, revenue less all variable costs.

- **RevPASH**: revenue per available seat-hour.
- **Seat fit**: guests ÷ seats at occupied tables. **Space fit**: guests ÷ banquet capacity.
- **Combinable tables**: tables that push together to seat larger parties.
- **Realisation factor**: the share of the arithmetic gain from shorter meals that actually arrives; 40% in the workbook.
- **Pickup** and **pace**: covers added before the service; covers on the books against the curve.
- **Regression**: a model that forecasts demand from its drivers. **Hold-out**: weeks kept back to test it.
- **MAPE** and **bias**: average absolute percentage error; average signed error.
- **Protection level** and **booking limit**: seats held for higher classes; the most a class may book, nested.
- **Displacement**: the contribution a group or event pushes out.
- **Critical ratio**: Cu ÷ (Cu + Co), the newsvendor's balance point.
- **Cover charge**: a prepaid, redeemable charge per guest on sell-out nights, forfeited on a no-show.
- **RevPASF, RevPAST, ConPAST**: revenue per square foot; revenue per square foot per session; contribution per square foot per session.
- **Prime cost**: cost of sales plus payroll.

Sources

Restaurant revenue management. Kimes, S. E., Chase, R. B., Choi, S., Lee, P. Y. and Ngonzi, E. N. (1998), "Restaurant revenue management: applying yield management to the restaurant industry", *Cornell HRA Quarterly* 39(3). Kimes, S. E. (1999), "Implementing restaurant revenue management: a five-step approach", *Cornell HRA Quarterly* 40(3). Kimes, S. E. (2004), "Restaurant revenue management: implementation at Chevys Arrowhead", *Cornell HRA Quarterly* 45(1). Kimes, S. E. and Thompson, G. M. (2004), "Restaurant revenue management at Chevys: determining the best table mix", *Decision Sciences* 35(3). Kimes, S. E., Cornell SHA705 and SHA708 course materials.

Tables and duration. Thompson, G. M. (2002), "Optimizing a restaurant's seating capacity: use dedicated or combinable tables?", *Cornell HRA Quarterly* 43(4). Thompson, G. M. (2009), "(Mythical) revenue benefits of reducing dining duration in restaurants", *Cornell Hospitality Quarterly* 50(1). Kimes, S. E., Wirtz, J. and Noone, B. M. (2002), "How long should dinner take? Measuring expected meal duration for restaurant revenue

management”, *Journal of Revenue and Pricing Management* 1(3). Noone, B. M., Kimes, S. E., Mattila, A. S. and Wirtz, J. (2007), “The effect of meal pace on customer satisfaction”, *Cornell HRA Quarterly* 48(3). Bertsimas, D. and Shioda, R. (2003), “Restaurant revenue management”, *Operations Research* 51(3).

Menus. Kasavana, M. and Smith, D. (1982), *Menu Engineering*. Miller, J. (1980), *Menu Pricing and Strategy*. Pavesic, D. (1983), cost-margin analysis. Yang, S. S. (2012), “Eye movements on restaurant menus: a revisitation on gaze motion and consumer scanpaths”, *International Journal of Hospitality Management* 31(3). Yang, S. S., Kimes, S. E. and Sessarego, M. M. (2009), “Menu price presentation influences on consumer purchase behavior in restaurants”, *International Journal of Hospitality Management* 28(1).

Forecasting, protection, reservations. Littlewood, K. (1972), “Forecasting and control of passenger bookings”, AGIFORS. Belobaba, P. (1989), “Application of a probabilistic decision model to airline seat inventory control”, *Operations Research* 37(2). Forst, F. G. (1992), “Forecasting restaurant sales using multiple regression and Box-Jenkins analysis”, *Journal of Applied Business Research* 8(2). Lasek, A., Cercone, N. and Saunders, J. (2016), “Restaurant sales and customer demand forecasting: literature survey and categorization of methods”. Kimes, S. E. (2009), guest attitudes to restaurant reservation policies, Cornell Center for Hospitality Research. SevenRooms (2025) and OpenTable, industry no-show data. Erlang, A. K. (1917), queueing.

Function space and finance. Kimes, S. E. and McGuire, K. A. (2001), “Function-space revenue management: a case study from Singapore”, *Cornell HRA Quarterly* 42(6). HVS (2023), Indian wedding market. HFTP and AHLA, *Uniform System of Accounts for the Lodging Industry*, 12th Revised Edition. CBRE (2025), hotel F&B trends. FHRAI and Horwath HTL, *Indian Hotel Industry Survey 2017–18*.

Companion files: the Practice Bible deck and *Excelsior Restaurant Revenue Management: The Mathematics*.

Non multa, sed multum.

Edition II additions. 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). Kimes, S. E. and Wirtz, J. (2002), “Perceived fairness of demand-based pricing for restaurants”, *Cornell HRA Quarterly* 43(1). Maister, D. (1985), “The psychology of waiting lines”. Kingman, J. F. C. (1961), heavy-traffic queue approximation. Restaurant revenue management: a systematic literature review and future challenges (2025), *British*

Food Journal 127(6). OpenTable Restaurant Impact Report, via Restaurant Online (6 October 2026). Engineering the wine list: price laddering as a revenue system, Hospitality Net (2026). Excelsior, June 2026 posts on queuing and the economics of the empty table.

Companion files: the Practice Bible, Edition II, and *Excelsior Restaurant Revenue Management: The Mathematics, Edition II*.