

METHODOLOGY V1.0

The methodology behind every Verdict.

5 layers. 6 classifications. 4 tier voice rules. One framework — versioned, cited, transparently the same for every property.

VERSION

v1.0

LAST VERIFIED

2026-05-15

METHODOLOGY AUTHOR

Sagar Sharma

FRAMEWORK CITATIONS

Kasavana 1982 · Cornell · STR · HVS

SECTION 01 · FRAMEWORK

What the Couverte Methodology Is

Every menu analysis you have ever seen runs on a single model. One model asks whether a dish is popular and profitable. Another asks how its profit compares to its food cost. A third asks how much revenue each seat earns per hour. Each one is right about the thing it measures and blind to everything else.

The Couverte Methodology does not pick a model. It runs the credible ones together, because each covers a gap the others leave open — and then it does the work no model does on its own: it tells you what to change, in language that fits the kind of restaurant you actually run.

The shape of the framework

The methodology has four moving parts.

Five layers. Three of them are analytical — they read what the numbers say about each dish, the seat it occupies, and the market around it. One is contextual — what the world outside your kitchen demands. One is prescriptive — the specific action, written so your F&B manager can execute it on Thursday.

Six item classifications. The classic menu-engineering grid sorts dishes into four boxes — Stars, Plowhorses, Puzzles, Dogs. We keep those four and add two of our own, so a dish that looks safe but is exposed to ingredient prices, or one that is one small adjustment away from earning its keep, gets named for what it actually is.

Four restaurant tiers. A local trattoria and a Michelin tasting room do not optimise for the same thing. The same finding produces a different recommendation — and a different tone of voice — depending on whether you are a casual neighbourhood spot, a hotel F&B operation, a fine-dining room, or a Michelin kitchen.

A critic pass. Before a recommendation is written into a Verdict, it is reviewed — cross-checked against the other models, sense-checked against kitchen reality, and held to the standard of advice you could hand a manager without a footnote.

What you get out of it

The output is not a dashboard. It is a short, ranked set of actions per item: a diagnosis, the single most impactful change, an implementation detail specific enough to brief a service team, and a window for when to look again.

Everything in this methodology is versioned and cited. The reasoning traces to decades of published menu-engineering, revenue-management, and pricing-psychology research — the full source list is in the provenance section. What is new is the synthesis: no published model combines these layers with weights that shift by restaurant type, and none of them carry the qualitative judgment of whether a dish belongs on *your* menu at all.

The nine sections that follow walk the framework end to end — the five layers, the six classifications, the four tiers, the critic and QA pass, and the provenance behind all of it.

Couverte Methodology v1.0.

SECTION 02 · THE FIVE LAYERS

5 layers

Layer 1 — Item Economics

SCORE WEIGHT (CASUAL TIER) **30%***“What does this dish actually earn us?”*

Layer 1 is the foundation. It is pure financial analysis of every item on your menu — no psychology, no market, just what the dish costs to put on the plate and what it brings back. In the Couverte Score for a casual restaurant, the economics carry the most weight of any single component, because for that kind of operation the numbers have to work above all else.

A one-sentence summary of the math

Take each dish’s selling price, subtract what the food costs to produce it, and you have its contribution margin — then weigh that margin by how often the dish actually sells.

What it measures

Layer 1 needs four inputs per item: the selling price, the recipe cost per portion, the units sold over a real period (four to eight weeks minimum — a single weekend lies to you), and the total number of items on the menu.

From those it calculates the figures every operator should know cold:

- **Contribution margin** — selling price minus food cost. The cash the dish leaves on the table.
- **Food cost percentage** — food cost divided by selling price. The classic margin-and-popularity grid ignores this entirely, which is how a dish with a fat contribution margin and a dangerous 42% food cost gets called a winner.
- **Menu mix percentage** — a dish’s share of total units sold. How much of the menu’s volume it carries.
- **Weighted contribution margin** — margin multiplied by units sold. A dish earning a small margin on heavy volume can out-earn a high-margin item nobody orders.

It then sorts each dish into the classification grid — and adds a **cost-risk flag** the classic grid does not have. A top performer whose food cost runs well above its category average is not simply a Star; it is exposed to the next ingredient price spike. A steady mid-margin seller with a low food cost is not simply a Plowhorse; it is a candidate to become a Star with one small move.

One concrete intervention

A steady seller — good volume, modest margin — comes back with a food cost percentage below its category average. That low food cost is headroom. Layer 1 flags it as a candidate to convert: a small price increase, on the order of **+€2.10**, lifts the contribution margin into Star range while the low food cost keeps the change low-risk. You are not inventing demand. You are repricing margin the dish was already strong enough to carry.

Layer 2 — Time-Space Efficiency

SCORE WEIGHT (CASUAL TIER) 10%

“How hard is this dish working for us relative to the seat and hour it occupies?”

A dish does not exist on a spreadsheet. It exists in a service period, at a table, in a dining room with a fixed number of seats. Layer 2 measures the dish against that reality — the seat it holds and the hour it holds it for. A tasting plate can score beautifully on margin and still be the wrong dish if it pins a four-top for ninety minutes during your busiest service.

A one-sentence summary of the math

Divide the profit a dish generates by the seats and hours it consumes — so a dish is judged not just on what it earns, but on what it earns *per seat, per hour*.

What it measures

Layer 2 needs the outlet’s seat count, the operating hours per service period, the average dining duration by course type, and seat occupancy data where the reservation system can provide it.

From those it calculates the time-and-space metrics:

- **RevPASH** — revenue per available seat hour. How much top-line each seat generates per hour of service.
- **ProPASH** — *profit* per available seat hour. The sharper figure: a busy seat turning low-margin covers can post a strong RevPASH and a weak ProPASH. Revenue and profit do not always move together, and ProPASH is where that gap shows.
- **Seat utilisation ratio** — covers served against the turns the room could physically deliver. A casual room running well below its possible turns is leaving capacity unsold every service.
- **Item ProPASH contribution** — each dish’s weighted margin expressed against the seats and hours it occupies.

What Layer 2 means depends on the restaurant. For a casual room, turns are the whole game. For hotel F&B, ProPASH by daypart shows where the real money is — breakfast can be high-volume and low-margin while dinner is the opposite. For fine dining, low turnover is the point, so the figure that matters is profit per cover. For a Michelin kitchen, the tasting menu fixes both duration and price, and time-space efficiency nearly drops out as a primary measure.

One concrete intervention

A casual lunch outlet is turning its tables noticeably below the number the room could deliver. Layer 2 reads that gap as unsold capacity and points the recommendation at speed, not price: move two slow-firing dishes off the express-lunch section and replace them with items the kitchen can plate without bottlenecking the pass. The covers were always available — the menu was holding the seats hostage.

Layer 3 — The Couverte Score

SCORE WEIGHT (CASUAL TIER) 25%

“One number that tells you how this item is really performing.”

Layer 3 is the proprietary part. It takes everything Layer 1 and Layer 2 produce and meshes it into a single score per dish — weighted by the kind of restaurant you run. It is the layer that turns a wall of metrics into one figure an operator can rank a menu by.

A one-sentence summary of the math

Normalise five components to a common 0–100 scale, multiply each by a weight that shifts with your restaurant tier, and add them up — so the same dish scores differently at a casual room than it does at a Michelin kitchen.

What it measures

The Couverte Score combines five components: contribution margin, popularity, food cost efficiency, ProPASH contribution, and brand alignment. Each is normalised to 0–100 within its category, then weighted.

The weights are not fixed. They shift across the four tiers:

- At a **casual** restaurant, the economics dominate — margin, velocity, and food cost efficiency together carry most of the score.
- At **hotel F&B**, ProPASH rises in weight, because profit per seat-hour across dayparts is where the real money is.
- At a **fine-dining** room, brand alignment climbs sharply.
- At a **Michelin** kitchen, brand alignment is the dominant weight — creative coherence and guest experience come first, profit second.

Brand alignment is the component no published model includes. It answers one question: does this dish belong on *this* restaurant’s menu? A truffle pizza can score high on margin and popularity at a casual Italian spot and still be wrong if the place positions itself as a simple, honest trattoria. A low-margin amuse at a Michelin restaurant can score poorly on economics and still be essential to the tasting narrative. It is assessed against the restaurant’s stated concept, its ingredient-sourcing identity, the kitchen’s technique level, and whether the dish matches guest expectations for its price tier.

One concrete intervention

A signature pasta at a fine-dining room scores in the mid-range — solid margin, modest velocity. At a casual tier, that score would push a recommendation to reprice or reposition for volume. But at the fine-dining tier the weight on brand alignment is high, and this dish scores at the top of that component: it is exactly what the restaurant is known for. The Score reads the dish as a brand anchor, not a margin problem, and the recommendation protects it — hold the price, keep its menu position — rather than chasing turns it was never meant to deliver.

Layer 4 — Market Context

SCORE WEIGHT (CASUAL TIER) 5%

“What does the world outside this kitchen demand?”

Layer 4 wraps the item-level analysis in market reality. The same dish, in the same kitchen, can earn a different recommendation depending on what is happening in the market around it. The first three layers can be done by anyone with a spreadsheet and patience. This one cannot — it needs data from beyond your four walls.

A one-sentence summary of the math

Compare each item’s price and each menu’s overall position against aggregated data from comparable restaurants in the same market — so a recommendation reflects not just your numbers, but where your numbers sit against everyone else’s.

What it measures

Layer 4 reads three things:

- **Competitive positioning** — where each item’s price sits relative to comparable dishes at comparable restaurants in the same market, whether the menu’s overall price point is aligned to its tier and location, and whether there are price-convergence risks, such as casual-dining items drifting close to fast-casual prices and triggering a value-perception problem.
- **Consumer-segment behaviour** — how spending is moving by income band. High-income guests are spending freely at fine dining and upscale casual; middle-income guests are pulling back from fast casual, trading down to value or up to “worth it” experiences. The casual price battleground is where that pressure concentrates.
- **Seasonal and demand patterns** — which items show real seasonal demand curves, where there are daypart opportunities, and how the local event calendar moves demand.

This is the layer where the **benchmark database** becomes critical. Saying “your grilled sea bass is priced below the market median for comparable hotel restaurants in this city, and items at the next price point up show stronger wine-pairing attach rates” requires aggregated data across many properties. It is the part of the methodology that grows sharper with every restaurant on the network — a dataset no single operator can assemble alone.

One concrete intervention

A 5-star hotel restaurant prices its grilled sea bass conservatively. Layers 1 through 3 see a healthy, well-behaved dish and leave it alone. Layer 4 sees something the in-house view cannot: the price sits below the market median for comparable rooms in that city, and the gap is pure margin left on the table. The recommendation closes part of that gap — a measured increase toward the market band — with the confidence that comparable guests are already paying it elsewhere.

Layer 5 — Prescriptive Actions

SCORE WEIGHT (CASUAL TIER) 30%

“Here’s what to do, in language that fits your operation.”

Layer 5 is where the four analytical layers converge into something you can act on. This is the layer that separates the methodology from a report. The key insight: the *same* analytical finding produces a *different* recommendation depending on the tier.

A one-sentence summary of the math

There is no formula here — Layer 5 is judgment: it takes the score, the classification, and the market read, and writes the single highest-impact change, calibrated to what your kind of restaurant can realistically execute.

What it measures

Layer 5 does not measure — it prescribes. Take a Puzzle: a dish with a high margin and low popularity. The finding is identical across every tier. The recommendation is not:

- At a **casual** restaurant, the move is mechanical and volume-focused — reposition the dish to a prime spot on the menu, sharpen its description, test its price elasticity.
- At **hotel F&B**, it becomes a service play — feature it as a chef’s recommendation, brief the team on a two-sentence spoken description, test it through room service.
- At a **fine-dining** room, it becomes a menu-architecture play — work it into the tasting menu, or position it next to the highest-priced item as an anchor.
- At a **Michelin** kitchen, it becomes a narrative question — does the dish serve the menu’s arc? If it does, treat it as a signature course; if its popularity is low because it sits at the wrong moment in the sequence, reposition it rather than reprice it.

Every action plan, whatever the tier, has the same backbone: a **diagnosis** (which classification, which score components are weak, and why), a **primary action** (the single most impactful change — price, position, description, portion, or removal), an **implementation detail** specific enough to hand to the F&B manager, an **expected impact**, and a **measurement window** for when to re-evaluate, typically four to eight weeks out. Premium tiers add an experience-enhancement step — a hospitality moment, not just a profit move.

Menu psychology is woven in here, not bolted on. How a price is displayed, how long a description runs, how a section is laid out — all of it is calibrated by tier inside the recommendation itself.

One concrete intervention

A high-margin starter is barely selling at a casual bistro — a textbook Puzzle. Layer 5 does not say “promote it.” It writes the action: move the dish to the top-right of a two-panel menu, give it a short description that leads with provenance, and trim the price by a small, testable amount to probe demand. Then it sets a four-week window to look again — because a recommendation without a review date is a guess, not a plan.

SECTION 03 · CLASSIFICATIONS

The Six Item Classifications

Classic menu engineering sorts every dish into a four-box grid: high or low popularity, against high or low contribution margin. It is a good start and a famous one. But four boxes hide two things that matter — a dish that looks like a winner while quietly carrying dangerous cost, and a dish that looks ordinary while sitting one small move from greatness.

So the Couverte Methodology keeps the classic four and adds two of its own. Six classifications instead of four. Each one comes with its own kind of action.

The classic four

Star — high popularity, high contribution margin. The dishes guests love that also pay well. The action is not to change them; it is to protect them. Hold the price, hold the menu position, and keep an eye on the food cost so a Star does not quietly drift into trouble. *Sample intervention: protect — lock menu position, no price change, monitor food cost monthly.*

Plowhorse — high popularity, low contribution margin. The crowd-pleasers that do not earn enough. They carry your volume but not your margin. The action is to lift the margin without breaking the demand — a modest price increase, a portion adjustment, or a cheaper-to-plate garnish. *Sample intervention: reprice +€2.10 to lift contribution margin, then watch volume for four weeks.*

Puzzle — low popularity, high contribution margin. The dishes that pay well but nobody orders. The margin is already there; the problem is visibility. The action is promotion — better menu position, a sharper description, a price test to find the demand. *Sample intervention: move to a prime menu position and rewrite the description to lead with provenance.*

Dog — low popularity, low contribution margin. Low demand, low margin. They earn little and occupy menu space and kitchen attention. The action is usually removal, or a rework significant enough to change both the demand and the economics. *Sample intervention: remove from the menu, or rework the recipe and price before the next print run.*

The two Couverte additions

Vulnerable Star — a dish classified as a Star, but with a food cost percentage running well above its category average. Profitable today, exposed tomorrow. It looks like a winner and will keep looking like one right up until the next ingredient price spike compresses its margin. The action is to defend the margin before the market forces the issue — re-cost the recipe, lock a supplier price, or adjust the portion while the dish is still strong. *Sample intervention: re-cost the recipe and trim the portion slightly to pull food cost back toward the category average.*

Conversion Candidate — a Plowhorse with a food cost percentage below its category average. That low food cost is headroom. A small price or portion adjustment can shift the dish into Star territory with very little risk, because the cost base gives it room to absorb

the change. The action is a low-risk conversion play — make the small move and watch the dish climb. *Sample intervention: reprice +€2.10 — the below-average food cost absorbs the increase and lifts the dish toward Star.*

SECTION 04 · RESTAURANT TIERS

The Four Restaurant Tiers

The same finding does not mean the same thing everywhere. A low table-turn number is a crisis at a casual lunch spot and a non-issue at a Michelin kitchen. A high-margin dish that nobody orders gets repriced at one kind of restaurant and reframed as a narrative problem at another.

So the methodology runs on four tiers. The tier you select changes the scoring weights, changes which metrics lead, and — just as importantly — changes the *voice* of the recommendation. Each tier has a way it speaks, drawn from how that kind of restaurant displays prices, writes menu descriptions, and makes decisions.

Tier 1 — Local / Casual

The economics come first and the room runs on table turns. Recommendations are direct, mechanical, and volume-focused.

Voice rule: Plain and practical. Charm pricing is acceptable. Menu descriptions run longer — roughly eight to fifteen words — and lead with value: house-made, locally sourced, served with. Currency symbols are optional. The recommendation reads like a clear instruction to a working operator who has five other jobs today.

Sample copy line: “Move the chicken Milanese to the top-right panel, stretch the description to twelve words, and test a small price cut — we’ll know in four weeks.”

Tier 2 — Hotel F&B

Profit per seat-hour across dayparts is where the real money is. The F&B experience shapes the whole property’s reputation, so recommendations balance the numbers with the guest.

Voice rule: Balanced — provenance and technique together. Whole numbers are preferred, no currency symbols on printed menus. Descriptions run twelve to twenty words. Recommendations often become service plays: chef’s recommendations, spoken descriptions for the team, room-service tests. The voice frames a change as something that serves the guest *and* the bottom line, never as cold optimisation.

Sample copy line: “Feature the sea bass as the chef’s recommendation at dinner, brief the team on a two-sentence description, and test it as a highlighted room-service item.”

Tier 3 — Fine Dining

Turnover is intentionally low — the figure that matters is profit per cover. Brand alignment climbs sharply in the score. Recommendations shift toward per-cover spend and perceived value.

Voice rule: Precise and restrained. No currency symbols, no decimals, no trailing zeros — the price follows the description in the same font, never in a column. Descriptions run short, four to ten words, with technique and ingredient doing the work and no decorative adjectives. The voice is composed and confident; it respects the room.

Sample copy line: “Position the langoustine course beside the highest-priced plate as an anchor, and lift the tableside element to raise perceived value without raising cost.”

Tier 4 — Michelin

Time-space efficiency nearly drops out as a primary measure — the tasting menu fixes both duration and price. Brand alignment is the dominant weight in the score. Creative coherence and guest experience come first; profit follows.

Voice rule: Spare and deferential to the kitchen’s intent. Three to eight words, or ingredient-only listings — the menu *is* the brand. No currency symbols. The voice never instructs the chef on creative matters; it raises questions about coherence and sequence and leaves the cuisine to the kitchen.

Sample copy line: “If this course serves the menu’s narrative arc, treat it as a signature moment; if its low draw is a question of sequence, reposition it rather than reprice it.”

SECTION 05 · CRITIC + QA

The Critic and QA Pass

The methodology's first principle is that no single model can be trusted on its own. Each classic model is right about one thing and blind to the rest — which is why the analytical layers run several of them together and let each one cover the others' gaps.

That same principle does not stop at the analysis. It applies to the recommendation itself. Before a Layer 5 action becomes a Verdict you act on, it goes through a critic and QA pass — a deliberate review step whose job is to catch the recommendation that is technically correct and operationally wrong.

What the critic checks

The critic holds a draft recommendation against four tests, in roughly this order.

Does it survive the cross-check? A finding is only as good as its agreement across models. A dish flagged on contribution margin alone, with no support from the food-cost or time-space view, is a weaker signal than one the models agree on. The critic looks for that agreement before it lets a strong-sounding action through.

Does it fit the tier? The same finding should produce a different action — and a different voice — depending on the restaurant. A recommendation that tells a Michelin kitchen to reprice for volume, or hands a casual operator a paragraph of fine-dining restraint, has failed the tier test even if the numbers are sound.

Does it survive contact with the kitchen? A recommendation that ignores kitchen reality gets ignored by the kitchen. The critic checks labour intensity and station load: an action that improves margin but demands an extra prep station during peak service is flagged, not shipped.

Does it belong on this menu? This is the brand-alignment test. A dish can score well on every economic measure and still be the wrong dish for the restaurant's stated concept. The critic carries the qualitative judgment that a spreadsheet cannot — whether the recommendation serves the restaurant's identity, not just its P&L.

What QA confirms

Where the critic asks whether the recommendation is *right*, QA confirms it is *complete*. Every action plan that ships has the same backbone, and QA checks each part is present and specific:

- a **diagnosis** — the classification, the weak score components, and why
- a **primary action** — the single most impactful change, not a list of five
- an **implementation detail** — specific enough to hand to the F&B manager, with a name and a day attached
- an **expected impact** — labelled as an estimate, never as a promise
- a **measurement window** — a date to look again, typically four to eight weeks out

A recommendation without a diagnosis is an opinion. One without an implementation detail cannot be executed. One without a measurement window is a guess that never gets checked. QA is the step that refuses to let any of those reach you.

Why the pass exists

The output of this methodology is meant to be consulting-grade — advice you could hand to a working F&B leader and trust them to act on without a second analyst in the room. The critic and QA pass is what earns that standard. It is the difference between a number on a screen and a decision you can stand behind.

SECTION 06 · VERSION + PROVENANCE

Provenance and Sources

The Couverte Methodology is a synthesis, not an invention. What is new is the way the pieces are combined and weighted; the pieces themselves come from decades of published research. This section gathers every framework citation in one place. The work is cited here, in full, so the methodology can be checked rather than taken on trust.

Foundation — item economics (the 1980s bedrock)

The base layer of item economics rests on three foundational works.

- **Menu Engineering: A Practical Guide to Menu Analysis** — Kasavana & Smith (1982), Hospitality Publications / Michigan State University. The original popularity-and-contribution-margin grid, the menu-mix concept, and the popularity threshold that the classification logic still uses.
- **Cost/Margin Analysis: A Third Approach to Menu Pricing and Design** — Pavesic (1985), *International Journal of Hospitality Management*. Weighted contribution margin that factors in food cost percentage — the correction that lets the methodology flag a profitable-looking dish whose food cost is quietly dangerous.
- **A Third Cost-Control Perspective** — Miller (1980). The food-cost-and-volume view, which keeps cost discipline in the picture alongside margin.

A comprehensive *FIU Hospitality Review* literature review sets these three side by side and explains why no single one of them is sufficient on its own.

Evolution — revenue and space management (1999–2017)

The time-and-space layer draws on the modern revenue-management literature.

- **Implementing Restaurant Revenue Management: A Five-Step Approach** — Kimes (1999), *Cornell Hotel and Restaurant Administration Quarterly*. The RevPASH framework — revenue per available seat hour — and the principle that revenue management applies to F&B, not just hotel rooms. The fuller treatment is published through Cornell eCommons.
- **New Performance Indicators for Restaurant Revenue Management: ProPASH and ProPASM** — Heo (2017), *International Journal of Hospitality Management*. Profit per available seat hour, and the critical proof that a rising RevPASH does not always mean rising profit.
- **Reevaluating the RevPASH Formula** — Kalan (2023), Hospitality Net. The seat-utilisation ratio, a capacity-efficiency measure that complements ProPASH.

Cornell's eCornell course materials on contribution-margin analysis bridge the academic theory and its practical teaching.

Pricing psychology (scoring and prescriptive actions)

How the methodology advises on price display, menu description, and visual hierarchy comes from the pricing-psychology literature.

- **Priceless: The Myth of Fair Value** — Poundstone (2010), Hill and Wang. Anchoring, decoy pricing, charm pricing against whole-number pricing, and the effect of removing currency

symbols.

- **Menu Engineering Methodology** — Rapp, menuengineers.com. Decades of applied menu psychology: visual hierarchy, two-panel menu layout, the effect of description length, and the diminishing returns of more than one highlighted item per section.
- Empirical research on the Michelin-star pricing premium, and the Chef's Pencil global analysis of Michelin dining costs, inform the price calibration for the upper tiers.

Hospitality philosophy (the framing of recommendations)

These works shape the tone of the recommendations, especially at the premium tiers — not “raise the price” but how to improve the guest experience and the bottom line together.

- **Unreasonable Hospitality** — Guidara (2022), Optimism Press. The experience-first philosophy and the practice of investing deliberately in memorable guest moments.
- **Setting the Table** — Meyer (2006), HarperCollins. Enlightened hospitality, and the case that hospitality-first businesses outperform transaction-first ones.

Operations and implementation

These keep the recommendations grounded in kitchen and operational reality.

- **Kitchen Confidential** — Bourdain (2000), Ecco. Not a methodology text, but essential for understanding station bottlenecks and the reality of line cooking under pressure — which is why labour intensity is weighed before a recommendation ships.
- **The Restaurant Manager's Handbook** — Brown and contributors, Atlantic Publishing Group. The operational reference behind the implementation detail in every action plan.

Market intelligence (refreshed as new editions publish)

The Market Context layer draws on living industry sources, refreshed as new editions appear: the McKinsey consumer restaurant research, the National Restaurant Association's State of the Industry report, and market-bifurcation analysis from Restaurant Dive and Placer.ai. These are the published market signals; the proprietary benchmark network that grows with every property on the platform is described in the Market Context layer.

Ongoing research

The methodology is maintained against current research, principally *Cornell Hospitality Quarterly* and the *International Journal of Hospitality Management* — the journals where the foundational menu-engineering work was first published.

Couverte Methodology v1.0. Last updated April 2026. This provenance section is the single, authoritative home for the methodology's framework citations.