

Second-generation menu engineering: from popularity matrices to *price elasticity modeling by segment*

By  **Diego F. Parra** · Updated 2026-07-07 · Menu & Menu Engineering

MASTERRESTAURANT®

White Paper

Ingeniería de menú de segunda generación: de la matriz de popularidad al modelado de elasticidad de precios por segmento

Método probado en +8.400 restaurantes · 43 países

costorestaurante.com

QUICK VERDICT

Verdict: the popularity-margin matrix (Star / Plowhorse / Puzzle / Dog) is a static diagnostic that no longer defends margin in a volatile-input environment. Second-generation menu engineering adds three layers the classic matrix ignores: price elasticity by segment, theoretical vs. actual cost per dish, and inflation scenario simulation. In a full-service restaurant with a 30 % average food cost, a well-calibrated elasticity model recovers 2-4 points of operating margin without losing traffic. Diego F. Parra and the Masterrestaurant method start with the 10 dishes that concentrate 60 % of the sales mix, on real data from 8,400 restaurants in 43 countries.

 **White Paper** · Technical document · C-Suite & multilateral banking · 17 min read · 2026-07-07

INTELLECTUAL PROPERTY OF MASTERRESTAURANT® — EXCLUSIVE FOR SECTOR LEADERS

Classic menu engineering was born in 1982 (Kasavana and Smith, Michigan State University) to classify dishes by popularity and contribution margin. It worked for four decades because inputs rose 2-3 % a year and an annual map was enough. In 2026, with double-digit volatility in proteins and dairy —the FAO food price index has swung with year-over-year peaks above 20 %— that static map leaves money on the table every week. This white paper documents, in six chapters, why the 1982 model expired and how to replace it without throwing it away.

The document follows a white-paper logic: first the problem diagnosis (why the matrix fails), then the second-generation framework with its three new layers —price elasticity by customer segment, theoretical vs. actual cost variance, and inflation scenario stress—, then a quantified mini-case with real anonymized cash figures, and finally the model's limitations, assumptions and failure conditions. The goal is a defensible operating margin under uncertainty, calibrated with the Masterrestaurant method, not a nice chart on the wall. Every figure is cited with its source in prose.

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	CLASSIC POPULARITY MATRIX	2ND-GEN ELASTICITY MODELING
Decision variables	✗ 2 (popularity + margin)	✓ 5 (popularity, margin, elasticity, variance, scenario)
Recalculation frequency	✗ 1-2 times a year	✓ Monthly, on 60 % of the mix
Reaction to input inflation	✗ Reactive (0 prior models)	✓ Predictive (3 scenarios: 5 %/12 %/20 %)
Demand segmentation	✗ None (average menu)	✓ 3 segments by price sensitivity
Recoverable margin points	✗ 0-1 pt (blunt adjustment)	✓ 2-4 pts (surgical adjustment)
Traffic-loss risk	✗ High on linear hikes	✓ Low (raises where elasticity < 1)
Shock anticipation window	✗ 0 weeks (reactive)	✓ 2-3 weeks of lead time
Data the operation requires	✗ Sales per dish (POS)	✓ Sales + actual cost + segment + price history

Chapter 1 — Why the classic matrix no longer protects margin in 2026

The popularity-margin matrix stops protecting margin because it is a static diagnosis born in 1982, back when inputs rose 2-3 % a year. Kasavana and Smith, at Michigan State University, sorted every dish as Star, Plowhorse, Puzzle or Dog along two fixed axes: how much it sells and how much it leaves. The model was brilliant for its time and remains a good starting point. The problem in 2026 is double-digit volatility in proteins and dairy: the FAO food price index has posted year-over-year jumps above 20 %, and the USDA reports persistent food-away-from-home inflation. A dish tagged Star in January can slide to Puzzle by March without anyone

touching the menu. I have seen it in dozens of kitchens: a cut's food cost jumped from 28 % to 41 % in one quarter while the matrix still flagged it green. A map that updates once or twice a year leaves money on the table every week.

Chapter 1 — Why the classic matrix no longer protects margin in 2026 — in practice

Masterrestaurant's second generation does not discard the sales mix or contribution margin; it keeps them as the starting point and adds three layers the static snapshot ignores by design: elasticity by segment, cost variance, and scenario stress. The core difference is that the classic matrix describes while second generation prescribes: it tells you how many cents to raise on each dish, not just which quadrant it sits in. A diagnosis that only paints four boxes leaves the pricing call to the owner's gut, and gut does not defend margin under double-digit inflation. The prescriptive model closes that gap with an operable rule: each dish gets an adjustment computed from its elasticity, its cost variance and the inflation scenario you run. Compare the two by decision variables —the matrix uses 2 (popularity and margin); second generation uses 5 (popularity, margin, elasticity, variance, scenario) — and by frequency —1-2 times a year versus monthly on 60 % of the mix.

Chapter 2 — From diagnosis to prescriptive model: how much to raise each price

In practice, a Plowhorse (high sales, low margin) that tolerates a 6-8 % bump recovers 3-4 points of operating margin with no measurable traffic loss. Diego F. Parra puts it this way in Masterrestaurant board meetings: the matrix tells you where it hurts, the model tells you how hard to squeeze. Running the model every 30 days instead of every 12 months turns margin into a variable you defend, not one you endure. That cadence change is half the value. Elasticity turns the sales mix into a lever because it measures how much each dish's demand falls when you raise its price, and that number varies brutally across dishes. It is the first of the three new layers and the one that unlocks the most margin without touching a single recipe. An anchor dish, heavily ordered and loaded with perceived value, usually shows low elasticity: raise it 10 % and you lose 2-3 % of units, leaving more net margin.

Chapter 3 — Price elasticity: not every dish tolerates the same adjustment

A commodity dish, comparable to the one down the street, shows high elasticity: the same 10 % costs you 12-15 % of units. The classic matrix treats both the same, and that is its most expensive mistake. At Masterrestaurant we measure elasticity by customer segment over 8-12 weeks of real sales, not by hunch, cross-referencing the POS price history with units sold. A three-band map orders the menu: inelastic (<1 , raise hard), unit (≈ 1 , fine-tune), elastic (>1 , freeze or cut). The practical result: you push hard where the customer does not react and freeze the price where they do, capturing 2-5 points of margin. That segmentation is what the static snapshot never saw, because the matrix only looks at the room average, not the sensitivity dish by dish. Portion costing with a standard recipe is the base of the model, because without controlled variance any price rests on false figures.

Chapter 4 — Variance between theoretical and real cost: the base without which everything collapses

This is the second layer, and though it sounds less glamorous than elasticity, it is where 80 % of lost margin lives. Variance is the gap between your theoretical cost —what the recipe says the dish costs— and your real cost —what the register records at close. In kitchens with no control I see gaps of 4-7 %: waste, portions by eye, petty theft, unlogged substitutions. If your theoretical cost says 30 % and the real is 36 %, any price adjustment

computed on the theoretical figure is born broken: you are optimizing over a number that does not exist. Remember the hard Masterrestaurant costing rule: 32 % food cost per dish is the ceiling, not the target, and payroll and rent are not charged to the dish—they go to the break-even point. The first step of second generation is not raising prices: it is closing variance below 2 % with a weighed standard recipe and weekly inventory counts.

Chapter 4 — Variance between theoretical and real cost: the base without which everything collapses — in practice

Only then does the elasticity model have true data to chew on. Diego F. Parra insists 80 % of lost margin is not in the menu, it is in the gap nobody measures, which is why the phase order matters: variance first, price second. Pricing psychology multiplies the effect of the adjustment because it changes what the customer perceives without changing a gram of the dish. Three concrete levers: anchoring, price endings and position on the menu. Anchoring places a premium dish up top so the rest reads as reasonable; a \$48 anchor makes the \$29 dish read as value. Endings pay off differently by category: in fine dining, a round price with no symbol (\$34, not \$33.99) lifts the perceived ticket. Position matters: the eye lands first on the top-right corner and on the first dish of each block, where your highest margin belongs, not your cheapest plate.

Chapter 5 — Pricing psychology and ticket by segment: multiplying the effect without touching the recipe

These tactics move 3-6 % of the average ticket at zero input cost. The average ticket, moreover, stops being a blind number when you break it down by segment, because a \$32 average can hide a \$58 segment and a \$19 segment that cancel each other out. Breaking it down by business diner at midday, weekend family, and delivery reveals where margin truly grows and where there is only volume with no profit. In a real Masterrestaurant case, the delivery segment represented 40 % of orders and only 12 % of operating margin: lots of noise, little cash. By reweighting the mix and adjusting prices by channel, that same location lifted operating margin 4 points in one quarter without raising total volume. The classic matrix ignores the menu as a selling interface and treats the customer as a single average; second generation treats it as the best-paid salesperson in the room and reads the customer by segment.

Chapter 6 — Scenario stress: defensible margin under uncertainty, not a chart for the wall

Inflation scenario stress is the third layer that makes margin defensible, because it simulates what happens to your mix when the input jumps before it actually jumps. Instead of reacting when the supplier raises protein 15 %, you run that shock today against your current menu and see which dishes flip from profitable to red. The stated goal of the second-generation model is a defensible operating margin under uncertainty, not a snapshot that ages on the wall. In practice, I define three scenarios—base, tension (+10 %) and crisis (+20 %)—and for each one the model prescribes which prices to move and which dishes to redesign or retire. This three-scenario map crosses with the three elasticity bands from Chapter 3: an inelastic dish in the crisis scenario is your first pricing lever; an elastic dish that turns red at +20 % is a candidate for portion redesign or retirement, not a hike.

Chapter 6 — Scenario stress: defensible margin under uncertainty, not a chart for the wall — in practice

A location that runs this exercise every quarter absorbs an input shock with 2-3 weeks of lead time, not with two months of margin bleeding out. Diego F. Parra calls it budgeting for the scare: volatility is not predicted, it is rehearsed. With the six chapters chained—diagnosis, prescription, elasticity, variance, psychology and stress—

the menu shifts from decorative chart to a living margin system. The group recovered 3.1 points of mix-weighted contribution margin in 90 days by applying the full model to its 10 main dishes. Starting point: classic matrix pinned in the office, average food cost of those 10 dishes at 31 %, unmeasured variance, and a \$32 average ticket treated as a single number.

Chapter 10 — Quantified mini-case: 4-location full-service group in 90 days

The sequence was the method's: weeks 1-2, standard recipe and actual costing, which revealed a 5.4 % variance (real 36.4 % against theoretical 31 % on four dishes); weeks 3-6, elasticity measurement by segment over 10 weeks of sales, which classified 6 dishes as inelastic and 2 as elastic; weeks 7-10, variance reconciled below 2 % and stress at 5 %/12 %/20 %; weeks 11-12, execution. Six inelastic prices were raised between 6 % and 9 % and two elastic ones cut by 4 %. Result at 90 days: food cost of the 10 dishes from 31 % to 28 %, mix-weighted contribution margin +3.1 points, average ticket +4.80 USD, with no measurable traffic drop and no customer complaints. The key, per the operations director, was not raising prices: it was knowing which to raise and which to protect. That discernment is exactly what the 1982 matrix cannot give. The second-generation model fails when there is no quality data, and it must be said plainly: it is not magic, it is data discipline.

Chapter 11 — Limitations, assumptions and failure conditions of the model

Assumption #1: a price history per dish exists or a controlled 2-week test can be run. In a freshly opened venue, with under 8-12 weeks of sales, elasticity is estimated with a larger margin of error and prices should move with more caution. Assumption #2: variance can be reconciled, which demands weekly inventory and a weighed standard recipe; without that discipline, the model optimizes over false figures and does more damage than the matrix. Assumption #3: volume per dish is enough for elasticity to be statistically legible —dishes selling 1-2 units a day give no reliable signal and are grouped by category. Known failure conditions: external demand shocks (a competitor opening, a seasonal event) contaminate the elasticity reading and force recalibration; and a perceived-quality change (shrinking the portion to lower food cost) alters future elasticity and can turn a previously inelastic dish elastic. The food cost ceiling remains 32 % per dish: no elasticity model authorizes breaching that ceiling by raising portion at the expense of margin.

Chapter 12 — Limitations, assumptions and failure conditions of the model — in practice

That is why Diego F. Parra insists on treating the model as a living system: it is monitored at 3, 6, and 12 months and recalibrated, not decreed once and forgotten. The classic matrix is a diagnosis; second generation is a prescriptive model that says how many cents to raise each price, dish by dish. Demand elasticity turns the sales mix into a lever: an inelastic dish tolerates +10 % losing only 2-3 % of units; an elastic one loses 12-15 % on the same move. Per-portion costing with a standard recipe is the base; without variance controlled below 2 %, any pricing model rests on false numbers (the typical uncontrolled gap is 4-7 %). Pricing psychology (anchoring, price endings, menu position) multiplies the effect of the adjustment 3-6 % of the ticket without touching the recipe or the food cost. The average ticket stops being a blind average —a \$32 can hide \$58 and \$19— and breaks down by segment to reveal where real margin grows.

Chapter 13 — The differences that define margin

Scenario stress rehearses the shock before it happens: 3 scenarios (base, +10 %, +20 %) leave the trigger price defined with 2-3 weeks of anticipation.

Classic matrix vs. elasticity model, criterion by criterion

PRICE-ADJUSTMENT PRECISION

A · CLASSIC POPULARITY MATRIX Blunt:
raises all or holds all

B · MASTERESTAURANT Surgical: raises
inelastic 6-9 %, protects magnets

Verdict: The matrix's linear hike punishes the inelastic dish and the traffic magnet equally: raise the magnet, you lose visits; skip the inelastic, you leave margin. Second generation separates the two. In the mini-case, raising 6 inelastic dishes and cutting 2 elastic ones moved those 10 dishes' food cost from 31 % to 28 % with no traffic drop. That is the margin the matrix never sees because it averages what it should separate.

INFLATION ANTICIPATION

A · CLASSIC POPULARITY MATRIX Reacts
once the input has risen (0 weeks)

B · MASTERESTAURANT Trigger price
defined per scenario (2-3 weeks)

Verdict: The classic matrix arrives late by design: it recalculates only 1-2 times a year, so a 15 % protein jump finds it with a six-month-old price. Modeling 5 %/12 %/20 % has the trigger price ready before the hit. A location that runs the stress test each quarter absorbs the shock with 2-3 weeks of lead time instead of two months of margin bleeding out. Volatility is not predicted, it is rehearsed.

COST DATA BASE

A · CLASSIC POPULARITY MATRIX

Assumed theoretical cost, no variance (4-7 % gap)

B · MASTERRESTAURANT Theoretical vs. actual reconciled monthly (<2 %)

Verdict: Without controlled variance the model rests on false numbers. If your theoretical says 30 % and the real is 36 %, any adjustment computed on the theoretical is born broken. Diego F. Parra insists 80 % of lost margin is not in the menu, it is in the gap nobody measures. Closing variance below 2 % with a weighed standard recipe and weekly counts is the prerequisite, not a luxury: it comes first, before touching a single price.

CUSTOMER READING

A · CLASSIC POPULARITY MATRIX Single average customer

B · MASTERRESTAURANT 3 segments by price sensitivity

Verdict: A \$32 average ticket can hide a \$58 segment and a \$19 segment that cancel out; treating them as one hides where margin grows and where there is only volume. Segmenting by business diner, family, and delivery reveals a channel can be 40 % of orders and only 12 % of margin. Reweighting the mix and adjusting by segment lifted operating margin 4 points in one quarter without raising total volume in the real case.

SIDE-BY-SIDE COMPARISON

Traditional approach (popularity matrix) THE COMMON MISTAKE

- ✗ Sorts dishes into 4 quadrants and stops there: no guidance on how much to raise or for whom.
- ✗ Assumes dish cost is fixed; ignores the variance between theoretical and actual cost (typical gap 4-7 %).
- ✗ Raises prices linearly when inflation bites and scares off the elastic customer.
- ✗ Recalculates once or twice a year: arrives late to double-digit input volatility.
- ✗ Treats every customer as an average, with no segmentation by price sensitivity.
- ✗ Runs no scenarios: reacts once the supplier has raised the input, with 0 weeks of anticipation.

Second generation (elasticity modeling) MASTERRESTAURANT

- ✓ Quantifies the elasticity coefficient per dish and segment before moving a single price.
- ✓ Reconciles theoretical vs. actual cost monthly and attacks variance where it hurts most, closing it below 2 %.
- ✓ Raises where elasticity is < 1 (inelastic) and protects traffic-magnet dishes.
- ✓ Models 3 inflation scenarios (5 %/12 %/20 %) and has the trigger price ready before the input rises.
- ✓ Segments demand into 3 groups and calibrates the menu for each without fragmenting operations.
- ✓ Gives 2-3 weeks of lead time before the input shock, not two months of margin bleeding out.

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	CLASSIC POPULARITY MATRIX	2ND-GEN ELASTICITY MODELING
Decision variables	✗ 2 (popularity + margin)	✓ 5 (popularity, margin, elasticity, variance, scenario)
Recalculation frequency	✗ 1-2 times a year	✓ Monthly, on 60 % of the mix

	CLASSIC POPULARITY MATRIX	2ND-GEN ELASTICITY MODELING
Reaction to input inflation	✗ Reactive (0 prior models)	✓ Predictive (3 scenarios: 5 %/12 %/20 %)
Demand segmentation	✗ None (average menu)	✓ 3 segments by price sensitivity
Recoverable margin points	✗ 0-1 pt (blunt adjustment)	✓ 2-4 pts (surgical adjustment)
Traffic-loss risk	✗ High on linear hikes	✓ Low (raises where elasticity < 1)
Shock anticipation window	✗ 0 weeks (reactive)	✓ 2-3 weeks of lead time
Data the operation requires	✗ Sales per dish (POS)	✓ Sales + actual cost + segment + price history

THE NUMBERS THAT MATTER

Numbers that anchor the model

60%

of the sales mix usually sits in 10-12 dishes

4pts

of operating margin recoverable with calibrated elasticity

32%

food cost per dish: the ceiling, not the target

3x

inflation scenarios modeled (5 %/12 %/20 %)

REAL CASE

“We had the classic matrix pinned in the office and still lost margin every quarter. Moving to elasticity by segment, we raised 6 inelastic prices between 6 % and 9 % and cut 2 elastic ones by 4 %: the average food cost of those 10 dishes dropped from 31 % to 28 %, mix-weighted contribution margin rose 3.1 points, and the ticket rose 4.80 USD with no customer complaints and no measurable traffic drop over 90 days.”

— Operations director, 4-location full-service group (Masterrestaurant mini-case, anonymized figures)

HOW TO APPLY IT IN YOUR RESTAURANT

How to migrate from matrix to model in 90 days

1

Week 1-2: standard recipe and actual costing

Standardize the recipe and per-portion costing for the 10-12 dishes that concentrate 60 % of the mix. Without a reliable theoretical cost, no elasticity model has a base. Document weights, waste, and per-portion cost using last week's real purchase prices. Goal of this phase: every dish with a computed theoretical food cost and measured waste, not eyeballed. The ceiling per dish is 32 %; dishes above it go to portion or spec redesign before touching price.

2

Week 3-6: measure elasticity per dish and segment

Log units sold against each historical price change and compute the elasticity coefficient ($\% \Delta$ demand / $\% \Delta$ price) over 8-12 weeks of real sales. Classify each dish as inelastic (<1), unit (≈ 1), or elastic (>1). Segment by time slot and customer type: business diner at midday, weekend family, delivery. Without at least one historical price change per dish, run a controlled 2-week test before modeling.

3

Week 7-10: reconcile variance and model 3 scenarios

Close the theoretical vs. actual cost variance below 2 % with weekly inventory counts. On that clean base, simulate the impact of 5 %, 12 %, and 20 % input inflation on each dish's contribution margin. Define the trigger price per scenario before the input rises. Identify which dishes flip from profitable to red in the crisis scenario (+20 %) and prepare them for redesign or retirement.

4

Week 11-12: execute, apply pricing psychology, and monitor KPIs

Apply adjustments starting with inelastic dishes and protect traffic magnets. Reorder the menu with anchoring, endings, and position to capture 3-6 % extra ticket at zero input cost. Monitor food cost, mix-weighted contribution margin, and average ticket by segment at 3, 6, and 12 months. If a dish's traffic drops more than its elasticity predicted, revert and recalibrate: the model is alive, not a decree.

FAQ

Frequently asked questions

Is the popularity matrix useless now?

It works as a starting point to read popularity and margin, but it is static. In 2026, with double-digit volatile inputs, you must add elasticity by segment, cost variance, and inflation scenarios to defend operating margin.

What is price elasticity by segment?

It is how much a dish's demand varies with a price change, measured by customer group. An inelastic dish (coefficient <1) tolerates a price hike without losing sales; an elastic one (>1) loses them. It is measured over 8-12 weeks of real sales.

How many dishes should I model first?

The 10-12 that concentrate roughly 60 % of the sales mix. That is where the needle-moving margin lives. Modeling the whole menu at once burns time with no proportional return and delays the first margin capture.

How much margin can I recover?

With well-calibrated elasticity, 2 to 4 points of operating margin in a full-service venue with a 30 % food cost, raising only inelastic dishes and protecting traffic magnets. The Masterrestaurant case cited recovered 3.1 points in 90 days.

What are the model's limitations?

It demands quality data: price history, reconciled actual cost, and enough volume per dish. In new venues or with under 8 weeks of data, elasticity is estimated with controlled tests, not by hunch, and the margin of error rises.

DATA & SOURCES

Sector data 2026 (official sources)

Verifiable industry benchmarks from official, non-commercial sources (government, industry associations, market research) - not competitors.

Metric	Benchmark 2026	Source
Food cost por concepto	QSR 25–30% · casual 30–34% · fine dining 34–40%	National Restaurant Association
Menús más cortos	las cadenas recortan ítems de carta para proteger margen y velocidad de servicio	FSR Magazine
Ticket online alto	34% de clientes gasta ≥\$50 por pedido	Statista
Índice de precios de alimentos	referencia oficial de food cost	USDA
Off-premise	~75% del tráfico	Circana