

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

By  **Diego F. Parra** · Updated 2026-07-09 · 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 classic four-quadrant matrix (star, plowhorse, puzzle, dog) diagnoses the past; it doesn't model how demand reacts when you move the price. Second-generation menu engineering keeps that matrix as a base layer and adds a second one: price elasticity of demand measured by guest segment and daypart. If your per-dish food cost already sits under the 32% ceiling but total contribution margin stays flat, the problem isn't ingredient cost — it's that you're setting one price for different demand curves. Modeling elasticity by segment recovers 2 to 4 margin points without touching the recipe.

 **White Paper** · Technical document · C-Suite & multilateral banking · 12 min read · 2026-07-09

INTELLECTUAL PROPERTY OF MASTERRESTAURANT® — EXCLUSIVE FOR SECTOR LEADERS

Menu engineering was born in 1982 with Kasavana and Smith as a 2x2 matrix crossing popularity (sales mix) against contribution margin. Forty-four years later it remains the tool most operators use — when they use any — even though 71% of guests decide their order based on menu design and placement (OneHubPOS, Menu

Engineering 2024). The matrix classifies; it doesn't predict. And in 2026, with input inflation still volatile, classifying isn't enough.

This whitepaper documents the leap to the second generation: keep the matrix as diagnosis and layer on an elasticity model that segments demand by guest type, occasion and channel. Diego F. Parra's thesis is blunt: most restaurants leave 2 to 4 margin points hidden because they charge one price to demand curves that look nothing alike. I've seen it in single units and in twenty-location chains.

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	CLASSIC MENU ENGINEERING (1ST GEN)	MENU ENGINEERING 2.0 (SEGMENT ELASTICITY)
Decision variable	✗ Popularity + contribution margin (2x2 matrix)	✓ Popularity + margin + segment price elasticity
Horizon	✗ Retrospective (closed-period mix)	✓ Predictive (expected demand response to price)
Price granularity	✗ One price per dish for everyone	✓ Price by segment/daypart/channel (same dish)
Cost starting point	✗ Per-portion food cost ≤ 32%	✓ Food cost + absolute contribution margin in \$
Margin lever	✗ Menu redesign and visual repositioning	✓ Elasticity modeling + redesign + description
Typical margin gain	✗ 1-2 pts on the mix	✓ 2-4 extra pts without touching the recipe
Data required	✗ Sales per dish (basic POS)	✓ Sales per dish × segment × hour (POS + CRM)

Chapter 1 — What is the classic menu-engineering matrix missing?

The Kasavana and Smith matrix classifies the past but never predicts how demand reacts when you move the price.

Born in 1982, it plots popularity against contribution margin across four quadrants —star, plowhorse, puzzle, dog — and still remains the only tool most operators use. The flaw is architectural: it tells you a dish sells poorly, not the price at which it would start selling for a specific segment. And menu layout matters more than most think: 71% of guests decide their order based on menu design and placement, according to OneHubPOS (Menu Engineering 2024). In 2026, with food-cost inflation still volatile, classifying yesterday does not protect tomorrow's margin. The photograph is well taken; the film is missing. Second generation keeps the matrix as its base

layer and overlays a price-elasticity model that segments demand by guest, occasion and channel. The first generation optimizes the photo of the menu; the second models the film of how demand shifts when you change the price.

Chapter 2 — What does second-generation menu engineering involve?

Instead of one contribution margin, you cross that margin with distinct sensitivity curves: the midday business diner barely notices a one-dollar increase, while the Saturday occasion diner clearly does.

The behavioral evidence backs the approach: 56% of guests chose dishes with descriptive labels in Cornell's study (Wansink, Food & Brand Lab). That same sensitivity to framing is what 2.0 models with data, not intuition. You do not replace the matrix: you turn it into the first layer of a living model. Most restaurants leave 2 to 4 margin points hidden because they charge one price to demand curves that look nothing alike. This is Diego F. Parra's central thesis at Masterrestaurant, and I have seen it identically in single-location operations and in twenty-unit chains. Willingness to pay varies by attribute with enormous gaps: 72% of guests would pay more for sustainability and 18% would accept a 6% to 10% premium, according to Toast (Restaurant Sustainability Survey 2025); 38% pay more for protein-rich dishes, according to Nation's Restaurant News (2025).

Chapter 3 — How much margin does a single price hide?

A flat price ignores both populations at once. The second generation does not raise everything: it identifies which dishes tolerate the increase for which segment and leaves the sensitive ones intact.

Those 2 to 4 points are the difference between surviving and growing. Engineering 2.0 requires crossing the POS with the CRM and the time slot, because the classic matrix works only with sales mix and that is not enough to segment the curve. The POS tells you what sold; the CRM and the hour tell you who bought it and on what occasion, which is where elasticity lives. Without that cross you cannot separate the midday diner from the weekend one, and 42% of younger guests share a main dish more frequently, according to Acosta Group (2025): a pattern that alters the ticket per person and that an aggregated POS hides. Adding the time slot reveals another layer: 37% of consumers seek quick bites instead of large meals, according to Circana (2024).

Chapter 4 — What data does the model need that the POS alone lacks?

Modeling without those axes is averaging three different restaurants inside the same location and setting one single price.

A dish the classic matrix flags as a 'dog' —low popularity, low margin— stops being one when you re-price or re-position it for the right segment, not when you delete it. The classic version tells you which dish is a dog; 2.0 tells you at what price that same dish stops being a dog for a specific diner. Framing moves the needle: 56% of guests chose options with descriptive labels at Cornell (Wansink), and more than half of consumers lean toward a dish labeled 'spicy' in 2025 versus 39% in 2015, according to Datassential (Spicy Food Trends 2025). A 'dog' with the right label and a price calibrated to its occasion can become the margin that rescues the menu. Deleting it destroys an option that was merely miscommunicated and mispriced, not dead.

Chapter 5 — Why does aggressive dynamic pricing destroy restaurant margin?

Airline-style dynamic pricing is the trap of a poorly executed second generation: it punishes the perception of fairness and drives the diner away.

The data is blunt: 52% of consumers see dynamic pricing in restaurants as 'price gouging' and 36% would order less frequently if it is applied, according to Capterra (2024 survey). A well-built second generation is not surge pricing: it is structural segmentation by menu, channel and occasion —lunch prices distinct from dinner, a delivery menu with its own architecture— that the diner perceives as coherent. Diego F. Parra sums it up this way at Masterrestaurant: you model elasticity to set the right price per segment, not to change it every hour. The same sensitivity that makes people pay more for value —44% are motivated by locally sourced ingredients, according to Toast 2025— resents anything that smells like manipulation. Elasticity yes; visible opportunism, never. Implementation is layered: first the classic matrix as diagnosis, then the data cross, then elasticity modeling and finally differentiated prices by segment.

Chapter 6 — How is it implemented in layers without collapsing operations?

Starting with surge pricing is the error that leads to the 52% who perceive it as abuse (Capterra 2024); starting with a clean diagnosis avoids that blow.

The data layer requires joining POS, CRM and time slot; the elasticity layer, reading willingness to pay by attribute —61% seek 'natural' items, according to Nation's Restaurant News (2024), and 1 in 3 would pay more for plant-forward dishes, according to Datassential (2024)—. Each layer is validated before the next, measuring real margin per segment against the model. This way the chef-owner does not risk the operation: they build a system that, layer by layer, recovers the hidden 2 to 4 points without touching the diner's sense of fairness or the coherence of the menu. The first generation optimizes the PHOTO of the menu; the second models the FILM of how demand reacts when you move price. The classic uses contribution margin as its axis; 2.0 crosses it with price elasticity, which differs for the weekday-lunch business guest and the Saturday-night occasion guest.

Chapter 7 — What actually changes between generations

The classic works off POS data; 2.0 needs to cross POS with CRM and daypart to segment the demand curve. The classic tells you which dish is a 'dog'; 2.0 tells you at what price that same dish stops being a dog for a specific segment.

POINT BY POINT

Criterion-by-criterion comparison

MIX DIAGNOSIS

A · CLASSIC MENU ENGINEERING (1ST GEN)

Classifies dishes into 4 quadrants

B · MASTERRESTAURANT

Classifies and also models price response

Verdict: 2.0 wins: the classic diagnosis is the starting point, not the destination.

PRICE SETTING

A · CLASSIC MENU ENGINEERING (1ST GEN)

One price for all guests

B · MASTERESTAURANT Price aligned to each segment's elasticity

Verdict: 2.0 recovers margin the classic hides in inelastic segments.

PERCEPTION RISK

A · CLASSIC MENU ENGINEERING (1ST GEN)

Low (single stable price)

B · MASTERESTAURANT Low if aggressive dynamic pricing is avoided

Verdict: Tie: stable segmentation doesn't trigger the 52% rejection of price gouging.

DATA DEMAND

A · CLASSIC MENU ENGINEERING (1ST GEN)

Basic POS

B · MASTERESTAURANT POS + CRM + daypart

Verdict: The classic is simpler; 2.0 demands data governance, and that's where its margin lives.

SIDE-BY-SIDE COMPARISON

Classic menu engineering 1ST GENERATION

- ✗ 2x2 matrix: star, plowhorse, puzzle, dog
- ✗ Retrospective diagnosis of sales mix
- ✗ One single price per dish
- ✗ Optimizes visual position and description
- ✗ Improvement ceiling: 1-2 margin points

Menu engineering 2.0 MASTERRESTAURANT

- ✓ Classic matrix + segment price-elasticity layer
- ✓ Models how demand responds to price
- ✓ Differentiated pricing by daypart, channel, occasion
- ✓ Segments guest, hour and absolute contribution in \$
- ✓ Improvement ceiling: 2-4 extra margin points

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	CLASSIC MENU ENGINEERING (1ST GEN)	MENU ENGINEERING 2.0 (SEGMENT ELASTICITY)
Decision variable	✗ Popularity + contribution margin (2x2 matrix)	✓ Popularity + margin + segment price elasticity
Horizon	✗ Retrospective (closed-period mix)	✓ Predictive (expected demand response to price)
Price granularity	✗ One price per dish for everyone	✓ Price by segment/daypart/channel (same dish)
Cost starting point	✗ Per-portion food cost $\leq$ 32%	✓ Food cost + absolute contribution margin in \$

	CLASSIC MENU ENGINEERING (1ST GEN)	MENU ENGINEERING 2.0 (SEGMENT ELASTICITY)
Margin lever	✗ Menu redesign and visual repositioning	✓ Elasticity modeling + redesign + description
Typical margin gain	✗ 1-2 pts on the mix	✓ 2-4 extra pts without touching the recipe
Data required	✗ Sales per dish (basic POS)	✓ Sales per dish × segment × hour (POS + CRM)

THE NUMBERS THAT MATTER

Numbers behind the model (industry sources, 2024-2026)

71%

of customers decide their order based on menu design and placement

56%

chose dishes with descriptive labels in the Cornell study

52%

sees restaurant dynamic pricing as 'price gouging'

36%

would order less often if dynamic pricing were used

75%

of customers prefer smaller portions for less money

72%

is willing to pay more at restaurants with sustainability

VISUALIZATION

The numbers, visualized

of customers decide their order based on menu design and placement



chose dishes with descriptive labels in the Cornell study



sees restaurant dynamic pricing as 'price gouging'



would order less often if dynamic pricing were used



of customers prefer smaller portions for less money



is willing to pay more at restaurants with sustainability



Sources: [OneHubPOS, Menu Engineering 2024](#) · [Cornell University Food & Brand Lab \(Wansink\)](#) · [Capterra, 2024 survey](#) · [National Restaurant Association, State of the Restaurant Industry 2024](#) · [Toast, Restaurant Sustainability Survey 2025](#)

Chart by [masterrestaurant.com](#)

REAL CASE

“71% of guests make their decision by looking at how the menu is designed and placed. That makes menu engineering one of the lowest-CapEx profitability levers there is: you don’t rebuild the kitchen, you reorganize information and price. But one single price over that information is leaving money on the table when your segments respond differently.”

— **Synthesis by Diego F. Parra (Masterrestaurant) on OneHubPOS, Menu Engineering 2024 data**

HOW TO APPLY IT IN YOUR RESTAURANT

90-day implementation roadmap

1

Days 1-30 · Standard recipe and theoretical cost

Lock the standard recipe and per-portion costing of every dish to know the exact theoretical cost. Without reliable theoretical cost there's no measurable food cost variance or real contribution margin. Keep per-dish food cost $\leq 32\%$; payroll, rent and utilities don't load onto the dish (they go to break-even). Close with last quarter's sales mix pulled from the POS.

2 Days 31-60 · Classic matrix + segmentation

Build the 2x2 matrix (popularity × contribution margin) and, in parallel, cross sales per dish with guest segment and daypart from POS and CRM. Here the segment demand curve appears: the same dish has different price elasticity at business lunch and at occasion dinner.

3 Days 61-80 · Elasticity and price modeling

Estimate segment price elasticity with controlled A/B price tests (not aggressive dynamic pricing: 52% see it as abuse per Capterra 2024). Raise price where demand is inelastic and protect volume where it's elastic. Reposition stars visually and add descriptions (measured sales lift by Cornell).

4 Days 81-90 · KPIs and margin governance

Install the tracking dashboard: food cost variance, absolute contribution margin in \$, average check by segment and prime cost. Set monthly mix review and quarterly elasticity review. Present to the board the ROI in recovered margin points with zero kitchen CapEx.

FAQ

Frequently asked questions

Does menu engineering 2.0 replace the classic Kasavana-Smith matrix?

It doesn't replace it: it uses it as a base layer. The matrix classifies each dish by popularity and contribution margin; the second generation adds segment price elasticity on top. Without the matrix you'd know how much each dish sells, but not at what price to maximize it per guest type.

Is segment pricing the same as dynamic pricing?

No. Dynamic pricing changes price in real time by instant demand, and 52% perceive it as abuse (Capterra 2024). Segment pricing sets stable, different prices by daypart, channel or occasion, aligned with each segment's willingness to pay. It's transparent and doesn't punish customer trust.

What data do I need to model elasticity by segment?

The viable minimum is sales per dish crossed with daypart and channel from the POS. The advanced level adds guest segment from the CRM. With that you run controlled A/B price tests to estimate real elasticity, without inventing theoretical curves.

How much margin can I really recover?

Diego F. Parra's reading of real operations puts the typical gain at 2 to 4 additional margin points over what the classic matrix already delivers, without touching the recipe or raising food cost. The lever is charging the right price to each segment, not cheapening the ingredient.

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
Reducción de calorías por el etiquetado en el menú	≈7,3% menos de calorías	US FDA / estudios de menu labeling
Menos calorías por transacción en una gran cadena de café (etiquetado)	-4,6% de calorías por transacción	American Journal of Preventive Medicine — estudio
Ahorro estimado al sistema de salud por el etiquetado de calorías (FDA)	≈USD 8 mil millones en 20 años	US Food and Drug Administration
Usuarios de fármacos GLP-1 que comen fuera con menos frecuencia (EE. UU.)	54% de los usuarios	Encuesta a 1.000 usuarios GLP-1 vía Fortune — 2025
Usuarios de GLP-1 que consumen menos snacks (EE. UU.)	≈70% de quienes reportan menos calorías	EY-Parthenon — encuesta 2025
Reducción del gasto de hogares con usuarios de GLP-1 (EE. UU.)	-10% en un año (100 categorías)	Numerator — 2025

Propiedad Intelectual de Masterrestaurant® — Exclusivo para Líderes de Sector · masterrestaurant.com