

Territorial Pre-Feasibility Models for Restaurant Chains: Cost-Stress Scenario Simulation Before Signing High-Value Leases

By  **Diego F. Parra** · Updated 2026-07-06 · Leadership & Team

MASTERRESTAURANT®

White Paper

Modelos de Prefactibilidad Territorial para Cadenas de Restaurantes: Simulación de Estrés de Costos antes de Firmar un Arrendamiento de Alto Valor

Método probado en +8.400 restaurantes · 43 países

meseros.ai

QUICK VERDICT

Verdict: signing a high-value lease on a single-scenario projection is the most expensive structural failure in restaurant expansion. The correct approach does not project an optimistic P&L: it subjects each location to *triaxial stress* —input inflation at 5%, 12% and 20%, staff turnover at 40% and 90%, and a 15% sales downside— and signs only where EBITDA stays positive in the adverse scenario. In our data across 8,400+ restaurants in 43 countries, 62% of branch closures within 24 months held a contract that would never have passed this stress test. Staff turnover, not rent, is the variable that breaks the model first. This 2026 white paper documents the full framework across six chapters, three quantified matrices and a real mini-case.

 **White Paper** · Technical document · C-Suite & multilateral banking · 22 min read · 2026-07-06

INTELLECTUAL PROPERTY OF MASTERRESTAURANT® — EXCLUSIVE FOR SECTOR LEADERS

A chain opening its twelfth branch faces a problem the second one never had: the high-value lease —\$18,000 to \$45,000 monthly in AAA locations— is signed for 5, 7 or 10 years. It is the least reversible CapEx commitment in the business. Yet most expansion directors evaluate it with a single cash-flow projection, almost always the optimistic one, built on labor and input-cost assumptions that 2024-2026 inflation already disproved. Diego F. Parra has audited dozens of these decisions at Masterrestaurant, and the pattern repeats with surgical precision.

The blind spot is not rent. It is staff turnover. A lease is fixed and known; a location's real labor cost depends on the local labor market, the zone's skills gap and how fast the team burns out. A branch in a location with 90% annual turnover pays three times the recruiting and training cost, runs understaffed and bleeds sales through poor service —all while the AAA rent keeps running. Per the U.S. Bureau of Labor Statistics, food-service annual turnover runs above 75%, far higher than the whole-economy average. Territorial pre-feasibility that ignores this variable projects a ghost.

This white paper presents the Masterrestaurant territorial pre-feasibility model: a quantitative framework Diego F. Parra developed after auditing dozens of failed expansions. Instead of a P&L, it produces a resilience matrix that subjects each candidate location to simultaneous stress scenarios of input cost, staff turnover and sales decline, and returns a binary verdict: sign or discard. The six chapters that follow develop the macroeconomic context, the failure of the traditional approach, the theoretical framework with its formulas, the solution architecture, the stress-scenario benchmark and a 90-day implementation roadmap with tracking KPIs. It closes with the honest limitations and assumptions of the analysis.

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	SINGLE-SCENARIO PROJECTION (TRADITIONAL APPROACH)	TRIAxIAL-STRESS PRE-FEASIBILITY (MASTERRESTAURANT)
Scenarios modeled per location	✗ 1 (optimistic)	✓ 9 (3 inflation × 3 turnover)
Labor variable in the model	✗ Fixed estimated cost (~22%)	✓ Turnover 40%/65%/90% dynamic
Decision threshold	✗ IRR > 18% at base	✓ EBITDA > 0 in adverse scenario
Input-inflation sensitivity	✗ Not modeled (0 levels)	✓ 5% / 12% / 20% simulated
24-month closures (MR data)	✗ 62% of failed sample	✓ 8% of failed sample
CapEx exposed per error	✗ \$420,000 avg/closure	✓ \$0 (discarded before signing)
Evaluation time per location	✗ 3-5 days	✓ 6-8 days

Chapter 1 — The macroeconomic context: why 2024-2026 broke the inherited expansion model

The 2024-2026 macroeconomic context made the expansion model chains used for a decade obsolete. Three indicators explain it. First, input inflation: USDA food-price indexes stacked double-digit increases no 2019 P&L anticipated, pushing food cost above the 32% ceiling in entire menu categories. Second, the labor market: the U.S. Bureau of Labor Statistics documents food-service annual turnover above 75%, far above the whole-economy average. Third, the cost of capital: financing a AAA site's CapEx is dearer than five years ago. Diego F. Parra puts it plainly: restaurant expansion stopped forgiving lazy assumptions, and territorial pre-feasibility that fails to absorb these three shocks projects a business that no longer exists. A high-value lease is the least reversible CapEx commitment across the entire chain. A AAA site runs \$18,000 to \$45,000 monthly and is signed for 5, 7 or 10 years; a ten-year contract at \$30,000 a month commits \$3,600,000 before selling the first plate.

Chapter 2 — The high-value lease as the least reversible CapEx in the business

Unlike kitchen equipment you resell or a campaign you cancel, the lease runs in full even if the branch closes. I've seen it across dozens of expansions: the expansion director signs on the optimistic projection, operations underperform, and the AAA rent keeps debiting on time while EBITDA sinks. That asymmetry —variable revenue against a fixed seven-figure obligation— is what demands modeling the worst case before signing, not after. At Masterrestaurant we treat every lease as irreversible, because financially it is. For the operator, the context's message is direct: any pre-feasibility template built before 2024 structurally underestimates labor and food cost and must be recalibrated against live sources. Before evaluating your next site, refresh three assumptions with verifiable external data —input inflation (USDA), sector turnover (U.S. Bureau of Labor Statistics) and cost of capital— and treat them as variables, not inherited constants. A model that isn't recalibrated yearly isn't conservative: it's blind.

Chapter 3 — Operator implications (Chapter 1)

Anchoring each assumption to a public index is also what makes your decision defensible to the board when it asks where each number came from. Signing a high-value lease on a single-scenario projection is the most expensive structural failure in restaurant expansion. The mistake I see again and again: an expansion director approves an \$18,000 to \$45,000 monthly site, locked in for 5, 7 or 10 years, on a single optimistic P&L. That CapEx never comes back. At Masterrestaurant we put every site through a triaxial stress test —input inflation, staff turnover and a sales drop, all at once— because the optimistic projection sells the project to the board, but the stress model prevents the closure. Signing ten years of AAA rent on a 2019 labor assumption is betting the whole chain's cash on the worst labor market never arriving. It arrives. And when it does, 62% of those contracts —in our sample— end in closure before 24 months.

Chapter 4 — The quantified cost of inaction: what the chain that skips the stress loses

The cost of not modeling stress is measurable and brutal: \$420,000 in irrecoverable CapEx on average per closing branch, plus the reputational cost and the toll on the leadership team. In Masterrestaurant data across 8,400+ restaurants, 62% of 24-month closures carried a contract that would not have passed triaxial stress; with the model applied, that share drops to 8%. In dollars: a chain opening five sites a year with the traditional approach can expect three fragile closures; with the stress model, under half of one. The difference —two saved locations — is over \$800,000 in protected CapEx a year, not counting the lost cash of two branches that should never

have opened. Inaction isn't neutral: it's the most expensive decision in the portfolio. The blind spot in territorial prefeasibility is not the rent: it is staff turnover. The lease is fixed and known; the real labor cost depends on the local job market, the area's skills gap and how fast the team burns out.

Chapter 5 — Staff turnover, not rent, is the blind spot that sinks the site

A branch in a market with 90% annual turnover pays up to three times the cost of recruiting and training, runs understaffed and bleeds sales through poor service—all while the AAA rent keeps running. I have seen it across dozens of expansions: site twelve faces a problem site two never had, because the founding team can no longer be cloned across twelve locations at once. A model projecting labor cost at 22% in a zone with real turnover of 85% is not projecting a restaurant: it is projecting a ghost. Turnover shifts up to 33% of prime cost over twelve months. For the operator, Chapter 2's lesson is that the single-scenario approach fails not from moderate optimism but from omitting the variable that moves most. Audit your last three openings: how many modeled turnover as a dynamic variable and how many as a fixed cost? If the answer is 'none', your expansion process is structurally exposed, and the next \$420,000 closure is already on its way.

Chapter 6 — Operator implications (Chapter 2)

The concrete action: ban your committee from approving any location whose P&L doesn't include at least one quantified adverse-turnover scenario. That single rule, with no added sophistication, reorders the priorities of the whole table. The theoretical framework rests on three input variables and one decision rule. The first formula is stressed prime cost: $\text{prime cost} = \text{food cost} + \text{labor cost}$, where food cost is stressed by input inflation (5%, 12%, 20%) and labor cost by turnover (40%, 65%, 90%). The second is site EBITDA: $\text{EBITDA} = \text{sales} - \text{prime cost} - \text{AAA rent} - \text{other fixed costs}$, computed in each of the nine cells. The central and most uncomfortable assumption is that food cost per dish must never exceed 32% even in the 20% inflation scenario; if it does, that site's menu engineering is unviable from the start. Diego F. Parra formalizes what many operators intuit but never write: a location's resilience is a measurable function, not a hunch.

Chapter 7 — How does turnover enter the model as a dynamic variable, not an assumption?

Staff turnover enters the Masterrestaurant model as a dynamic variable, not a fixed assumption nobody revisits. Every turnover point above 40% loads three costs:

recruiting, the learning curve and margin loss from poor service. In a site with a high skills gap, that drag can turn a projected 14% EBITDA into an operating loss—without the rent changing by a single dollar. The model doesn't use a brochure turnover number; it uses the zone's real figure, cross-referenced with local minimum wage, density of competitors chasing the same talent and employment seasonality. A 30-point turnover gap between two sites with the same rent can mean \$90,000 a year in extra labor cost. That number, not the optimistic cash flow, is what decides whether the ten-year lease gets signed. Every resilience matrix rests on explicit assumptions, and hiding them is the trap that sinks restaurant financial models.

Chapter 8 — The honest assumptions that hold up the resilience matrix

The Masterrestaurant model declares four. First: AAA rent is fixed and non-renegotiable over the lease term. Second: input inflation and turnover are independent of each other, though in a crisis they correlate. Third: average ticket holds stable in the 15% sales-downside scenario—the drop is traffic, not price. Fourth: a new employee's learning curve lasts 30 to 90 days by role. Each assumption is debatable, which is exactly why it's written

down: a model that doesn't expose its assumptions isn't rigorous, it's internal marketing dressed up as finance. Methodological honesty is part of the rigor, not an appendix. For the operator, the framework translates into a practical demand: write your formulas and assumptions before entering a single number. The concrete action is a nine-cell spreadsheet —three inflation levels by three turnover levels— where each cell computes EBITDA against fixed rent. Document the four assumptions beside it and who validated them.

Chapter 9 — Operator implications (Chapter 3)

When the board asks why you discarded a 'gorgeous' site, you don't answer with intuition: you show the worst-case cell in the red and the assumption that holds it. That traceability turns a defensive decision into corporate governance and shields the expansion director from hindsight review. The Masterrestaurant territorial pre-feasibility architecture replaces the P&L with a three-component resilience matrix. The input component calibrates zone variables: real turnover, skills gap, projected input inflation and contracted rent. The simulation component crosses the three inflation levels with the three turnover levels and computes EBITDA across the nine resulting cells. The decision component applies the binary threshold: sign only if EBITDA stays positive in the most adverse cell (90% turnover + 20% inflation). The system doesn't ask 'how much will I earn if everything goes well?'; it asks 'do I survive if turnover jumps to 90% and inputs rise 20%?'.

Chapter 4 — Solution architecture: the resilience matrix component by component

The first question looks good in committee; the second protects the whole chain's cash. Diego F. Parra designed this architecture after auditing dozens of failed expansions: each component closes a distinct mode of self-deception. The resilience matrix returns a binary verdict —sign or discard—, not a reassuring range of scenarios. Each site is scored on its ability to survive the worst simultaneous cross: inputs +20%, turnover at 90%, sales -12%. If the site stays in positive EBITDA under that cross, it gets signed; if it drops to an operating loss, it is discarded even if the optimistic scenario promises an 18% margin. The discipline is hard on purpose: I have watched chains close three of the five sites they opened in a year because each looked viable on its own P&L, but none held up under the triaxial stress. A 10-year AAA lease at \$30,000 a month is \$3,600,000 committed.

Chapter 11 — The resilience matrix: a binary verdict, not a reassuring range

No committee should approve that figure on a projection that cannot survive its own worst case. The binary verdict exists so nobody fools themselves with averages. A mini-case quantifies the model's value. A 14-unit restaurant group was about to sign ten years in a AAA mall at \$32,000 monthly —\$3,840,000 committed. The traditional P&L showed 16% EBITDA with labor cost assumed at 22% and everyone applauded. Running the site through triaxial stress with the zone's real turnover (85%), real labor cost jumped to 31% and EBITDA fell to -4%: an operating loss. The site was discarded after six days of analysis. Six months later, a rival chain signed that same unit and closed in fourteen months. The direct saving: \$460,000 of CapEx never committed, plus fourteen months of negative cash never suffered. Six days of matrix against nearly half a million dollars.

Chapter 12 — Quantified mini-case: the AAA site the matrix saved from a \$460,000 closure

That's the model's ROI, and it isn't theoretical: it's the math the expansion director actually ran. For the operator, the architecture teaches that the solution isn't costly software or a permanent consultant: it's a three-component discipline that fits in a well-built spreadsheet. The concrete action is to institutionalize the binary threshold in the committee minutes, so no site advances without its worst-case cell computed and signed off. Complement the matrix with the method's tools —the canvas to structure the site as a business unit and the cash-flow projection

to verify each scenario's liquidity— and you have a repeatable process. The chain that turns the hunch into a matrix stops losing \$420,000 per error and starts compounding margin site after site. The benchmark is built by simulating the three input-inflation scenarios —5% conservative, 12% base, 20% stress— crossed with the three turnover levels.

Chapter 5 — Benchmark and stress-scenario simulation (5%, 12% and 20% inflation)

In a typical site with \$1,200,000 annual sales, \$360,000 rent and a 30% base food cost, the conservative scenario (5% inflation, 40% turnover) yields a healthy EBITDA near 15%. The base scenario (12% inflation, 65% turnover) compresses it to 6-8%: still viable, but with no margin for error. The stress scenario (20% inflation, 90% turnover) drives food cost above 34% and labor cost above 30%, pushing EBITDA negative. The reading is unambiguous: a site that only works in the conservative scenario is a bet. The benchmark doesn't average the nine cells —the average hides tail risk— it reads the worst cell and decides on it. Input inflation and staff turnover don't add: they multiply on prime cost. A 30% base food cost rising to 34% on 20% inflation already presses margin; if turnover simultaneously pushes labor cost from 22% to 31%, combined prime cost jumps from 52% to over 65%.

Chapter 14 — How input inflation and turnover combine to break prime cost

With a healthy prime-cost ceiling near 60%, that site operates in structural deficit. This is the mechanism the single-scenario approach never sees, because it models each variable separately and at its optimistic level. USDA food-price indexes confirm input shocks aren't hypotheses: they're the recent norm. Modeling prime cost under combined stress —not under each isolated variable— is what separates an honest projection from an exercise in complacency. A chain that audits its twelfth site with the triaxial stress framework avoids the mistake of cloning site two's P&L into a labor market that looks nothing alike. The founding branch may have opened with 35% turnover and a hand-built team; site twelve enters a zone with 85% turnover and a severe skills gap, and there the same concept yields different cash. With the Masterrestaurant model, the expansion director walks into committee with an actionable number: this site survives the worst case at 6% EBITDA, this other one drops to -4% and is dropped.

Chapter 15 — What does a chain gain by auditing site number twelve with this framework?

Replacing the hunch with the resilience matrix has meant, across the expansions I audited with Diego F. Parra, going from 40% of problem sites a year to under 10%.

Expansion stops being a bet and becomes engineering. For the operator, the benchmark leaves three operating rules. First: if your site only clears the 5% conservative inflation scenario, don't sign it; the recent market doesn't deliver conservative scenarios. Second: watch combined prime cost, not food cost and labor cost separately — the trap lives in the interaction. Third: use the base scenario (12% inflation, 65% turnover) as your realistic planning line, not the optimistic one. The chain that plans on the base scenario and verifies survival on the adverse one builds a portfolio that withstands the cycle. The one that plans on the conservative one builds a portfolio that collapses with the first shock, and in 2026 the shocks arrive on time.

Chapter 6 — Implementation: 90-day roadmap, KPIs and ROI for the board

Implementation deploys over a 90-day roadmap. In days 1-30 the turnover variable is calibrated by zone for each pipeline site and the data is checked against external sources —sector turnover from the U.S. Bureau of Labor Statistics and input inflation from the USDA. In days 31-60 the nine-cell matrix is built per site and the four assumptions are documented with their validator. In days 61-90 the binary threshold is institutionalized in the committee minutes and approved sites are shielded with Open Badges micro-credentials and a shift-leadership PDA to cut real entry turnover. Diego F. Parra recommends running the first cycle on an already-signed site as calibration: if the model would have predicted its real performance, it's ready to decide the next ones. Tracking KPIs translate the matrix into operating governance across the first year. At 3 months the board should demand each new site's real entry turnover against the modeled level: an upside deviation is the first alarm.

Chapter 17 — The 3-, 6- and 12-month tracking KPIs the board should demand

At 6 months the central KPI is real combined prime cost against the base cell's projection, with food cost per dish watched under the 32% ceiling. At 12 months the cycle closes with real site EBITDA against the base and adverse cells: if the site lives where the matrix predicted, the model is calibrated; if not, recalibrate before the next opening. The ROI for the board is direct: each avoided \$420,000 closure pays the six days of analysis per site thousands of times over. Expansion stops being faith and becomes a control dashboard. This analysis has limits worth declaring honestly. First, the data across 8,400+ restaurants in 43 countries is Masterrestaurant proprietary, not a random sample audited by a third party; the closure percentages (62% vs 8%) reflect field experience, not a controlled study. Second, the model assumes input inflation and turnover are independent, when in real crises they tend to correlate and worsen jointly —the worst case could be worse.

Chapter 18 — Limitations and assumptions of the analysis

Third, the food-cost (32%) and prime-cost (60%) ceilings are MASTERRESTAURANT targets and vary by format and country. Fourth, the simulation assumes fixed rent; a sales-linked variable-rent clause would change the matrix. Fifth, the external sources (U.S. Bureau of Labor Statistics, USDA) are floor references, not exact proxies for each local zone. No model replaces judgment: the model orders the decision, it doesn't replace it. The traditional approach asks 'how much will I make if everything goes well?'; the stress model asks 'do I survive if staff turnover spikes to 90% and inputs rise 20%?'. The first question sells the project to the board; the second prevents the closure. Diego F. Parra insists: AAA lease CapEx is not recoverable, so the only honest projection is the one that assumes the worst possible labor market. In Masterrestaurant data, skipping this question explains 62% of 24-month closures; adding it drops them to 8%.

Chapter 19 — The differences that decide whether the branch survives

Staff turnover enters the Masterrestaurant model as a dynamic variable, not an assumption. Every turnover point above 40% loads recruiting cost, learning curve and service-driven waste. In a location with a high skills gap, that drag can turn a projected 14% EBITDA into an operating loss —without rent moving a single dollar. A 30-point turnover gap between two sites with identical rent can mean \$90,000 a year in extra labor cost. The traditional standard closes the deal and waits; the correct model closes the location in the model before signing it. Discarding a fragile location on the spreadsheet costs six days of analysis; discovering its fragility after the lease costs \$420,000 average in irrecoverable CapEx plus the reputational cost of the closure. A 10-year AAA lease at \$30,000 a month commits \$3,600,000. The resilience matrix turns a bet into an engineering decision.

Comparative analysis: traditional approach vs. triaxial stress

SCENARIOS MODELED

A · SINGLE-SCENARIO PROJECTION
(TRADITIONAL APPROACH)

A single optimistic P&L

B · MASTERRESTAURANT Nine triaxial-
stress scenarios

Verdict: Triaxial stress reveals the breaking point the single scenario hides: in MR data, the single-scenario projection lets 62% of the contracts that end up closing slip through, while the nine-cell matrix catches them before signing.

TURNOVER TREATMENT

A · SINGLE-SCENARIO PROJECTION
(TRADITIONAL APPROACH)

Fixed estimated labor cost (~22%)

B · MASTERRESTAURANT Dynamic variable
calibrated by zone

Verdict: Calibrated turnover captures the first cause of closure; fixed cost ignores it. A site with 85% real turnover modeled as a 22% fixed labor cost hides up to \$90,000 a year in overrun — enough to flip a 14% EBITDA into a loss.

DECISION RULE

A · SINGLE-SCENARIO PROJECTION
(TRADITIONAL APPROACH)

Base IRR above 18%

B · MASTERRESTAURANT Positive EBITDA
in adverse scenario

Verdict: The adverse threshold protects CapEx; base IRR rewards optimism. A \$3,600,000 ten-year contract cannot hang on a scenario that only works if nothing goes wrong; the binary threshold forces proof of survival in the worst cross.

CAPEX EXPOSURE

**A · SINGLE-SCENARIO PROJECTION
(TRADITIONAL APPROACH)**

\$420,000 average per closure

**B · MASTERRESTAURANT \$0: discarded
before signing**

Verdict: Discarding on the spreadsheet costs days; discovering it after the lease costs hundreds of thousands. Six days of analysis against \$420,000 irrecoverable is the best cost-benefit ratio in the entire expansion.

24-MONTH OUTCOME

**A · SINGLE-SCENARIO PROJECTION
(TRADITIONAL APPROACH)**

62% closures in the sample

**B · MASTERRESTAURANT 8% closures in
the sample**

Verdict: The stress model reduces branch closures almost eightfold. It's not that the market improves: fragile locations never get signed, so the sample that does open is pre-selected for resilience.

SIDE-BY-SIDE COMPARISON

The traditional approach: single-scenario projection FRAGILE

- ✗ A single P&L, almost always optimistic, on 2019 assumptions.
- ✗ Labor cost treated as constant (~22%), detached from the local labor market.
- ✗ Staff turnover absent from the model or as a footnote.
- ✗ Decision by base IRR with no adverse stress test.
- ✗ The lease is signed; the error surfaces 18 months later.

The correct model: triaxial cost-stress MASTERESTAURANT

- ✓ Nine scenarios per location: 3 inflation levels × 3 turnover levels.
- ✓ Staff turnover as an input variable calibrated by zone.
- ✓ Signing threshold: positive EBITDA in the most adverse scenario.
- ✓ Resilience matrix returning a binary sign/discard verdict.
- ✓ CapEx protected: fragile locations discarded before the lease.

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	SINGLE-SCENARIO PROJECTION (TRADITIONAL APPROACH)	TRIAxIAL-STRESS PRE-FEASIBILITY (MASTERESTAURANT)
Scenarios modeled per location	✗ 1 (optimistic)	✓ 9 (3 inflation × 3 turnover)
Labor variable in the model	✗ Fixed estimated cost (~22%)	✓ Turnover 40%/65%/90% dynamic
Decision threshold	✗ IRR > 18% at base	✓ EBITDA > 0 in adverse scenario
Input-inflation sensitivity	✗ Not modeled (0 levels)	✓ 5% / 12% / 20% simulated
24-month closures (MR data)	✗ 62% of failed sample	✓ 8% of failed sample
CapEx exposed per error	✗ \$420,000 avg/closure	✓ \$0 (discarded before signing)
Evaluation time per location	✗ 3-5 days	✓ 6-8 days

THE NUMBERS THAT MATTER

The numbers of territorial stress (Masterrestaurant data 2026)

62%

of 24-month closures held contracts failing the stress test

90%

annual turnover in high-skills-gap locations

420

K USD

average irrecoverable CapEx per branch closure

33%

of prime cost that turnover shifts over 12 months

9

scenarios simulated per location before signing

8%

of closures when the location passed triaxial stress

VISUALIZATION

The numbers, visualized

Kitchen turnover — 2026 industry benchmark



Industry net margin — 2026 industry benchmark



Optimal food cost — 2026 industry benchmark



Off-premise operation — 2026 industry benchmark



Labor cost — 2026 industry benchmark



Sources: [National Restaurant Association](#) · [Statista](#) · [Circana](#) · [U.S. Bureau of Labor Statistics](#)

Chart by [masterrestaurant.com](#)

REAL CASE

“We were about to sign ten years in a AAA mall. The P&L showed 16% EBITDA and everyone applauded. We ran the location through Diego's triaxial stress and at 85% turnover—the zone's real figure—the model showed a 4% operating loss. We discarded it. Six months later a rival chain opened there and closed in fourteen months. We saved \$460,000 of CapEx for six days of analysis.”

— Expansion Director, 14-unit restaurant group

HOW TO APPLY IT IN YOUR RESTAURANT

How to build your territorial pre-feasibility model

1. Calibrate the turnover variable by zone

Before projecting any revenue, measure the candidate location's labor market: sector annual turnover, availability of skilled staff and local skills gap. Set three input levels—40% (base), 65% (mid) and 90% (adverse)—with their associated recruiting, training and service-waste cost. This is the variable that breaks the model first; don't estimate it, measure it. Benchmark your zone figure against sector turnover from the U.S. Bureau of Labor Statistics so you don't fool yourself with a brochure number.

2. Simulate triaxial cost-stress

Cross the three turnover levels with three input-inflation levels (5%, 12%, 20%) to generate nine prime-cost scenarios per location. In each cell compute the resulting EBITDA against the fixed AAA rent. The matrix reveals the exact point where the location stops being viable. Don't average scenarios: the average hides the tail risk that causes closure. Anchor input inflation to a verifiable index—USDA food-price data serves as a floor—not to a purchasing hunch.

3. Apply the binary signing threshold

The Masterrestaurant rule is hard: sign only if EBITDA stays positive in the most adverse scenario (90% turnover + 20% inflation). If the location works only at base, it's a bet, not a decision. The binary threshold removes the temptation to rationalize a bad contract because 'the location is gorgeous'. The model decides, not the enthusiasm. Log the verdict and its assumptions in the committee minutes: that traceability protects the expansion director when the board reviews the call at 24 months.

4

4. Shield the approved location with micro-credentials

For locations that do pass, reduce real entry turnover by training management with Open Badges micro-credentials and a shift-leadership PDA. Cutting turnover from 65% to 45% improves the mid-scenario EBITDA by 3-5 points. Training is not expense: it's the lever that moves the location from the adverse scenario to the viable one. Measure the effect at 3, 6 and 12 months with the same matrix to verify the management investment translates into lower real labor cost.

FAQ

Frequently asked questions about territorial pre-feasibility

Why does staff turnover break the model before rent does?

Because rent is fixed and known at signing, while turnover is variable and multiplies in high-skills-gap locations. Every point above 40% loads recruiting, learning curve and service waste, and can turn a 14% EBITDA into a loss without rent changing at all.

How many scenarios should I simulate per location?

At least nine: three input-inflation levels (5%, 12%, 20%) crossed with three turnover levels (40%, 65%, 90%). Never average the scenarios; the average hides the tail risk that causes closure. Decide on the most adverse scenario, not on the mean.

What threshold do I use to decide whether to sign the lease?

The Masterrestaurant threshold is binary: sign only if EBITDA stays positive in the adverse scenario (90% turnover plus 20% inflation). If the location works only in the optimistic base scenario, it's a bet, not a defensible expansion decision before the board.

Does management training change the model's result?

Yes, measurably. Cutting real turnover from 65% to 45% with Open Badges micro-credentials and a shift-leadership PDA improves the mid-scenario EBITDA by 3-5 points. Training is the lever that moves a location from the adverse scenario to the viable one.

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
Rotación de sala (FOH)	>70% anual	U.S. Bureau of Labor Statistics
Costo por cada salida	\$1,500–3,000 por empleado	Nation's Restaurant News
Tendencias laborales del sector	presión salarial al alza desde 2020	McKinsey (insights)
Cultura y retención	cultura y desarrollo interno figuran como palanca #1 de retención en pymes	Inc.
Rotación de cocina	~50% anual	National Restaurant Association

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