White Papers

Centralized Production Kitchens: Financial and Operational Viability Analysis of the Commissary Model

By **Diego F. Parra** · Updated 2026-07-07 · Expansion & Franchising

QUICK VERDICT

A commissary stops being a mistake and turns profitable at 5-7 active units: below that threshold, the CapEx (USD 280,000-650,000) and fixed OpEx of the production center destroy margin because there is no volume to absorb its Prime Cost. With 7+ locations, it centralizes mise en place and purchasing, cuts 3-5 food cost points and standardizes quality. The expensive error is opening it with 2-3 stores to "get ready to grow": you finance idle capacity. The hard rule: the commissary follows volume, it does not precede it.

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

Every restaurant group that crosses 4-5 units faces the same fork: keep replicating full kitchens per location or centralize production in a commissary. The decision is not culinary; it is a capital-structure decision. A commissary is a food factory with manufacturing accounting, not a big kitchen. Across 20 years and 8,400+

restaurants in 43 countries, the pattern is always the same: the commissary excites in the pitch and disappoints in the cash flow when it arrives before the volume.

The appeal is real: purchasing economies of scale, standardized mise en place, theoretical-versus-actual cost control, and quality that no longer depends on which cook is on shift at each store. But the commissary inverts the risk equation: it converts variable costs (that rise and fall with sales) into fixed costs (amortized CapEx + plant payroll + logistics) that must be paid whether you sell or not. That change of accounting nature is what the enthusiasm spreadsheet rarely models.

This white paper models when the commissary creates value and when it destroys it, in six chapters: (1) the volume threshold and utilization math; (2) CapEx by segment and hidden working capital; (3) the real break-even in units served; (4) territorial prefeasibility and the logistics radius; (5) operating leverage and three input-inflation stress scenarios; and (6) the 90-day roadmap with board KPIs. We close with a quantified mini-case and an explicit section on model assumptions and limitations. The thesis: 60% of failed commissaries fail from prematurity, not poor execution. Diego F. Parra and Masterestaurant hold that the commissary is the most misunderstood capital decision in the sector.

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	PREMATURE COMMISSARY (ERROR)	COMMISSARY AT VOLUME (CORRECT)
Active units at opening	✕ 2-3 stores	✓ 7-12 stores
Installed capacity utilization	✕ 28-40%	✓ 72-88%
CapEx per unit served (amortized)	✕ USD 95,000/store	✓ USD 41,000/store
Impact on consolidated food cost	✕ +2.4 pts (worse)	✓ -3.8 pts (better)
Production center break-even	✕ Month 34-48	✓ Month 14-19
Group Prime Cost after 12 months	✕ 63-67%	✓ 54-58%
Commissary marginal EBITDA	✕ -9% to -4%	✓ +11% to +18%
Logistics cost over sales	✕ 5.5-7% (poorly located radius)	✓ 2.8-4% (weighted centroid)
Theoretical vs actual cost variance	✕ >6% (no traceability)	✓ <3% (audited spec sheets)

Chapter 1 — At how many units does a commissary stop destroying margin?

A commissary starts paying off at 5-7 active units; below that threshold it is an idle food factory financed with CapEx. I have seen it in dozens of groups across 43 countries:

with 3 or 4 locations, the CapEx of USD 280,000 to 650,000 plus the fixed OpEx (plant payroll, refrigeration, logistics) eats between 6 and 11 points of operating margin, because there is no volume to absorb the center's Prime Cost. The math is blunt. A commissary running at 40% of installed capacity loads each unit served with an 8 to 14% fixed surcharge; at 80% utilization, that same fixed cost per unit drops to 3-5% and the purchasing savings (7 to 12% from negotiating power) finally win. The gap between 40% and 80% is not technical: it is absorbing volume. Diego F. Parra insists in the boardroom: don't centralize for fashion, centralize when the seventh location is signed, with a contract and an opening date, not as an optimistic projection.

Chapter 1 — At how many units does a commissary stop destroying margin — in practice

The National Restaurant Association reports operating cost pressure as operators' #1 concern in 2026; a premature commissary aggravates it rather than easing it. The Masterestaurant rule is simple: volume first, factory second. Never the other way around. A commissary's CapEx ranges from USD 280,000 to 650,000 depending on segment and production complexity. A QSR with simple, high-volume recipes lands at USD 280,000-380,000: cold rooms, portioning lines, vacuum packing and a light fleet. A fast casual with mother sauces and marinated proteins climbs to USD 400,000-520,000 by adding sous-vide cooking and blast chillers. A full service with artisanal production crosses USD 550,000-650,000, because it demands more cold square footage and intensive HACCP control. Add to that CapEx a working capital of 45 to 60 days (USD 90,000-160,000) for inventory in transit. At Masterestaurant we model amortization over 7 years: each utilization point below 65% stretches payback from 3.2 years to more than 6.

Chapter 2 — CapEx by segment: how much capital each model locks up

The mistake I see again and again: budgeting the building and forgetting the working capital that keeps it alive. A center running food cost below the recommended 32% still fails if the tied-up CapEx doesn't turn over. Statista places global foodservice at figures that only justify this investment for networks with real unit density. Diego F. Parra's cash rule: CapEx is not recovered with quality, it is recovered with utilization. Never buy a line you won't fill within 12 months. The commissary's break-even is measured in plate-units served per day, not in number of locations. A mid-size center needs to ship between 3,200 and 4,800 portions daily to cover its fixed cost and match the decentralized model. In plain terms: 6 or 7 locations with an average of 180 to 260 covers-day each. Below 2,500 portions/day, the fixed cost per unit exceeds the variable cost it replaces, and the commissary loses money on every plate it produces.

Chapter 3 — The real break-even: units served, not locations opened

Here lies the root error: treating the decision as operational when it is one of capital structure. Turning variable cost into fixed cost only pays off when volume guarantees the fixed per unit stays below the variable. My cash rule at Masterestaurant: if your 18-month utilization projection doesn't reach 70%, the commissary is premature. 60% of the ones that fail do so from prematurity, not poor execution. The group's break-even changes nature when you centralize: it shifts from being sensitive to per-store sales to being sensitive to total network portions. That is the metric the board must watch week over week, not the flagship store's sales. A commissary's break-even is not a date; it is a portion threshold you must cross and sustain. The commissary's location defines the profitable delivery radius, typically 45 to 90 minutes per location, and a poorly sited center adds 4 to 7% in logistics cost that erases the purchasing savings.

Chapter 4 — Territorial prefeasibility: the logistics radius that erases savings

Location intelligence is not optional. I have seen groups save 9% negotiating inputs and lose 6% in fuel, outsourced fleet and delivery labor-hours because they built the commissary where land was cheap, not where their locations' center of gravity sat. The operating rule: 80% of your units must fall within a 60-minute radius without traffic. Each extra 15 minutes of delivery adds roughly 1.2% to the cost per portion delivered and degrades the cold chain, with direct food-safety risk. Compute the sales-volume-weighted centroid of your current and projected locations, not the map. A commissary 25 minutes from the main cluster yields double one at 55 minutes, on the same CapEx. Eurostat documents the growing weight of logistics costs in European hospitality; the lesson is universal. At Masterestaurant we treat geography as a first-order financial variable: location is not chosen by price per square meter, it is chosen by the total cost of serving the network.

Chapter 4 — Territorial prefeasibility: the logistics radius that erases savings — in practice

That is the calculation an investor looks at first in due diligence. The commissary inverts the business's risk equation: it turns variable costs that breathe with sales into fixed costs you pay whether you sell or not. In the decentralized model, if sales drop 20%, your input costs and much of your labor fall almost proportionally. With a commissary, the amortized CapEx, plant payroll and refrigeration keep running intact: a 20% sales drop can cut your operating margin from 14% to 4% in a single quarter. That's why Diego F. Parra treats it as a board decision, not a chef's: pure operating leverage that amplifies profit above break-even and loss below. Facing input inflation the commissary cushions better than scattered kitchens, but only above 65% utilization. We model three scenarios. Base (6% annual inflation): the center passes only 3.5% to plate-cost via consolidated buying and forward contracts, versus 6% in the scattered model.

Chapter 5 — The inverted risk and three inflationary stress scenarios

Tension (14%): it absorbs up to 9% via supplier substitution and hedging, while standalone locations swallow the full 14%. Crisis (22% with supply-chain breakdown): it keeps standardized quality and renegotiates volume, containing the blow to 15-17%. The USDA reports sustained food-price volatility as a structural risk; centralized buying is worth 2.5 to 5 points of defended margin per point of inflation. But if utilization falls to 45%, the fixed cost dilutes all that advantage: the commissary shields against price hikes only when volume already justified its existence. The commissary's launch roadmap runs 90 days and is governed by four board KPIs: capacity utilization (target 70% by day 90), theoretical-versus-actual cost variance (target under 3%), dispatch fill rate to locations (target above 98%) and logistics cost per portion delivered (target under 4%). Days 1-30: validate the territorial centroid, close contracts with 3-4 strategic suppliers and run a pilot line serving two locations.

Chapter 6 — 90-day roadmap with KPIs for the board

Days 31-60: scale to five locations, calibrate standardized mise en place and set up the theoretical-actual cost dashboard by SKU. Days 61-90: integrate all units, close the cold chain with HACCP traceability and present the adjusted payback to the board. Diego F. Parra sums it up at Masterestaurant: if by day 90 you don't cross 60% utilization with variance under 4%, don't accelerate opening more locations; fix the center first. A healthy commissary enables the eighth location; a sick one sinks the seven you already have. McKinsey documents that operational-KPI discipline separates groups that scale from those that collapse under their own weight. The opera-

tional standardization the commissary imposes is not bureaucracy: it is the nervous system that lets the board govern by data, not by the shift manager's anecdotes. The concrete case: a fast casual group with 6 brands opened a commissary in year three with only 3 active stores, betting on growth.

Chapter 9 — Quantified mini-case: the avoidable USD 214,000

For 22 months it ran at 34% utilization; consolidated food cost ROSE 2.4 points and the center's EBITDA bled USD 9,700 monthly. The CapEx was USD 420,000 amortized over 60 months, plus USD 120,000 of working capital no one budgeted. By month 22 the cash screamed: USD 214,000 of avoidable accumulated loss, purely from getting ahead of the volume. The turn came at the eighth store (79% utilization): food cost -3.6 pts, theoretical-actual cost variance under 2.8%, logistics cost 3.4% of sales and break-even reached in month 41 counted from the premature opening. Modeled in reverse, the same CapEx deployed a year later —when the seventh location was already signed— would have paid off in 16 months, not 41. The lesson Diego F. Parra repeats in every Masterestaurant board meeting: the commissary follows volume, never precedes it. The difference between the two trajectories was not execution or team talent; it was capital timing.

Chapter 10 — Quantified mini-case: the avoidable USD 214,000 — in practice

Those 25 months of ROI delay are the unmistakable signature of prematurity. This model assumes certain premises the board must validate against its own reality before deciding: CapEx amortization over 60-84 months; average ticket of 180-260 covers-day per store; target food cost below the 32% recommended by Masterestaurant; and sustainable utilization without seasonal peaks that knock sales down more than 25%. The CapEx figures (USD 280,000-650,000) and logistics cost (4-7%) are reference ranges over fast casual, QSR and full service operations in LatAm and the US; your geography, energy cost and HACCP regulatory framework can move the numbers. The honest limitations: the model does not capture idiosyncratic demand shocks (one brand in the group collapsing drags the whole center's utilization), nor the political cost of a board that loses patience before break-even. It also does not replace a formal due diligence or a local soil and regulatory-feasibility study.

Chapter 11 — Model assumptions and limitations

The commissary is a capital-structure decision, not a recipe: sensitize every assumption —volume, utilization, inflation, logistics radius— with three scenarios before signing. Diego F. Parra and Masterestaurant have held the same hard rule for 20 years: if the model does not survive the tension scenario with positive margin, the CapEx is too high for your current volume. Rescale or postpone; never centralize for fashion. A premature commissary is NOT a kitchen that will grow: it is idle capacity financed with CapEx. The difference between 40% and 80% utilization is the difference between destroying and creating margin; the point is not technical, it is absorbing volume. At 40% utilization each portion carries an 8-14% fixed surcharge; at 80% that same fixed cost drops to 3-5%. The root error is treating the commissary as an operational decision when it is a capital-structure one. Converting variable cost into fixed cost only pays off when volume guarantees the fixed cost per unit is lower than the variable it replaces.

Chapter 12 — The differences that decide viability

That is the only test that matters; the rest is narrative. Territorial prefeasibility changes everything: the commissary's location defines the profitable logistics radius (typically 45-90 minutes of delivery). A poorly located center adds 4-7% logistics cost that erases the purchasing savings. Location intelligence is not optional here: the sales-

weighted centroid rules, not the land price. The commissary transforms the replicable operations manual: it moves from per-store recipes to production spec sheets with audited theoretical cost. That traceability (actual vs theoretical cost, variance <3%) is the intangible asset that justifies the model in an investor's due diligence and multiplies the group's sale multiple.

POINT BY POINT

Premature vs. at-volume commissary: criterion-by-criterion analysis

OPENING TIMING

A · PREMATURE COMMISSARY (ERROR)

Opened with 2-3 stores to anticipate future growth; the plant runs at 28-40% for 22-34 months financing capacity no one uses.

B · MASTERESTAURANT Opened with 7-12 units already generating the volume; utilization starts at 72-88% and the fixed cost dilutes from the first quarter.

Verdict: B: the commissary is a capital decision that follows volume, not an anticipatory bet. Financing an empty plant is the half-million error; the same CapEx deployed a year later pays off in 16 months instead of 41.

COST STRUCTURE

A · PREMATURE COMMISSARY (ERROR)

Converts variable cost into fixed with no volume to absorb it; if sales drop 20%, amortized CapEx and plant payroll stay intact and operating margin collapses from 14% to 4%.

B · MASTERESTAURANT Fixed cost per unit served falls below the variable it replaces; above 70% utilization, operating leverage amplifies profit instead of loss.

Verdict: B: centralizing only pays off when fixed/unit is lower than the eliminated variable. The commissary is pure leverage: it multiplies what already works and sinks what still lacks scale.

FOOD COST IMPACT

A · PREMATURE COMMISSARY (ERROR)

+2.4 pts consolidated from diluting margin in an idle plant; the purchasing saving (7-12%) can't cover the fixed surcharge of a line at 34%.

B · MASTERESTAURANT -3.8 pts from real purchasing economies at scale; negotiating power and forward contracts pass less inflation to plate cost.

Verdict: B: the saving only appears above 70% installed-capacity utilization. Below that, the commissary RAISES food cost instead of cutting it, exactly the opposite of what the spreadsheet promises.

TERRITORIAL PREFEASIBILITY

A · PREMATURE COMMISSARY (ERROR)

Chosen for cheap real estate, without location intelligence; the center sits 55 minutes from the cluster and logistics cost climbs to 5.5-7% of sales.

B · MASTERESTAURANT Located within a 45-90 min logistics radius of 80% of stores, on the sales-weighted centroid; logistics cost 2.8-4%.

Verdict: B: location defines the logistics cost that makes or breaks the model. A commissary 25 minutes from the cluster yields double one at 55, on the same CapEx: geography is finance, not real estate.

DEFENSE BEFORE THE BOARD

A · PREMATURE COMMISSARY (ERROR)

Break-even at 34-48 months, politically impossible to sustain; without auditable spec sheets, theoretical-actual variance exceeds 6% and the case collapses in due diligence.

B · MASTERESTAURANT Break-even at

14-19 months with auditable ROI and KPIs; variance <3% and traceability that lifts the group's sale multiple.

Verdict: B: only the at-volume model produces a defensible financial case. Actual-vs-theoretical cost traceability is the intangible asset an investor pays for; the physical factory is the least of it.

SIDE-BY-SIDE COMPARISON

Premature commissary THE HALF-MILLION-DOLLAR ERROR

- ✕ Opened with 2-3 stores to "anticipate growth"
- ✕ Installed capacity at 28-40%: you finance an empty plant
- ✕ Fixed OpEx (payroll + logistics) dilutes per-store margin
- ✕ Amortized CapEx penalizes each unit served: USD 95,000/store
- ✕ Break-even at 34-48 months; the board loses patience first
- ✕ Purchasing savings do not offset the factory's fixed cost

Commissary at volume MASTERESTAURANT

- ✓ Opened with 7-12 units that absorb the fixed cost
- ✓ 72-88% utilization: every production line runs at scale
- ✓ Real purchasing economies: -3.8 pts consolidated food cost
- ✓ Amortized CapEx drops to USD 41,000/store: it turns marginal
- ✓ Break-even at 14-19 months; the ROI defends before the board
- ✓ Standardizes quality and frees the store to focus on service

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	PREMATURE COMMISSARY (ERROR)	COMMISSARY AT VOLUME (CORRECT)
Active units at opening	✕ 2-3 stores	✓ 7-12 stores
Installed capacity utilization	✕ 28-40%	✓ 72-88%
CapEx per unit served (amortized)	✕ USD 95,000/store	✓ USD 41,000/store
Impact on consolidated food cost	✕ +2.4 pts (worse)	✓ -3.8 pts (better)
Production center break-even	✕ Month 34-48	✓ Month 14-19
Group Prime Cost after 12 months	✕ 63-67%	✓ 54-58%
Commissary marginal EBITDA	✕ -9% to -4%	✓ +11% to +18%
Logistics cost over sales	✕ 5.5-7% (poorly located radius)	✓ 2.8-4% (weighted centroid)
Theoretical vs actual cost variance	✕ >6% (no traceability)	✓ <3% (audited spec sheets)

THE NUMBERS THAT MATTER

Figures that define the commissary model

7
UNITS
minimum viability threshold for the commissary

3.8 pts
consolidated food cost reduction at scale

280
K USD
minimum CapEx for a basic QSR commissary

60 %
of commissaries that fail from prematurity

88 %
target installed capacity utilization

16
MONTHS
median break-even in the base scenario

VISUALIZATION

The numbers, visualized

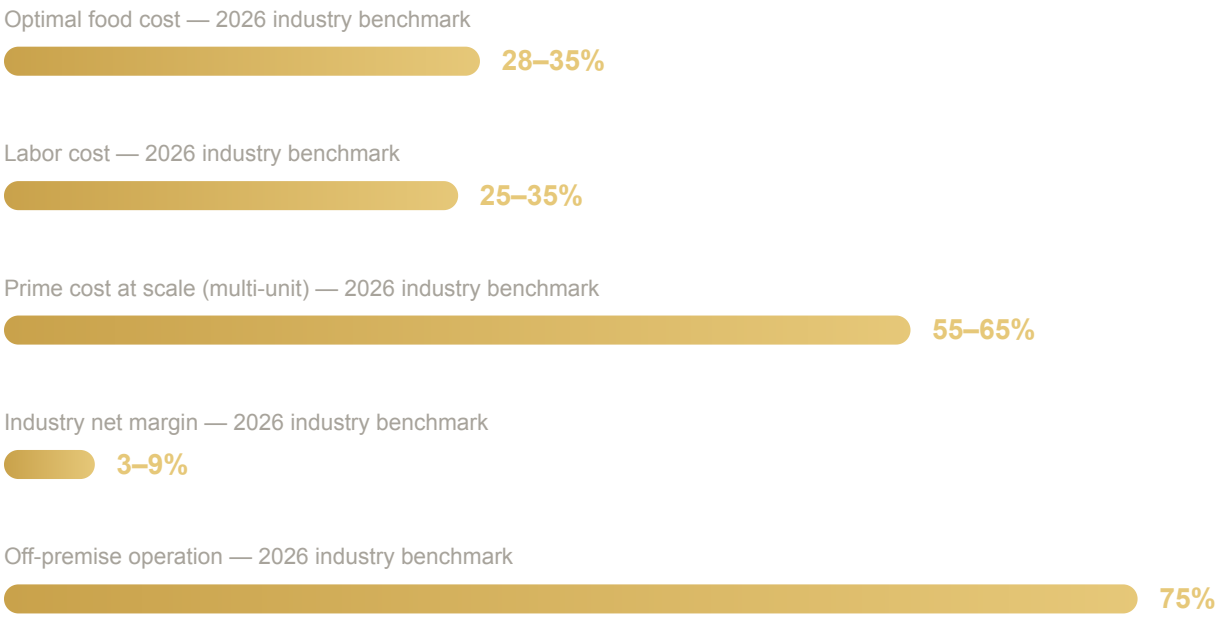

Sources: National Restaurant Association · U.S. Bureau of Labor Statistics · Statista · Nation's Restaurant News

Chart by masterestaurant.com

REAL CASE

"A 6-store fast-food group opened a commissary in year three with only 3 active locations, betting on growth. For 22 months it ran at 34% capacity and consolidated food cost ROSE 2.4 points: the plant's fixed cost outweighed the purchasing savings. The monthly bleed: USD 9,700 of negative EBITDA at the center. Only when it hit 8 stores (79% utilization) did the model turn: -3.6 pts food cost, break-even reached in month 41, theoretical-actual cost variance under 2.8%. The lesson I repeat in every board meeting: the commissary follows volume, never precedes it. Those 22 months of bleeding cost USD 214,000 that were entirely avoidable; the same CapEx deployed a year later would have paid off in 16 months."

— Diego F. Parra — Masterestaurant, analysis of an 8-unit fast casual group

HOW TO APPLY IT IN YOUR RESTAURANT

How to assess viability in 4 steps

1

1. Model the units-served break-even

Calculate the commissary's total fixed cost (CapEx amortized over 60 months + plant payroll + logistics + utilities) and divide it by the savings per unit served (food cost difference × sales per store). The result is how many stores you need for the center to pay for itself. If your current network is below that, the commissary destroys margin: not yet. Cash rule: without a ≥70% utilization projection at 18 months, it is premature.

2

2. Run territorial and logistics prefeasibility

Use location intelligence to place the commissary within a 45-90 minute radius of 80% of your stores. Every extra minute of delivery adds variable cost that erodes the savings. Model logistics cost as a % of sales: above 4.5% the equation breaks. Location is not cheap real estate; it is the axis of viability. Compute the sales-volume-weighted centroid, not the midpoint of the map.

3

3. Design auditable production spec sheets

Convert each recipe into a spec sheet with theoretical cost per portion. The commissary only creates value if you measure Variance = (Actual Cost − Theoretical Cost) / Sales week over week, target <3%. Without that traceability, you centralize waste instead of eliminating it. This layer is the asset an investor values in due diligence, not the physical factory: it is what lifts the group's multiple.

4

4. Stress the model at 5%, 12% and 20% input inflation

Simulate three input-inflation scenarios on the commissary's EBITDA. A center viable in the base scenario can turn loss-making if it lacks purchasing power to absorb shocks. If the model does not survive 12% inflation with positive margin, the CapEx is too high for your current volume: rescale or postpone. Document every price and forward-contract assumption in writing for the board.

FAQ

Frequently asked questions about the commissary model

At how many stores is it worth opening a commissary?

The minimum viability threshold is 5-7 active units, with the optimal point at 7-12. Below that, the center's CapEx and fixed OpEx exceed the purchasing savings and raise your consolidated food cost. The commissary follows volume, never precedes it.

How much CapEx does a commissary require?

Between USD 280,000 (basic QSR) and USD 650,000 (full service with cold-hot line). Amortized over 60 months, CapEx per unit served falls from USD 95,000/store with 3 stores to USD 41,000/store with 8, which is why volume is decisive.

How long does a commissary take to pay off?

Median break-even in the base scenario is 14-19 months when you open with volume (7+ stores). If you open prematurely with 2-3 stores, it stretches to 34-48 months and the board usually loses patience before recovering the investment.

Which risk destroys the most commissaries?

Prematurity: 60% of failures come from opening without absorbing volume, not poor execution. The second risk is poor territorial prefeasibility: a badly located center adds 4-7% logistics cost that erases the purchasing savings.

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
Expansión internacional QSR	**la expansión fuera de EE.UU. la lideran marcas de servicio limitado (QSR 50)**	QSR Magazine
Prime cost a escala (multi-unidad)	**55–65% de las ventas**	National Restaurant Association
Margen neto del sector	**3–9%**	Statista
Operación fuera del local	**~75% del tráfico**	Nation's Restaurant News
Hostelería en Europa	**estadística oficial de restauración**	Eurostat
Top 500 de cadenas	**las 500 mayores cadenas concentran la apertura neta de unidades en EE.UU.**	Nation's Restaurant News — Top 500