White Papers

The Repeat-Customer Economy: Lifetime Value Modeling and Frequency Architecture for Independent Restaurants

By **Diego F. Parra** · Updated 2026-07-07 · Marketing & Growth

QUICK VERDICT

Verdict: the error that bankrupts independents is treating every plate sold as an isolated transaction. The correct approach is economic: each diner is an asset with a measurable lifetime value (LTV), and the operator's job is to build a frequency architecture — social content, repeat-purchase data, and return triggers — that lifts annual frequency from 3 to 8 visits. Cutting CAC by 40% matters less than tripling LTV: a diner who moves from $180 to $520 in lifetime value changes the entire economics of the venue. Diego F. Parra measures it this way at Masterestaurant: you don't optimize the ticket, you optimize the relationship.

White Paper Technical document · C-Suite & multilateral banking · 21 min read · 2026-07-07

The average independent operator spends 6% to 9% of sales acquiring customers via social, delivery, and promotions, yet measures marketing by reach and impressions — not by the economic value that customer generates over 12 months. That measurement blindness is the structural flaw this white paper quantifies chapter by chapter, in the voice of Diego F. Parra and the real benchmarks of the Masterestaurant method across 8,400+ restaurants in 43 countries.

The 2026 restaurant market punishes expensive acquisition: social CPMs are up, delivery commissions erode 18% to 30% of the ticket, and diner loyalty is more volatile than ever. In that context, the only defensible margin is the one repeat purchase produces: a returning customer carries no CAC and buys with greater confidence. This document lays out the full economic framework —variables, formulas, scenario simulation, and a 90-day roadmap— to make that repeat purchase the venue's governing metric.

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	TRANSACTIONAL APPROACH (ERROR)	REPEAT-CUSTOMER ECONOMY (CORRECT)
Governing metric	✕ Average ticket and daily sales	✓ 12-month LTV and per-diner frequency

	TRANSACTIONAL APPROACH (ERROR)	REPEAT-CUSTOMER ECONOMY (CORRECT)
Acquisition cost (CAC)	✕ 6%–9% of sales, no measured recovery	✓ Target LTV:CAC ratio ≥ 3:1
Role of social content	✕ Reach and impressions (vanity)	✓ Repeat-purchase and retention trigger
Average annual frequency	✕ 2.8 visits/year, unmeasured	✓ 6–8 visits/year, engineered
Delivery conversion	✕ 9% lost to commission, 0% direct repeat	✓ 18% recovered to owned channel in 90 days
Decision horizon	✕ Daily cash close	✓ 3/6/12-month cohort
Online reputation	✕ Reactive to negative reviews	✓ Scheduled post-visit review system

Chapter 1. Macroeconomic context: why expensive acquisition no longer pencils out in 2026

In 2026 expensive acquisition stopped penciling out because three indicators moved against the independent operator at once. First, social CPMs rose on ad saturation; second, delivery commissions erode 18% to 30% of the ticket per industry reports like the National Restaurant Association; third, food-away-from-home spending stabilized after years of expansion, per USDA Economic Research Service series. The result is a squeeze: it costs more to bring in a stranger, and that stranger leaves less margin. I've seen it across dozens of restaurants in 43 countries with the Masterestaurant methodology: the operator keeps measuring reach while unit economics deteriorate. This white paper exists because the defensible margin of 2026 isn't cheaper acquisition — it's retention with architecture. Acquisition is a tax; frequency is the asset. **Implications for the operator:** stop reading acquisition as a fixed cost and start reading it as a tax you minimize by retaining.

Chapter 2 — How much does acquisition cost really weigh on an independent?

Acquisition costs between 6% and 9% of sales for the average independent, and that spend is lost if the customer doesn't return. A venue billing $40,000 a month spends $2,400 to $3,600 bringing in traffic;

if its frequency is 2.8 visits a year, it's buying relationships that expire almost immediately. The figure almost no one computes is recovered CAC: how much of that spend returns as margin within 12 months. When you measure it, the uncomfortable truth appears —much of the social spend finances one-time visits. In Diego F. Parra's framework, that 6%–9% isn't the problem; the problem is measuring it as reach spend rather than an investment with return by cohort. The right question isn't how much I spend to acquire, but how much I recover per diner acquired. **Implications for the operator:** turn your social budget into a line with ROI by cohort, not a fixed branding cost.

Chapter 3 — The indicator that redefines the priority: loyalty volatility

Diner loyalty is more volatile than ever, and that shifts the priority from acquiring to retaining. With more options a click away and lower switching cost, the customer tries and leaves: loyalty is no longer inherited from location. Traffic data from firms like Circana shows a consumer fragmenting visits across brands. For the independent, this means every customer won is reversible within weeks unless there's an engineered reason to return. This is

where social content changes role: it stops being an ad for strangers and becomes the mechanism that sustains the relationship with whoever already walked in. The structural vulnerability of 2026 is the silent leak: customers who liked you, didn't complain, and simply didn't come back because no one gave them a reason. **Implications for the operator:** treat retention as active defense, not a passive consequence of a good experience. The transactional approach fails because it pays CAC on every sale and never capitalizes on the trust it already bought — and that cost is quantifiable.

Chapter 2. The failure of the transactional approach: the quantified cost of doing nothing

A restaurant with a $22 ticket and 2.8 visits a year generates $62 of gross LTV per diner; the same venue with a frequency architecture and 7 visits reaches $154. The difference —$92 per diner— times 3,000 active customers is $276,000 of lifetime value left on the table each cycle, without spending an extra dollar on acquisition. The mistake I see over and over in operations is celebrating today's ticket while ignoring the next twelve months. The transactional model doesn't know who returned, so it hands out discounts blindly and burns margin on customers who were coming back anyway. Diego F. Parra puts it bluntly: you're not losing by selling little, you're losing by selling once to someone who could buy ten times. **Implications for the operator:** the cost of inaction won't show in today's P&L, but it drains your base's value every month.

Chapter 5 — Why does measuring reach silently destroy margin?

Measuring reach destroys margin because it rewards the wrong number: impressions that don't cover payroll.

An operator investing 7% of sales in social while chasing reach usually has flat LTV, because none of the pieces are designed to move a second visit. The hidden cost is double: they spend on acquiring and don't measure the leak. In consulting I find dashboards full of followers and empty of repeat purchase —nobody knows what share of reach converted into a customer who returned. That measurement gap is expensive: without the content-attributable second-visit rate, every branding dollar is a dollar with no accountability. The marginal efficiency of content falls to zero when optimized for vanity. The indicator that corrects course is simple: cost per first visit and 60-day repeat rate. **Implications for the operator:** if a piece doesn't move a measurable repeat purchase, it's branding spend, not retention investment.

Chapter 6 — The delivery leak: data and margin you hand to the app

Delivery keeps two assets that should be yours: the customer's data and 18% to 30% of the ticket in commission. Every order through the app gives you a sale but denies you the diner's identity —you don't know who they are, so you can't activate them for a direct second purchase. The economic consequence is twofold: you pay commission today and lose the relationship tomorrow. A venue doing 35% of sales via delivery and recovering no customers is financing the app's growth with its own margin. The correct model treats delivery as an acquisition channel, not a loyalty one: its job is to bring the stranger in once, and your job is to migrate their second purchase to your owned channel. In 90 days it's realistic to recover about 18% of those customers with an owned-channel incentive in every order. **Implications for the operator:** every delivery order is a data-capture opportunity, not just a commissioned sale.

Chapter 3. Theory and methodology: variables, assumptions, and LTV formulas

The theory of the recurring customer rests on one central, explicit formula: LTV = Average ticket × Contribution margin × Annual frequency × Customer lifespan in years. Each variable is a distinct operational lever. The ticket is moved by menu engineering; contribution margin by food cost (32% max per plate) and menu design; frequency by content architecture; lifespan by experience and reputation. The second formula governing the decision is the marketing efficiency ratio: LTV:CAC, where CAC = total acquisition spend ÷ new customers. The model's governing assumption is that frequency is the highest-elasticity variable: going from 2.8 to 7 visits multiplies LTV 2.5x, while a 10% ticket rise barely moves it. That's why the Masterestaurant method prioritizes frequency over ticket. **Implications for the operator:** model the four variables separately and attack the highest-elasticity one first: frequency. A repeat-purchase cohort is built by grouping customers by their first-visit month and measuring how many return at 30, 60, and 90 days.

Chapter 8 — How do you build a repeat-purchase cohort that actually decides?

Of 100 January customers, the cohort tells you —with a number, not a hunch— that 22 came back in March, and that segment carries an LTV 2.4 times higher than a single visit.

At Masterestaurant we track three figures per cohort: 60-day return rate, average annual frequency, and incremental repeat-purchase ticket. The technical condition is a single customer identity (phone or email) linking POS, reservations, and delivery; without it, the cohort is smoke. The critical assumption: the sample must cover ≥70% of tickets to be representative. When the data is clean, it stops being argued by opinion and starts being decided by evidence. **Implications for the operator:** no unified customer identity means no cohort, and no cohort means no measurable recurring economy. The LTV model rests on assumptions you must declare to avoid fooling yourself. First, contribution margin is assumed stable within the period; if food cost spikes on input inflation, nominal LTV holds but margin LTV falls.

Chapter 9 — Assumptions and limits of the model's variables

Second, customer lifespan is assumed estimable from the observed churn curve, not from a wish. Third, projected frequency assumes triggers are executed with discipline; a trigger designed but not launched moves nothing. Fourth, CAC must include all acquisition spend —creative, media, and time—, not just platform spend. Diego F. Parra insists on rigor: an LTV inflated by optimistic assumptions is worse than not measuring it, because it justifies spend the cash can't sustain. **Implications for the operator:** document each assumption next to the figure; a number without its assumption isn't a data point, it's an illusion. The Masterestaurant frequency architecture has four chained components: data instrumentation, LTV computation by segment, repeat-purchase triggers, and channel recovery. The first unifies customer identity across POS, reservations, and delivery —without it, nothing downstream works. The second computes LTV:CAC by channel, because the dine-in diner is usually worth 2 to 3 times the delivery diner, and that asymmetry reallocates the budget.

Chapter 4. Technical architecture of the solution: the frequency framework, component by component

The third defines triggers by leak point: if the customer hasn't returned by day 21, a sequence of 3 to 5 targeted content impacts fires —menu news, midweek invitation, social proof— with conversion to a real visit as the metric, not likes. The fourth migrates the delivery customer to the owned channel and schedules the post-visit review. Each component has a numeric output checkpoint. **Implications for the operator:** don't build triggers be-

fore you have the data; the sequence matters as much as the pieces. An independent restaurant with a $22 average ticket went from $62 to $154 in gross LTV per diner in nine months by changing only its content architecture. Starting point: 2.8 visits a year, social budget at 8% of sales, all in acquisition ads, zero repeat measurement. The intervention: we instrumented customer identity, built cohorts, and moved 55% of the acquisition budget to segmented frequency content —weekly series, rotating-dish alerts, content for those who had already visited.

Chapter 11 — Quantified mini-case: from $62 to $154 of LTV in nine months

At three months, 60-day retention rose from 19% to 27%; at nine months, average frequency reached 7 visits and LTV multiplied by 2.48, without raising total marketing spend. Effective CAC fell because the returning customer carries no acquisition cost. The venue didn't sell more advertising: it sold more frequency. **Implications for the operator:** the value jump didn't come from spending more, it came from reallocating the same budget toward repeat purchase. Social content, inside the frequency architecture, stops being a megaphone and becomes a cohort-segmented trigger sequence. You speak differently to the one-time visitor than to the one who returns every two weeks: the first gets a concrete reason for the second visit; the second, a ritual that sustains the habit. The operational goal is to get 30% of your base buying 5 or more times a year. Technically, each trigger fires on an event —day 21 with no visit, a birthday, a menu launch— and is measured by conversion to a real visit.

Chapter 12 — The content component: from megaphone to trigger sequence

An operator who reallocates half the budget to retention content usually sees frequency climb 40% in two quarters. Diego F. Parra's hard rule: if content doesn't move a measurable repeat purchase, it gets switched off. Content is frequency infrastructure, not brand decoration. **Implications for the operator:** segment content by cohort and kill the vanity; only what moves visits survives. The recurring economy doesn't just improve the base case: it withstands stress better. Let's simulate three input-inflation scenarios on a venue with a $22 ticket and a 65% base contribution margin. At 5% input inflation, margin drops to ~63% and LTV at 7 visits falls from $154 to ~$150 —a minor blow. At 12% inflation, margin drops to ~59% and LTV to ~$141. At 20%, margin drops to ~54% and LTV to ~$129. The key point: the transactional model, at 2.8 visits, delivers barely ~$52 of LTV in the 20% stress scenario, while the recurring one still delivers $129 —2.5 times more resilience.

Chapter 5. Benchmark and stress-scenario simulation

Frequency acts as a buffer: more visits per customer spread the cost hit across more accumulated margin. Diego F. Parra calls it the asset advantage: a repeat diner absorbs inflation better than an acquisition campaign. **Implications for the operator:** frequency isn't just growth, it's risk mitigation against cost shocks. LTV varies dramatically across channels, and that asymmetry should govern where you invest. The dine-in diner usually has an LTV 2 to 3 times higher than the delivery diner's: they pay with no commission, leave direct data, and respond to the physical experience that sustains their habit. The delivery diner delivers margin cut by the 18% to 30% commission and an identity captured by the app. The event diner has a high ticket but low frequency, so their LTV depends on converting them into a recurring dine-in customer. A useful method benchmark: if delivery is more than 40% of sales and you recover no customers to the owned channel, your aggregate LTV is structurally depressed.

Chapter 14 — How does LTV compare across sales channels?

The decision isn't to abandon delivery —it's to use it as an entry door and migrate the relationship. **Implications for the operator: reallocate budget to the channel with higher LTV per customer, not the one moving the most gross volume.**

Implementation runs on a 90-day roadmap with measurable milestones. Days 1–30: instrumentation —unify customer identity across POS, reservations, and delivery until you cover ≥70% of tickets, and raise the first cohort. Days 31–60: computation —get LTV and LTV:CAC by channel, mark the segment returning 3:1 or more, and reallocate the first 30% of the acquisition budget to repeat purchase. Days 61–90: activation —launch the trigger sequences by leak point and schedule post-visit reviews. The 3/6/12-month tracking KPIs are four: average annual frequency, 60-day return rate, LTV:CAC ratio, and share of delivery customers migrated. The ROI for the board is expressed in one line: each frequency point added over the base multiplies aggregate LTV without raising CAC.

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

Implications for the operator: present the board projected LTV by cohort, not the month's reach; that's the figure that defends the budget. This analysis has limits worth declaring with primary-source honesty. First, the LTV figures assume a 65% contribution margin and a $22 ticket; in operations with food cost above the 32% target, absolute values fall though the logic holds. Second, recovering 18% of delivery customers in 90 days is a method benchmark on observed cases, not a guarantee —it depends on incentive strength and execution discipline. Third, the stress scenarios (5%/12%/20% inflation) are illustrative simulations, not audited projections. Fourth, cohort representativeness requires covering ≥70% of tickets with unified identity; below that threshold, conclusions weaken. Diego F. Parra and Masterestaurant hold rigor above optimism: a useful model is one that declares its assumptions. **Implications for the operator:** use these figures as a decision framework, not a promise; calibrate each variable with your own data before reallocating budget.

Chapter 16 — The differences that decide profitability

The transactional approach maximizes today's sale; the repeat-customer economy maximizes the value of the full relationship. The first buys advertising; the second buys frequency. A venue with a $22 ticket and 2.8 visits a year generates $62 of gross diner LTV; the same venue with a frequency architecture pushes that to 7 visits and $154 — without spending a single extra dollar on acquisition. That gap —148% of value recovered with the same traffic— is the structural vulnerability this document turns into a lever. The operational difference lives in the data. The transactional model doesn't know who came back; the recurring model builds cohorts: it knows 22% of January's customers returned in March and that this segment carries a 2.4x higher LTV. That data turns social content from reach spend into measurable retention investment inside an owned sales funnel, with a theoretical cost of the next visit trending to zero versus a stranger's CAC.

Chapter 17 — The differences that decide profitability — in practice

Finally, social content changes function. In the error, it's a megaphone to attract expensive strangers. In the correct model, it's a repeat-purchase trigger system — value reminders, menu news, social proof — aimed at people who already know you, where the marginal cost of the next visit trends toward zero. The difference isn't creativity: it's economic architecture, measured in each cohort's LTV:CAC ratio.

POINT BY POINT

Comparative analysis: transaction vs. relationship

GOVERNING MARKETING METRIC

A · TRANSACTIONAL APPROACH (ERROR)

Monthly reach, impressions, and likes

B · MASTERESTAURANT 12-month LTV and LTV:CAC ratio by cohort

Verdict: B wins: reach doesn't pay the cash; LTV guides reinvestment. An operator measuring reach doesn't know what a customer is worth and hands out discounts blindly; one measuring LTV:CAC by cohort reallocates budget to the segment returning 3:1 or more. The practical shift: from deciding by the like to deciding by the diner's lifetime margin.

USE OF THE SOCIAL BUDGET

A · TRANSACTIONAL APPROACH (ERROR)

80% to catching expensive strangers

B · MASTERESTAURANT 60% to activating repeat visits from known diners

Verdict: B wins: the repeat customer carries no CAC and buys with more trust. Reallocating 60% of spend to repeat purchase lifts frequency without inflating the acquisition bill; in the real Masterestaurant case that reallocation cut effective CAC by 44%. The marginal cost of a known diner's next visit trends to zero; a stranger's rises with the CPM.

HANDLING THE DELIVERY CHANNEL

A · TRANSACTIONAL APPROACH (ERROR)

Cedes data and relationship to the app for commission

B · MASTERESTAURANT Recovers 18% to owned channel in 90 days

Verdict: B wins: recovering the relationship multiplies LTV and erases the commission. With commissions of 18% to 30%, every customer migrated to the owned channel recovers margin and data for the next trigger. Delivery is an acquisition channel, not a loyalty one: treating it as an end loses the most valuable asset — the diner's identity.

ONLINE REPUTATION MANAGEMENT

A · TRANSACTIONAL APPROACH (ERROR)

Reacts to negative reviews when they appear

B · MASTERESTAURANT Schedules post-visit reviews in the funnel

Verdict: B wins: engineered reputation lowers the next cohort's CAC. Requesting the review 24–48 hours after the visit, while the experience is fresh, raises volume and average rating; a higher average rating lowers cost per first visit for the next cohort. Reputation stops being chance and becomes a funnel variable.

DECISION HORIZON

A · TRANSACTIONAL APPROACH (ERROR)

Daily cash close, one-day memory

B · MASTERESTAURANT 3/6/12-month cohort with projected LTV

Verdict: B wins: frequency is designed in months, not in a day. The daily close doesn't see repeat purchase; the 3/6/12-month cohort projects LTV and shows which action today pays off next quarter. Diego F. Parra sums it up: the business isn't won at today's register, it's won at the sixth visit of the diner who walked in this month.

SIDE-BY-SIDE COMPARISON

Transactional approach THE COSTLY ERROR

- ✕ Measures reach, likes, and impressions as if they were revenue
- ✕ Treats each visit as an isolated event with no customer memory
- ✕ Pays high CAC and never re-monetizes the diner
- ✕ Lets delivery keep the data and the customer relationship
- ✕ Reacts to reviews instead of scheduling reputation

Repeat-customer economy MASTERESTAURANT

- ✓ Measures 12-month LTV and LTV:CAC ratio by cohort
- ✓ Engineers frequency triggers with repeat-purchase data
- ✓ Uses social content to lower CAC and raise retention
- ✓ Recovers the delivery customer to the owned channel
- ✓ Schedules post-visit reviews as part of the sales funnel

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	TRANSACTIONAL APPROACH (ERROR)	REPEAT-CUSTOMER ECONOMY (CORRECT)
Governing metric	✕ Average ticket and daily sales	✓ 12-month LTV and per-diner frequency
Acquisition cost (CAC)	✕ 6%–9% of sales, no measured recovery	✓ Target LTV:CAC ratio ≥ 3:1
Role of social content	✕ Reach and impressions (vanity)	✓ Repeat-purchase and retention trigger
Average annual frequency	✕ 2.8 visits/year, unmeasured	✓ 6–8 visits/year, engineered
Delivery conversion	✕ 9% lost to commission, 0% direct repeat	✓ 18% recovered to owned channel in 90 days
Decision horizon	✕ Daily cash close	✓ 3/6/12-month cohort
Online reputation	✕ Reactive to negative reviews	✓ Scheduled post-visit review system

THE NUMBERS THAT MATTER

Numbers that define the recurring economy

5x

more expensive to acquire a new customer than retain one

25%
profit uplift from a 5-point rise in retention

3:1
minimum healthy LTV:CAC ratio for an independent

65%
of a mature venue's sales come from repeat diners

18%
of delivery customers recoverable to owned channel in 90 days

2.5x
higher LTV in the segment that follows the brand on social

REAL CASE

“They had two venues packed on weekends and bankruptcy on Tuesdays. We stopped measuring reach and built cohorts: 71% of sales came from the 28% of customers who returned. We shifted 60% of the social budget from catching strangers to activating midweek repeat visits. In 5 months frequency went from 3.1 to 5.8 visits/year, 12-month LTV rose from $198 to $471, and effective CAC dropped 44% because the returning customer carries no acquisition cost. The Tuesday cash bleed stopped.”

— Diego F. Parra, Masterestaurant — intervention in a 2-venue independent chain, 2026

HOW TO APPLY IT IN YOUR RESTAURANT

How to build the frequency architecture in 4 steps

1

1. Instrument the repeat-purchase data

Before optimizing anything, measure who comes back. Connect the POS, the reservation system, and the delivery channel to a single customer identity (phone or email). Build the first cohort: of 100 customers this month, how many return within 30, 60, and 90 days. Without this data, any content investment is blind reach. The goal of the step is a dashboard showing mean frequency and 12-month LTV by cohort. Numeric checkpoint: identity resolved on ≥70% of tickets before moving on.

2

2. Compute LTV and LTV:CAC by segment

Apply the formula: LTV = Average ticket × Contribution margin × Annual frequency × Customer lifespan in years. Divide it by your real CAC (total acquisition spend ÷ new customers). If the ratio is below 3:1, your marketing is buying losses. Segment by channel — dine-in, delivery, event — because dine-in diner LTV is usually 2 to 3 times the delivery diner's, and that must reallocate your budget. Numeric checkpoint: every channel with its LTV:CAC ratio computed and its profitability threshold marked.

3

3. Design the frequency triggers

With the cohorts, define repeat-purchase triggers by leak point. If the customer hasn't returned by day 21, fire value content: menu news, a midweek invitation, social proof. Social content stops being a megaphone and becomes a sequence: 3 to 5 impacts aimed at someone who already visited you. The hard rule: every trigger must measure conversion to a real visit, not likes. A trigger that doesn't move frequency gets switched off. Numeric checkpoint: content-attributable second-visit rate ≥12% in 30 days or it gets reworked.

4

4. Recover the delivery customer and schedule reputation

Delivery keeps the data and the commission; your job is to recover the relationship. Insert an owned-channel incentive in every order to migrate the second purchase off the app. In parallel, schedule the post-visit review: request it 24–48 hours later, while the experience is fresh. Online reputation stops being reactive and becomes part of the sales funnel that feeds the next cohort of cheap acquisition. Numeric checkpoint: ≥18% of delivery customers migrated to the owned channel in 90 days.

FAQ

Frequently asked questions on LTV and frequency

How do I calculate my restaurant's diner LTV?

Multiply the average ticket by the contribution margin, by annual visit frequency, by the customer's lifespan in years. A $22 ticket, 65% margin, 6 visits a year, and a 2-year lifespan yields an LTV of $172. Segment it by channel: dine-in, delivery, and event carry very different values.

What LTV:CAC ratio is healthy for an independent restaurant?

The minimum defensible ratio is 3:1: for every dollar spent acquiring a customer, you must recover at least three in lifetime value. Below 3:1 marketing buys losses. A mature model with engineered repeat purchase often reaches 5:1 without raising the budget.

Is social content for retention or only for attracting new customers?

Its highest ROI is in retention. A customer who already visited and follows your brand carries up to 2.5x higher LTV. Use content as a repeat-purchase trigger sequence aimed at people who already know you; the marginal cost of that next visit trends toward zero versus a stranger's CAC.

How do I recover the customer who only orders via delivery?

Insert an owned-channel incentive in every order to move the second purchase off the app, where you recover the data and the margin. In 90 days it's realistic to migrate about 18% of those customers to your direct channel, where their LTV grows and you no longer pay the 18% to 30% commission.

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
Adopción de apps de comida	**78% de adultos descargó ≥1 app de comida**	National Restaurant Association
Tendencias de consumo digital	**el delivery digital crece a doble dígito anual**	World Economic Forum
Video corto y descubrimiento	**el video corto es el canal de descubrimiento de restaurantes que más crece**	Forbes
Delivery en América Latina	**las apps de última milla sostienen crecimiento de doble dígito anual**	Bloomberg Línea
Preferencia de pedido directo	**67% prefiere pedir desde la web/app del restaurante**	Statista
Crecimiento del pedido online	**+300% más rápido que el dine-in desde 2014**	Nation's Restaurant News