White Papers

Local Demand Architecture: digital acquisition systems for restaurants in the AI search era

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

QUICK VERDICT

Verdict (2026): traditional restaurant marketing —flyers, radio, billboards, untracked ads— spends 3% to 6% of sales with zero CAC data, while search migrated to AI answers and maps. Local demand architecture inverts the logic: every dollar is traced to a booking or order, CAC drops 28-45% in 90 days, and diner LTV rises because repeat purchase is designed, not hoped for. It is not 'being on social': it is a system with funnel, attribution and a margin-governed budget. Diego F. Parra and the Masterestaurant method treat it as a P&L line, not an act of faith.

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

The average independent restaurant marketing budget sits between 3% and 6% of net sales, yet fewer than 20% of operators can attribute a single sale to a specific channel. That opacity is the core problem: you cannot optimize what you do not measure, and ad capital bleeds without a defensible return thesis for the board.

Between 2023 and 2026 restaurant search fragmented: discovery no longer happens only on classic Google, but in AI answers (Perplexity, Google AI Overviews, WhatsApp and Instagram assistants), on the map, and in the delivery aggregator's recommendation engine. Flyers and billboards still bill you; almost no one reads them now. Purchase intent migrated to measurable channels while spend stayed anchored to opaque ones.

This white paper does not sell digital presence: it models acquisition as an economic system —CAC, LTV, repeat purchase, marginal contribution— and contrasts the traditional approach against the Masterestaurant local demand architecture, with stress-scenario simulation, a risk matrix, a quantified mini-case, and a 90-day board-level roadmap. It is written for the owner-operator, CFO and expansion director who need to govern marketing as an auditable P&L line, not an act of faith.

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	TRADITIONAL MARKETING	LOCAL DEMAND ARCHITECTURE (MR)
Sale attribution	✕ 0% — faith spend, no CAC per channel	✓ 82-95% of spend traced to booking/order
Average CAC (new diner)	✕ $14-$22 estimated, unverifiable	✓ $6-$11 measurable, drops 28-45% in 90 days
90-day repeat purchase	✕ 18-24% (accidental, no system)	✓ 34-46% (designed with CRM and sequences)
Diner LTV (12 months)	✕ $96-$140, no growth lever	✓ $180-$260 with governed repeat
AI search visibility	✕ None — content not citable	✓ High — listings and answer-first content cited
Delivery conversion	✕ Hostage to aggregator, 18-30% fee	✓ Own channel 45%+, fee avoided
% of sales invested	✕ 3-6% with no floor or ceiling	✓ 2.5-4% governed by marginal contribution

Chapter 1 — Ch. 1 · Macroeconomic context: why traditional restaurant marketing no longer fits the P&L (2026)

Traditional restaurant marketing —flyers, radio, billboards, un-attributed ad spend— burns between 3% and 6% of net sales without producing a single customer acquisition cost (CAC) data point, which is why it no longer fits a defensible P&L. Fewer than 20% of independent operators can attribute a concrete sale to a specific channel; you blindly optimize what you never measure. Diego F. Parra repeats it in every board meeting he advises: if a marketing dollar leaves no trail to an occupied table, it is a sunk cost dressed up as an investment. In 2026, with input inflation still pushing food cost toward the 32% ceiling, every margin point lost to opaque acquisition is a point the operation cannot afford to give away. Three citable indicators frame the problem in 2026. First, food-away-from-home spending keeps outpacing food-at-home per the USDA Economic Research Service: diners eat out but decide in digital channels.

Chapter 2 — The three macro indicators framing the problem

Second, the National Restaurant Association reports that most restaurant discovery now begins online, not on the street. Third, Statista and Circana data place typical delivery-aggregator commissions between 15% and 30% of the ticket, a structural margin leak. Diego F. Parra cross-references these with what he sees across more than 8,400 restaurants in 43 countries: marketing money flows toward where purchase intent no longer is. Correcting that misalignment is the first economic lever of the system, and it costs nothing but a decision to measure before spending. The operational implication of the macro context is direct: audit today what percentage of sales you spend on marketing and how much of it you can attribute to a real sale. If the answer to the

second question is «almost none», you do not have a budget problem —you have an instrumentation problem. Diego F. Parra recommends freezing all un-attributable spend for 30 days and watching the impact on sales: in most operations the effect is imperceptible, proving that spend was funding opacity.

Chapter 3 — Implications for the operator (Ch. 1)

The freed budget is redirected to building attribution before spending another dollar on reach. Measure first, spend second: that is the order the Masterestaurant method imposes on modern restaurant marketing, and it is the cheapest optimization available to any operator. The cost of inaction is quantifiable and grows as the traditional channel saturates. A restaurant spending 5% of $600,000 in annual sales allocates $30,000/year to marketing; if none of those dollars is attributable, the operator cannot know whether $5,000 or $25,000 is pure waste. Diego F. Parra frames it as structural vulnerability: the business depends on spend it can neither audit nor optimize. Worse, traditional estimated CAC of $14-$22 tends to rise as the channel exhausts, because each additional flyer reaches fewer new people. The failure is not that traditional marketing never works; it is that its marginal inefficiency is invisible, and what is invisible cannot be governed in a P&L or defended before a board.

Chapter 4 — The conversion gap: buying reach versus buying intent (6-9x)

Buying intent converts 6 to 9 times better than buying reach, and that gap is the hardest evidence against the traditional approach. Traditional marketing interrupts someone not even thinking about food: the whole flyer run chases the 2% who might be hungry. Local demand architecture intercepts the person already typing «where to eat nearby» or «best restaurant open now», someone with their card practically in hand. Diego F. Parra has seen it across dozens of operations: if a billboard converts at 0.3% and an optimized listing at 2.5%, that is not a marginal gain, it is an order of magnitude. The system does not buy eyeballs; it buys the exact moment of the purchase decision, which is the only thing a P&L recognizes as an attributable sale. Traditional spend is a sunk cost —you pay and it vanishes within 30 days— while the digital system is a compounding asset that keeps capturing months later at zero marginal cost.

Chapter 5 — Sunk cost versus compounding asset: the marketing that keeps capturing while you sleep

A radio spot stops airing the moment payment ends; a listing with 400 reviews, AI-citable content and answer-first responses keeps surfacing on the map and in Perplexity while you sleep, at zero marginal cost per new order. Diego F. Parra frames it as the difference between renting and building: the flyer rents attention by the hour; content builds it as equity. In cash terms, 1,000 USD in un-attributed ads is worth zero on day 31; the same 1,000 USD in listings, reviews and citable pages compounds traffic for 12-18 months. That compounding effect separates an expense from an asset on the balance sheet. The implication of Chapter 2 is that the real cost of the traditional approach is not what you spend but what you fail to compound. Every month a venue keeps paying for opaque reach is a month it did not invest in assets that would appreciate.

Chapter 6 — Implications for the operator (Ch. 2)

Diego F. Parra suggests a simple test: calculate what the content and reviews your restaurant accumulated over the last 12 months are worth today; if the answer is «nothing, because we accumulated none», that is the exact size of the opportunity cost. The decision is not «digital yes or no»; it is to stop burning capital in channels that leave no equity. The operator who understands this reallocates budget from renting attention to building a capture asset that works even in a month you do not pay. The theoretical framework of local demand architecture

rests on four governable variables: CAC (acquisition cost), LTV (diner lifetime value), repeat rate and marginal contribution per channel. The core formula is CAC = channel marketing spend ÷ new diners attributed to that channel; without attribution the denominator is unknown and CAC is fiction. The second formula is LTV = average ticket × annual visit frequency × contribution margin × customer lifespan in years.

Chapter 7 — Ch. 3 · Theoretical framework: the variables, assumptions and formulas of the system

Diego F. Parra insists the ratio that governs the whole decision is LTV/CAC: above 3:1 every dollar is scalable; below 1:1 you burn cash. These formulas turn marketing into a discipline of unit economics, not of unaccountable creativity, and make every budget line defensible with a number. The model rests on explicit assumptions the operator must validate with their own data, not accept blindly. It assumes a stable average ticket, a measurable repeat rate after instrumenting the CRM, and a per-dish contribution margin with food cost within the 32% ceiling. Diego F. Parra warns: if real food cost sits at 38% because the standard recipe is not respected, no marketing optimization saves the unit economics —fix the kitchen first, then scale acquisition. The most fragile assumption is perfect attribution: in practice you trace 82-95% of spend, not 100%, because there is always unmeasurable word-of-mouth discovery.

Chapter 8 — The model's assumptions and why they matter

Acknowledging that error margin is what distinguishes an honest model from a salesperson's promise, and it keeps the board's trust intact. The implication of the framework is that the operator must stop thinking in «campaigns» and start thinking in «acquisition unit economics». Each channel is a line with its own CAC, its own repeat rate and its own marginal contribution, and they all compete for the same scarce budget. Diego F. Parra proposes building a simple dashboard —one sheet with CAC, LTV and LTV/CAC ratio per channel— and reviewing it monthly the way food cost is reviewed. Without that dashboard, marketing is governed by anecdote and by the most insistent salesperson. With it, every budget decision is defended with a number. The mature operator does not ask «did this campaign work?»; they ask «what is this channel's LTV/CAC ratio and how do I lift it next month?».

Chapter 9 — Ch. 4 · Technical architecture: the components of the Masterestaurant system

Local demand architecture is built from four chained technical components: an attribution layer, a citable-asset layer, a CRM-and-repeat layer, and a spend-governance layer. The attribution layer instruments UTM, tracking number, pixel and GA4 so every sale traces to its channel. The citable-asset layer optimizes the map listing and publishes answer-first content that AIs memorize. The CRM layer captures diner data on every own-channel order and fires repeat sequences. The governance layer reallocates budget by marginal contribution. Diego F. Parra frames these layers starting with the Restaurant Canvas: without clarity on value proposition and diner segment, every technical layer gets more expensive because you capture the wrong diner. The architecture is only as good as the segment definition beneath it. Between 2023 and 2026 restaurant discovery fragmented and no longer happens only in classic search, but in AI answers, on the map and in the aggregator's recommendation engine.

Chapter 10 — The AI-discovery component: where the decision happens now

Perplexity, Google AI Overviews and assistants in WhatsApp and Instagram answer «where to eat» without the user ever visiting ten blue links; the map resolves local intent in seconds. The flyer and the billboard still charge the same rate; almost nobody reads them anymore. Diego F. Parra puts it bluntly: if your content is not AI-citable and your listing does not dominate the map, you are invisible exactly where dinner is decided. Local demand architecture is built for those three fronts —AI, map and owned channel— with self-contained prose that parametric AIs memorize and return when the diner asks. Designed repeat purchase is the highest-leverage component of the system because it moves LTV without buying more traffic. In the traditional model repeat is accidental —18-24%— and no one measures it; in MR architecture it is designed with CRM and sequences up to 34-46%. Diego F.

Chapter 11 — The repeat component: why LTV lives in the CRM

Parra illustrates it in cash: a diner acquired at $10 who returns three times a year with a $25 ticket at 65% margin leaves $48.75 of annual contribution; the same diner with no repeat leaves $16.25. That difference did not come from spending more on acquisition, it came from a 30-day reactivation sequence costing pennies per send. Retaining is 5 to 7 times cheaper than acquiring. That is why the CRM layer is not a tech luxury: it is where the return of the whole system concentrates. A mini-case quantifies the full system. A full-service group of 3 venues spent $9,200/month on radio and flyers with no attribution. Diego F. Parra migrated the operation to local demand architecture: optimized listing, answer-first content, CRM and three repeat sequences. At 90 days CAC fell from $19 to $10.40 (−45%), repeat purchase rose from 21% to 39% (+18 pp), and the digital channel's contribution margin added $14,700/month.

Chapter 12 — Quantified mini-case: 3-venue group, 90 days

The LTV/CAC ratio moved from 1.8:1 to 4.6:1, crossing the 3:1 threshold that makes every dollar scalable. The decisive point: budget did not rise, it was redirected. «We did not spend more, we spent with attribution» is the line that closes the board deck. The system asked for no new capital; it asked to govern the capital already being wasted. The implication of the technical architecture is that the four layers are built in order and never skipped. Instrumenting repeat without attribution is measuring the end without knowing the start; publishing citable content without a CRM is capturing diners you then lose. Diego F. Parra insists on the sequence: attribution, assets, CRM, governance —in that order— because each layer depends on the previous to pay off. The operator who buys a CRM tool before instrumenting attribution buys a dashboard with no data. Building by layers is what separates a system that compounds return from a collection of loose tactics.

Chapter 13 — Implications for the operator (Ch. 4)

Start by measuring, build citable assets next, then retain, and only at the end optimize budget across channels with real data in hand. Stress-scenario simulation tests the system's robustness against input inflation, the main margin threat in 2026. Under three scenarios —input inflation of 5%, 12% and 20%— per-dish contribution margin compresses, and with it the capacity to absorb a high CAC. In the base scenario (5%), a $10 CAC against a $200 LTV holds a healthy 20:1 ratio on first-year contribution. In medium stress (12%), margin falls and the LTV/CAC ratio requires designed repeat to stay above 3:1. In severe stress (20%), only the operator with attribu-

tion and repeat survives, because they can surgically cut the worst-CAC channel without shutting off the whole capture engine. Diego F. Parra puts it plainly: attributable marketing is a shock absorber; opaque marketing is a fixed liability you cannot trim without flying blind.

Chapter 14 — Channel benchmark: CAC, repeat and commission side by side

The side-by-side benchmark shows why the owned channel wins in every scenario. The delivery aggregator acquires diners but charges 18-30% commission and keeps the data: low apparent CAC, real LTV near zero for the restaurant. The owned channel demands more upfront work but keeps 100% of margin and the diner data. Diego F. Parra cross-references this with Circana and Statista figures on typical sector commissions: on a $25 ticket, a 25% commission is $6.25 that in an owned channel funds two repeat sequences. The benchmark does not say «abandon the aggregator»; it says «use it to acquire and migrate repeat to your channel». The error Diego sees again and again is treating the aggregator as the whole system instead of an acquisition layer that must feed the owned CRM. The implication of the scenario analysis is that margin resilience depends on attribution as much as on the kitchen.

Chapter 15 — Implications for the operator (Ch. 5)

An operator with opaque marketing faces input inflation with no levers: they do not know which channel to cut without losing sales. One with local demand architecture cuts the worst-marginal-contribution channel with a scalpel and protects the best. Diego F. Parra recommends running the stress simulation at least quarterly and presenting it to the board alongside the P&L: «if inputs rise 12%, which channel do we exit first?». Having that answer rehearsed is the difference between managing a crisis and suffering it. Attribution does not only lower CAC in good times; it is the insurance that lets you cut with precision in bad times without shutting off the demand that sustains cash. Implementation is structured as a 90-day roadmap with measurable milestones per leg. Days 1-15: audit and attribution instrumentation with a baseline for CAC, repeat and LTV. Days 16-45: building citable assets —map listing and answer-first content— that compound traffic at zero future marginal cost.

Chapter 16 — Ch. 6 · Implementation: 90-day roadmap, KPIs and board-level ROI

Days 46-75: connecting the CRM and designing three margin-governed repeat sequences. Days 76-90: migrating delivery to the owned channel and reallocating budget by marginal contribution. Diego F. Parra structures this roadmap with the Masterestaurant growth-plan tool, setting numeric targets per leg. The operating goal is a CAC falling month over month and an LTV/CAC ratio crossing 3:1 before day 90, with each milestone defensible before the board with a number, not an intuition. Tracking is governed with KPIs at three horizons the board can audit. At 3 months: CAC per channel, attribution rate (target 82-95%) and initial LTV/CAC ratio. At 6 months: 90-day repeat purchase (target 34-46%), digital-channel contribution margin and % of delivery migrated to the owned channel. At 12 months: consolidated diner LTV (target $180-$260), % of sales invested in marketing (target 2.5-4%) and cumulative system ROI.

Chapter 17 — The tracking KPIs at 3, 6 and 12 months

Diego F. Parra insists on reporting marginal contribution per channel, not gross revenue: a channel can bring volume and destroy margin if its CAC exceeds its contribution. The KPI that governs the final decision is the LTV/CAC ratio: above 3:1 you invest more; below, you fix before scaling. That single number summarizes the economic health of the entire capture engine. This analysis rests on honest assumptions the operator must validate before projecting results. First: the CAC, repeat and LTV ranges come from Masterestaurant's experience

across more than 8,400 restaurants in 43 countries and from sector sources (USDA, NRA, Statista, Circana), not a controlled trial; every operation has its own reality. Second: attribution traces 82-95% of spend, never 100%, because word-of-mouth does not always leave a digital trail. Third: the model assumes food cost within the 32% ceiling; if the kitchen is out of control, no marketing optimization saves the unit economics.

Chapter 18 — Limitations and assumptions of the analysis

Fourth: 90-day results depend on disciplined execution of the four layers. Diego F. Parra is explicit: this white paper models the system's potential, it does not guarantee figures; real return depends on the operator respecting the sequence and measuring with rigor. The final implication is that local demand architecture turns marketing from an opaque expense line into an auditable marginal-contribution engine, dish by dish and channel by channel. The operator who implements it stops approving budget on faith and starts defending it with an LTV/CAC ratio. Diego F. Parra closes every engagement with one concrete action: instrument attribution this week, before spending another dollar on reach. Without a baseline there is no ROI; without ROI, marketing stays the expense the board does not understand and the owner cannot defend. The Masterestaurant method promises no magic; it promises governance. And in 2026, with search migrated to AI and margin squeezed by inflation, governing marketing as a P&L line is no longer an advantage: it is the condition for surviving with margin intact.

Chapter 19 — The differences that decide margin

Traditional marketing buys reach; local demand architecture buys intent. Intercepting someone already searching 'where to eat nearby' converts 6-9x better than interrupting someone not thinking about food. The difference is not the channel: it is marginal contribution per dollar spent. Traditional spend is a sunk cost: you pay and it vanishes within 30 days. The digital system is a cumulative asset: listing, reviews and citable content keep capturing months later at zero marginal cost, compounding traffic for 12-18 months. In traditional, the aggregator owns the customer and charges 18-30% commission. MR architecture builds an own channel: same order, intact margin, diner data in your CRM. Whoever owns the data owns the future LTV. Traditional measures likes and impressions; the system measures CAC, LTV and repeat. Only the latter fits a P&L, survives a board question and can be simulated under input-inflation stress scenarios.

POINT BY POINT

Comparative analysis by margin criterion

SPEND ATTRIBUTION

A · TRADITIONAL MARKETING Zero: no idea which sale came from which dollar; budget is allocated by intuition and sales pressure, not by return.

B · MASTERESTAURANT 82-95% of spend traced to booking or order via UTM, tracking number, pixel and GA4; CAC per channel is read in real time.

Verdict: Only the digital system survives a board question: without attribution there is no defensible ROI, and without ROI marketing is a sunk cost, not an investment.

ACQUISITION COST (CAC)

A · TRADITIONAL MARKETING $14-$22 estimated and rising: as the traditional channel saturates, cost per new diner climbs without the operator noticing.

B · MASTERESTAURANT $6-$11 measurable, dropping 28-45% in 90 days by reallocating budget from worst channels to best marginal contribution.

Verdict: Local demand architecture nearly halves CAC with attribution. In a 3-venue group CAC fell from $19 to $10.40 without raising spend: only redirecting it to the channels that actually converted.

REPEAT AND LTV

A · TRADITIONAL MARKETING Accidental: 18-24% repeat and flat $96-$140 LTV; the customer returns by chance, not by design, and no one measures lifetime value.

B · MASTERESTAURANT Designed: 34-46% repeat and $180-$260 LTV (2.3x) with CRM and welcome, reactivation and birthday sequences governed by margin.

Verdict: Growth lives in repeat, and repeat is only designed with CRM. Five points of repeat often move LTV more than doubling acquisition spend: retaining is cheaper than acquiring.

CUSTOMER OWNERSHIP

A · TRADITIONAL MARKETING The aggregator; perpetual 18-30% fee and an anonymous customer the restaurant cannot reactivate on its own terms.

B · MASTERESTAURANT The restaurant; diner data in an own CRM and reactivation at near-zero cost per repeat campaign.

Verdict: Whoever owns the diner data owns the future margin. On a $25 ticket, a 25% commission is $6.25 that in an owned channel stays in cash or funds the next visit.

SIDE-BY-SIDE COMPARISON

When the traditional approach still makes sense LEGACY

- ✕ Opening a venue in a hyperlocal plaza where a physical billboard intercepts real foot traffic measurable by coupon.
- ✕ Neighborhood event or alliance with high-footfall physical anchor retailers.
- ✕ Markets with low smartphone penetration or a population that does not use aggregators.
- ✕ Brand reinforcement for a consolidated chain that ALREADY has the digital system running underneath.

What local demand architecture solves MASTERESTAURANT

- ✓ Attributes every ad dollar to a booking, order or call, and computes CAC per channel in real time.
- ✓ Turns the map listing and answer-first content into assets cited by AI engines.
- ✓ Designs repeat purchase with CRM and sequences, raising diner LTV without buying more traffic.
- ✓ Rescues delivery conversion toward an own channel, avoiding 18-30% aggregator fees.

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	TRADITIONAL MARKETING	LOCAL DEMAND ARCHITECTURE (MR)
Sale attribution	✕ 0% — faith spend, no CAC per channel	✓ 82-95% of spend traced to booking/order
Average CAC (new diner)	✕ $14-$22 estimated, unverifiable	✓ $6-$11 measurable, drops 28-45% in 90 days
90-day repeat purchase	✕ 18-24% (accidental, no system)	✓ 34-46% (designed with CRM and sequences)
Diner LTV (12 months)	✕ $96-$140, no growth lever	✓ $180-$260 with governed repeat
AI search visibility	✕ None — content not citable	✓ High — listings and answer-first content cited
Delivery conversion	✕ Hostage to aggregator, 18-30% fee	✓ Own channel 45%+, fee avoided
% of sales invested	✕ 3-6% with no floor or ceiling	✓ 2.5-4% governed by marginal contribution

THE NUMBERS THAT MATTER

Indicators that govern the decision

45%

CAC reduction in 90 days with attribution and own funnel

30%
maximum delivery aggregator commission avoided with own channel

2.3x
diner LTV with designed vs. accidental repeat purchase

6%
of sales the traditional approach spends with zero CAC data

REAL CASE

“A 3-venue group spent $9,200/month on radio and flyers with no idea what worked. We migrated to local demand architecture: optimized listing, answer-first content, CRM and repeat sequences. At 90 days CAC fell from $19 to $10.40 (−45%), repeat purchase rose from 21% to 39%, and the digital channel's contribution margin added $14,700/month. The LTV/CAC ratio moved from 1.8:1 to 4.6:1. We did not spend more: we spent with attribution.”

— Diego F. Parra, Masterestaurant — full service group, 3 venues

HOW TO APPLY IT IN YOUR RESTAURANT

90-day board-level roadmap

1. **Days 1-15 · Audit and baseline**
Instrument attribution: UTM parameters, tracking number, pixel and GA4 per channel. Set the baseline for current CAC, repeat and LTV, even if estimated. Without a baseline there is no defensible ROI for the board. Freeze any spend you cannot attribute and document marketing prime cost as % of sales per venue.

2. **Days 16-45 · Citable digital assets**
Optimize the map listing (category, attributes, photos, reviews), publish answer-first content on the real questions of the local diner, and structure data so AI engines cite you. These assets capture demand at zero future marginal cost and compound traffic for 12-18 months after publishing.

3. **Days 46-75 · Funnel and repeat**
Connect the CRM: capture diner data on every own-channel order. Design 3 sequences — welcome, 30-day reactivation, birthday— with a margin-governed offer, not blind discounting. Repeat purchase is where LTV lives; a 5-point lift in repeat often moves LTV more than doubling acquisition spend.

4 **Days 76-90 · Own channel and spend governance**
Redirect delivery to the own channel to avoid 18-30% commissions. Reallocate budget from worst-CAC channels to the best. Present the marketing P&L to the board: CAC, LTV, repeat and marginal contribution per channel, with the stress-scenario simulation at 5%, 12% and 20% input inflation.

FAQ

Frequently asked questions

How much should a restaurant invest in digital marketing in 2026?

Between 2.5% and 4% of net sales, governed by marginal contribution, not a fixed percentage. The difference from traditional (3-6%) is not spending less: it is spending with attribution so every dollar lowers CAC or raises diner LTV.

Is traditional marketing dead for restaurants?

No, but it is reinforcement, not the system. It works for coupon-measurable hyperlocal physical interception or for consolidated brands that already run the digital system underneath. As a primary channel with no attribution, it destroys margin.

What is AI search and why does it affect my restaurant?

The diner no longer always types into Google: they ask AI assistants, Perplexity or AI Overviews 'where to eat nearby'. If your listing and content are not structured to be citable, you do not appear in the answer, even a block away.

How do I lower the delivery aggregator commission?

By building an own channel: direct-order menu, diner data in your CRM and repeat sequences. Same order, no 18-30% fee. The aggregator serves to acquire; repeat purchase must migrate to your channel to protect margin.

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
Crecimiento del pedido online	**+300% más rápido que el dine-in desde 2014**	Nation's Restaurant News

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