 FOR CMOS OF 200+ UNIT FRANCHISED RESTAURANT BRANDS

The Franchise Marketing *Field* *Guide.*

*Your sales are up. Your traffic is not. What to fund,
what to measure, and what to stop doing before the
2027 budget locks.*

SECOND EDITION · UPDATED SEPTEMBER 2026

Contents.

A PRACTICAL REFERENCE,
WRITTEN FOR THE STRATEGIC
AND OPERATIONAL DECISIONS
IN FRONT OF YOU THIS QUARTER
AND NEXT.

SECOND EDITION. THE FIRST
EDITION IS SUPERSEDED BY
THIS ONE.

<i>00</i>	Foreword: the job in front of you
<i>I</i>	The visibility nobody owns
<i>II</i>	Agentic ordering, and the menu as a media asset
<i>III</i>	The channel stack, re-priced
<i>IV</i>	Your aggregator is now a media network
<i>V</i>	AI that moves the number
<i>VI</i>	Measurement that survives the board meeting
<i>VII</i>	The national-to-local handoff
<i>VIII</i>	The next 90 days, and the 2027 budget
<i>09</i>	What we got wrong
<i>10</i>	About Happy Cog

● FOREWORD

The job in front of you.

You run marketing at a restaurant brand with two hundred or more units. Some you operate. Most you don't. This year your sales are up and your traffic is not, and the distance between those two facts is a conversation you are about to have with your CFO.

Total restaurant and foodservice sales are tracking up roughly four percent in 2026. Strip out menu price increases and real growth is closer to one percent, with traffic flat to down across most segments. A marketing team can post a number that looks like growth without being able to say whether it bought a single additional visit.

This is a field guide for that problem.

It is the second edition. The first one argued that channel fragmentation had rewritten the marketing math, which was true and has since become a description of the weather. What changed after we published is more specific and carries deadlines. Discovery started moving into AI assistants. The delivery platforms turned themselves into media networks holding your purchase data. And Google began folding local advertising into automation you do not control, on a timetable that runs through 2027.

We have kept what held up, rewritten what did not, and added a short chapter at the back naming the calls we got wrong the first time. The opinions here are the ones we share with new clients in the first meeting. Some of them will challenge what your current vendors are telling you. We still think that is the point.

The Happy Cog Paid Media Team

Growth is priced, not trafficked. That is the sentence your CFO is about to say back to you.

I.

— PART ONE · DISCOVERY

The visibility nobody owns.

A customer asks an assistant where to eat before they open a browser. For a two-hundred-unit brand, that answer is assembled from two hundred separate location records, and the weakest one is usually not yours to fix.

Something changed in how people decide where to eat, and it did not arrive with a platform launch or a rate card. A customer opens ChatGPT, or Gemini, or the answer panel above their search results, and asks a question shaped like a conversation. *Somewhere to take my parents near Cherry Hill, not too loud, open past nine, not Italian again.* The assistant returns two or three names. Your brand is one of them, or it isn't.


For an independent restaurant this is a marketing problem with a tractable fix. For a brand with two hundred units it is an operations problem wearing marketing's clothes.

An assistant answering that question is not reading your national campaign. It is assembling an answer from structured location data (hours, address, price band, cuisine, service attributes, menu items), from review text and how recent it is, and from what your records say across Google Business Profile, Apple Business Connect, and the directories and aggregators that syndicate from them. It resolves all of it per location. So your visibility in AI search is not one number. It is two hundred numbers, and the median matters more than the best.

Now ask who maintains those records. At most franchised brands the honest answer is the franchisee, or a vendor the franchisee hired in 2022, or nobody since the last remodel. Corporate owns the brand page. The franchisee owns the location record. The assistant reads the location record.

This is the largest unfunded gap in franchise marketing right now, and it is unfunded precisely because it does not belong to anyone. It is too technical for the field team, too local for the brand team, and too unglamorous for the agency that sold you a campaign.

*Your AI visibility is not one number.
It is two hundred numbers, and the
median matters more than the best.*



One square per location record.
Schematic.

What the assistants actually weigh

01 Structured attributes, not prose. The assistant matches a query against fields it can read: price band, cuisine tags, service options, accessibility, hours by day, and increasingly individual menu items. A beautifully written location page with no structured data underneath it is close to invisible.

02 Review recency and specificity, not the star rating. A 4.3 with fifty reviews this quarter beats a 4.6 with nothing since last year, because the assistant is reading what people said. Repeated concrete specifics are what let it match a query: gluten free options, parking, quiet, good for groups, fast at lunch. Star rating alone no longer decides anything.

03 Consistency across sources. Conflicting hours between your site, your Google listing, and a delivery platform do not average out. They lower confidence, and a model with low confidence recommends somebody else.

04 Machine access to your menu. If your menu is an image, a PDF, or rendered by JavaScript that a crawler cannot execute, your menu does not exist to an assistant. This is the same requirement as the next chapter, arriving from a different direction.

VENDOR DATA

One 2026 vendor study put the share of restaurants effectively invisible in AI search at 83 percent. We cite it as vendor-sourced and note the methodology caveat: the figure comes from the vendor's own crawl of a sample it selected, and it should be read as an order of magnitude rather than a measurement. Your own prompt panel, described overleaf, produces a number you can defend.


A diagnostic you can run before your next leadership meeting

You do not currently have a number for this, and you can have one in an afternoon. Pick your ten largest markets. Write twenty-five prompts a real customer would use, phrased conversationally rather than as keywords, spread across occasions (dinner with family, quick lunch near the office, late night, dietary constraint, group booking). Run each one in ChatGPT, Gemini, and Perplexity. Score two things: whether your brand appears at all, and whether it holds the first recommendation.


Report it as two percentages by market. That is your baseline, it is repeatable monthly, and it is the first marketing number in years that a franchisee will immediately understand and care about.

10	25	3	2
MARKETS	PROMPTS EACH	ASSISTANTS	SCORES PER MARKET

THE SCORECARD



One block per assistant, one cell per prompt. Filled cells are appearances. Schematic.



Appearance rate and primary-recommendation rate, reported by market and re-run monthly. The series matters more than the first reading.

II.

—— PART TWO · TRANSACTION

Agentic ordering, and the menu as a media asset.

The brands moving first all made the same decision. Let discovery happen inside the assistant. Keep the checkout.

In April 2026, Starbucks launched an app inside ChatGPT that takes a description or a photograph and returns drink recommendations. Little Caesars launched a comparable app across the United States, Canada, and Mexico in the same month, pulling live menu, pricing, and location data into the conversation. In July, Square went further and automatically enrolled eligible United States food and beverage sellers into ordering through ChatGPT and Claude, at standard card processing rates rather than marketplace commission.

Look at what those three moves have in common. In every case discovery happens inside the assistant and checkout happens inside the brand's own app or site. Starbucks was explicit about why, and the reason is the loyalty program. Handing checkout to a third party is the aggregator trade repeated in a newer channel, and the brands who lived through the first one are not volunteering for the second.

We think that is the right call. We would add that most brands are not currently able to make it.

What the channel actually requires

For an assistant to recommend you, it needs to read your menu. For an assistant to transact, it needs a live feed of items, modifiers, prices, and availability, resolved to a specific location, exposed through an interface a machine can call. And for the resulting order to be worth taking, the money needs to settle to the correct franchise entity without a person reconciling it by hand at month end.

If you run one point-of-sale system inside a single corporate entity, that is a project. If you run four systems across three entity structures, which describes most two-hundred-unit brands we meet, that is *the* project, and the assistant channel is simply the thing that finally forces it. The work has a useful property, though: it is the same work Chapter I asks for. Clean structured menu and location data decides whether an assistant recommends you at all, long before it decides whether an assistant can transact with you.



(We built our own ordering platform partly for this reason, with multi-entity fund routing through Stripe Connect and live integrations across the major point-of-sale, payment, delivery, and loyalty systems. We would rather you understood the requirement than bought ours.)

The honest part

Order volume through assistants is small today and will stay small for a while. Nobody should reallocate a media budget on the strength of it this year.

The reason to move now is arithmetic and timing. The fee structure is the best in the off-premise mix by a wide margin, which means every order that arrives this way is worth several that arrive through a marketplace. And when volume does come, the brands with a clean feed will be the only ones the assistants can actually transact with. This is a channel where the build takes longer than the demand takes to arrive, which is the opposite of how most channels are timed.

Discovery in the assistant. Checkout in your app. That is the whole strategy, and the hard part is the feed.

III.

—— PART THREE · CHANNELS

The channel stack, re-priced.

Most of this table has moved since the first edition. One row is new. Two channels no longer need arguing for at all.

The first edition spent its longest chapter making the case that connected television is a performance channel and that digital out-of-home deserves a line in a franchise plan. Both arguments are settled now, so this chapter spends the space differently: less on whether to buy these channels and more on how to buy them without getting a report you cannot defend.

FIGURE 01 *The franchise marketing channel stack.*

CHANNEL	PRIMARY VALUE	COST RANGE	MINIMUM BRAND-LEVEL INVEST	MEASUREMENT MATURITY
<i>Connected TV</i>	Addressable reach with footfall measurement	\$28–\$51 CPM	\$5.6K / month, and see the note below	Mature
<i>Programmatic DOOH</i>	Geo-targeted presence with dynamic creative	\$4.50–\$9.00 CPM	\$3.4K / market	Directional
<i>Streaming audio</i>	The commute decision, targetable by daypart	\$17–\$28 CPM	\$3.4K / month	Mature
<i>Host-read podcasts</i>	Highest-trust format in the modern stack	\$28–\$56 CPM	\$11.3K / pilot	Directional
<i>Aggregator sponsored placement</i>	Defensive share inside off-premise	CPC model	\$5.6K / month	Native to platform
<i>Aggregator offsite extension</i>	Reaching marketplace audiences on the open web, measured in a clean room	\$12–\$22 CPM	\$11.3K / pilot	Emerging, and see Chapter IV
<i>Retail media (grocery and convenience)</i>	Targeted shopper reach, co-branded opportunities	\$9–\$17 CPM	\$5.6K / pilot	Maturing
<i>Direct mail</i>	Underperforming-store rescue, trade-area lift	\$0.45–\$1.01 per piece	Per-store basis	Mature, on offer codes
<i>Opted-in SMS</i>	Daypart promotion, weather and event triggers	\$0.011–\$0.034 per send	List-build phase first	Mature
<i>Mobile rewards apps</i>	Acquisition with clean attribution	CPA model	Offer underwriting	Native to platform

Cost and investment ranges reflect typical media costs in restaurant categories at the time of writing. Measurement maturity reflects the channel's ability to support footfall and incrementality measurement at modest investment levels. The offsite extension row is new to this edition and its range is directional: no first-edition baseline exists, and pricing varies widely by platform and clean-room arrangement.

Connected television, and how to buy it now

Buy footfall measurement as a purchase criterion, written into the plan before launch, rather than as a report you receive afterwards. That is the whole change from the first edition. The platforms will all tell you they measure store visits. What varies enormously is whether they will commit to a design that reaches statistical significance inside your actual footprint, and that depends far more on how geographically concentrated your units are than on how much you spend.

A brand with sixty units across four contiguous states can get a defensible read at a modest budget. A brand with two hundred units spread across forty-seven markets often cannot, at any budget it would sensibly commit, without narrowing to a test region first. We wrote a dollar floor in the first edition and we were wrong to. Ask instead for the number of exposed households per trade area the design requires, then check whether your footprint can supply it.

Digital out-of-home

Still the most underused channel in restaurant marketing, and the argument from the first edition stands without amendment: programmatic activation and dynamic creative have made it a different product from the billboard buys of a decade ago. A rainy day in Seattle pushes a soup creative. A ninety-degree afternoon near a high school pushes a milkshake thirty minutes before dismissal. Measurement remains directional rather than precise, and the channel pairs unusually well with CTV where you want brand presence and frequency at moderate budget.

Streaming audio and podcasts

Unchanged, and still underweighted. A customer in a car fifteen to forty-five minutes before lunch or dinner is a customer making a meal decision. Host-read podcast advertising remains the highest-trust format available, and the right targeting is by audience and trust rather than by topic: a drive-time news podcast in a city where you have units, read by a local anchor, will outperform a national food podcast by a measurable margin. Test in one or two markets before scaling.

The channel we still recommend against scaling

Paid placement on content-recommendation networks. We tested it across multiple operators before the first edition, we said pass, and nothing since has changed our view. The audience quality is poor, the format rewards clickbait, and the brand-safety profile is inconsistent. It works for direct-response e-commerce. It does not move same-store sales. Pass.

IV.

—— PART FOUR · OFF-PREMISE

Your aggregator is now a media network.

The first edition told you to hold share inside the app while you built the bridge off-platform. That was right, and it is now half the job.

On 4 June 2026, DoorDash consolidated DoorDash, Wolt, and Deliveroo into a single commerce media platform spanning more than four hundred thousand advertisers, and shipped four things at once: a premium homepage display format, offsite audience extension through Symbiosys, a LiveRamp clean room for measurement, and auto-bidding against minimum return-on-ad-spend targets. Read that as a company telling the market it is now in the media business.

The consequence for you is specific. The platform that takes a share of every order it delivers can now reach the customers who placed those orders on the open web, using purchase data you generated, on behalf of whoever is buying. Defending your position inside the app was the whole job when the app was a marketplace. It is not the whole job now that the app is also a demand-side platform.

The platform taking a share of your order can now reach your customer on the open web for a competitor. That is a contract conversation, not a media one.

Three things to do about it

- 01** Keep the defensive spend, and cap it. Sponsored placement inside the marketplace is the price of being competitive in the channel, not a growth lever. Budget it as rent, set the cap deliberately, and review it quarterly against your direct-order share rather than against its own return-on-ad-spend figure, which will always look good because it is measuring the last click before an order that was already going to happen.

- 02** Get into the clean room, on symmetrical terms. You should be able to measure the incrementality of your own sponsored placement, and you should be able to learn whether your customers are being reached off-platform. Access is available. It is rarely offered unprompted, and it is much easier to obtain at renewal than mid-term.

- 03** Renegotiate the agreement with a marketing person in the room. Data reciprocity, location-level reporting granularity, and permitted offsite use of your transaction data are now commercial terms with marketing consequences. Most brands renew marketplace agreements as an operations or finance exercise, and this is the finding in this chapter that is worth acting on this quarter.

The bridge off-platform still matters more

None of the above changes the underlying position, which is that a marketplace order is a customer you do not own: no email address, no order history, no path into loyalty, and twenty to thirty percent of revenue gone before it reaches you. Everything in this chapter is damage control on a channel you should be working to depend on less. The two things that actually reduce the dependency are a direct ordering experience worth returning to and paid media built to drive direct conversion rather than reach. Chapter II is the newest version of that same argument.

V.

— PART FIVE · AUTOMATION

AI that moves the number.

Most of what is written about AI in advertising is unhelpful to a multi-location operator. The applications that move same-store sales still fall into four categories. The ranking has changed.

The first edition put hyperlocal creative generation at the top of this list because, at the time, being able to produce hundreds of market-specific creative variants at a sensible cost was an edge. It is now table stakes. Every competent partner can do it, the tooling is commodity, and the lift it produces is real but no longer differentiating. Fund it, expect it, and stop treating it as a strategy.

The edge moved to operations. Two applications are doing more for the brands we work with than creative generation is.

Local content and review response at two-hundred-location scale

Chapter I described a data problem that nobody owns. The reason nobody owns it is that it has always been too labour-intensive to own: two hundred location records, each needing accurate attributes, current hours, fresh photography, and timely responses to reviews written in a local voice. That is now an automatable workflow with a human reviewer on top, which turns the largest unfunded gap in franchise marketing into a line item you can actually approve. It is the highest-return AI investment available to a franchise brand this year, and it is the least discussed.

Feed hygiene as a marketing function

Your menu, price, and availability feed decides whether assistants recommend you, whether they can transact with you, whether your Performance Max campaigns have anything useful to optimise against, and whether your marketplace listings are accurate. That makes feed quality a marketing dependency, not an IT chore, and it needs an owner on the marketing side of the house with the authority to escalate when it breaks.

Predictive budget allocation

Unchanged from the first edition, and the recommendation still stands: for a two-hundred-unit brand, buy it rather than build it. The capability is available off the shelf and from a handful of agencies, the models learn which markets respond to which channels over time, and the efficiency gains compound across quarters. Building it in-house is a nine-month project to arrive somewhere a vendor will take you in six weeks.

Inventory-aware bidding

Still the most powerful application on the list and still gated on infrastructure rather than on algorithms. Suppressing paid media for a limited-time item in markets where it has sold out, and lowering bids in markets already running at capacity, requires exposing inventory and capacity state through a clean interface. For a brand without a single point-of-sale system across all units, consolidating that reporting is a six-to-nine-month project on its own. Fund it after the infrastructure, not before.

FIGURE 02 *Where AI moves the same-store sales number, and what each application requires.*

APPLICATION	WHERE IT MOVES THE NUMBER	DATA PREREQUISITE	RECOMMENDATION
<i>Local content and review operations</i>	Directly improves assistant and local search visibility per unit	Location record ownership, review platform access, a human reviewer	Fund now, highest return available
<i>Feed hygiene and menu data</i>	Gates assistant recommendation, agentic ordering, and platform bidding quality	Live item, price, and availability data resolved per location	Fund now, and give it a marketing owner
<i>Hyperlocal creative generation</i>	Lift on click-through and attributed footfall in dynamic-creative channels	A master creative system, plus market and weather feeds	Table stakes, expect it
<i>Predictive budget allocation</i>	Compounding efficiency gains across quarters	Weekly store-level sales and marketing spend feed	Buy, do not build
<i>Inventory-aware bidding</i>	Removes wasted spend, protects margin on limited-time offers	Unified point-of-sale and capacity feed through an interface	Fund after the infrastructure
<i>Generative creative production</i>	Production cost, and same-store sales only alongside the above	Brand guardrails and a review workflow	Use with discipline
<i>Platform AI bidding</i>	Real when fed clean signal, punishing when fed thin data	Server-side conversions and enhanced conversions	Audit the data first

This edition publishes the figure without lift ranges. The first edition's ranges could not be re-supported against our current measurement work, and directional honesty is the whole reason people quote this table.

The thin-data problem, sharpened

One uncomfortable truth ran through the first edition's version of this chapter and it has aged better than anything else we wrote. AI systems are only as good as the data they are given, and most franchise restaurant brands have data that is too thin, too fragmented, or too dirty to support the applications they want to fund. Customer identity is split across the website, the ordering app, the loyalty program, and the marketplaces, with no reliable join key. Point-of-sale data reaches corporate weekly rather than continuously. Franchisee co-op spending runs outside corporate visibility. Online-to-offline attribution rests on a conversion event that does not fire half the time.

What changed is the cost of that condition. A year ago, thin data meant your bidding was less efficient than it could have been. Now, with pay-per-lead automation, ads inside AI answer surfaces, and minimum-return auto-bidding on the marketplaces, a brand that cannot supply clean signal is increasingly excluded from the inventory rather than merely charged more for it. *Thin data used to cost you efficiency. It now costs you access.*

This is not a marketing problem. It is a data infrastructure problem, and it is the reason the next chapter is the most important one in this guide.

The CMO who funds AI on top of a weak data foundation will produce a system worse than the processes it replaced. Fund the infrastructure first.

— PART SIX · MEASUREMENT

Measurement that survives the board meeting.

VI.

The fix is not more data. It is one source of truth that connects marketing spend to same-store sales lift at the unit level, with a methodology defensible to every audience around the table.

The team running paid media has metrics that work for their daily decisions. The CFO has metrics that work for the quarterly profit and loss. The franchise advisory council has metrics that work for their local concerns. None of those sets reconcile cleanly, which produces a measurement conversation in which everyone is right and nothing gets decided.

The north-star metric

Every franchise marketing program needs a single primary metric, and we recommend the same one we recommended in the first edition: incremental same-store sales lift per dollar of marketing spend, measured at the unit level, attributed at the channel level, and rolled up to the brand. It solves three problems at once. It is the metric the CFO already cares about, which means it survives the board meeting. It is the metric the franchisee cares about, which means it survives the council. And it is the metric the marketing team can directly influence, which means it drives the right behaviour.

Building it takes three feeds: a clean unit-level sales feed from the point of sale, a clean unit-level marketing spend allocation including the often-invisible co-op spending franchisees are doing locally, and an incrementality methodology (geographic holdout testing for channels without native incrementality measurement, platform-native reporting where it exists). The methodology takes three to six months to build properly and it is the single highest-value investment a franchise CMO can make in this category.

Price, mix, and traffic: the decomposition your CFO is about to ask for

This section is new, and it is the one that connects to the premise of this edition.

A same-store sales number is three things wearing one coat. Sales equal transactions multiplied by average check, and average check moves for two different reasons: because you raised prices, and because customers bought a different mix of items. Marketing owns transactions. Pricing owns price. Menu strategy and operations own mix. When a brand reports a two percent same-store sales increase without separating those components, marketing gets credit in a good year and blame in a bad one, and neither is earned.

Report it decomposed and the conversation changes completely. You can say that marketing delivered a traffic increase inside a quarter where pricing carried the headline, or admit that it did not, and either way you are now discussing something specific. In an environment where nominal growth is running at four percent and real growth near one, this is not a reporting refinement. It is the difference between a marketing budget that gets defended and one that gets trimmed.



Schematic. Marketing is accountable for the first band only.

FIGURE 04 *Decomposing a same-store sales number.*

COMPONENT	WHAT MOVES IT	WHO OWNS IT	WHAT MARKETING SHOULD REPORT
<i>Transactions</i>	Traffic, frequency, and new customer acquisition	Marketing	Incremental transactions per dollar of spend, by channel and market
<i>Price</i>	Menu price increases and promotional depth	Pricing and finance	Nothing, but state its contribution so it is not credited to marketing
<i>Mix</i>	Item selection, attachment, upsell, and daypart shift	Menu strategy and operations	Attachment and average items per order where marketing drove the offer

Three components, one number. Marketing is accountable for the first column and should be reporting against it specifically.

Measuring demand that leaves no click path

The demand described in Chapters I and II arrives without a pixel. An assistant recommends you, the customer opens your app directly, and nothing in any channel report will ever connect those two events. We would rather say plainly that this measurement is immature than pretend otherwise, since pretending otherwise is exactly what this guide criticises vendors for doing.

What works today is a three-part read. The prompt panel from Chapter I is your leading indicator, run monthly, reported as share of recommendations by market. Direct and branded search volume is your lagging indicator, and it is the series that will move first if assistant visibility improves. And geographic holdout design remains the only method that produces a number capable of surviving scrutiny, which means the markets where you improve listings and feeds should be chosen so that comparable markets are left alone.

The franchisee-level dashboard

Most CMOs report marketing performance at the brand level. Franchisees see a national number that may or may not match what is happening in their stores, and the mismatch erodes trust over time. A franchisee-level dashboard fixes it by showing, for each unit, the marketing spend allocated to that unit's trade area, the measured footfall and same-store sales lift, and a comparison against what that unit's own history would predict.

The build is not technically difficult. It is a reporting layer on top of data the marketing team already holds. The hard part is political, because the dashboard exposes where marketing spend is and is not working at unit granularity, which creates pressure to defend allocation decisions in detail. Most marketing teams resist building it for exactly that reason. The CMOs who build it anyway find that franchisee trust improves measurably and that ad-fund debates get shorter and easier to win.

IN PRACTICE

A 1,200-unit QSR brand saw franchise advisory council ad-fund vote participation move from 41 percent to 78 percent within six months of a \$180,000 dashboard pilot. The CMO described it as the most useful money the marketing team spent that year.

FIGURE 03 *One metric, five audiences: how the same data answers each question.*

AUDIENCE	WHAT THEY CARE ABOUT	WHAT TO SHOW THEM	CADENCE
<i>CEO and board</i>	Total same-store sales growth, marketing return, defence of the ad fund	Incremental lift per dollar of marketing spend, rolled up to brand level, with the price and traffic decomposition	Quarterly
<i>CFO</i>	Margin, cash conversion, capital allocation	Incremental lift per dollar plus contribution margin by channel	Monthly
<i>Franchise advisory council</i>	Where their stores' ad-fund dollars went and what they produced	Unit-level dashboard with trade-area spend, measured footfall and lift, and market-level share of assistant recommendations	Monthly
<i>Marketing team</i>	Channel and campaign performance for tactical decisions	Channel-level incrementality, attributed conversions, click-through, cost per thousand, footfall	Weekly
<i>Operations</i>	Traffic surges, inventory implications, daypart shifts	Forecast footfall lift by daypart and market, two weeks forward	Weekly

Every audience is looking at the same underlying data. The view changes. The source of truth does not.

Server-side tracking and consented data

A measurement stack built on browser-side pixels will fail a little more every year. The platforms have already made the decision to move server-side, so you are only choosing whether to migrate on your own schedule or have it forced on you. Implemented properly through a tag manager's server-side container, it improves conversion attribution accuracy by capturing events browser pixels miss, improves site performance by removing tag weight, and gives you control over what data goes to which platform. Budget three to four months, and expect the benefits to compound from the day it goes live.

VII.

—— PART SEVEN · FRANCHISEES

The national- to-local handoff.

A great national program produces national impressions. Local execution turns them into store visits. The handoff is where most franchise brands lose twenty to thirty percent of their program's potential value.

Marketing development funds, co-op programs, and brand-fund contributions all express the same underlying tension. The brand wants consistent execution and wants funds spent on activities the brand controls. The franchisee wants flexibility and wants funds spent on activities they control. *Neither side is wrong.*

The brands that resolve it well share three habits. They publish a clear allowed-use list, so a franchisee knows before spending which activities are reimbursable, at what rate, and under what creative-approval process, since ambiguity is the single biggest source of friction in a typical co-op program. They provide turnkey local creative that a franchisee pulls from a portal and customises within tightly bounded parameters. And they report on co-op performance with the same rigour as national, which is what earns them permission to be more prescriptive about future allowed uses.

The local digital playbook, and the migration you did not ask for

The strongest program in this category is a brand-managed local digital playbook: the franchisee opts in, pays a contribution, and the brand runs Google local advertising against their trade area at scale. That recommendation stands, and the mechanics underneath it changed in August 2026.

Google began migrating Local Services Ads out of their own console and into Google Ads as a pay-per-lead Performance Max campaign type, in phases running through 2027. The product works the same way (ads still appear only on Search and Maps, campaigns remain keywordless and driven by Google Business Profile data, and you still pay for valid leads rather than clicks), but three practical things change. The old dashboard stops being your management home. Historical campaign-level performance reporting does not migrate, so export it before your account moves. And because the campaigns now sync from Google Business Profile in real time, the listing hygiene described in Chapter I becomes a paid media dependency as well as an organic one. The same franchisee-maintained record that decides whether an assistant recommends you now also decides what your paid local ads say.

Add to that the broader direction of the platform, where search ads are now eligible inside AI answer surfaces and matching is increasingly handled by automation rather than by keyword. The practical instruction is unchanged and more urgent: the quality of the signal you send determines what your spend can buy.

Two more programs that work

The geofenced grand-opening package. A brand-managed bundle of geofenced display, digital out-of-home, and connected television around new unit openings, deployed automatically when the brand learns of an opening through its operations system. Predictable trigger, standardised package, real efficiency.

The weather and events response system. Pre-built creative responses to predictable events (weather, sporting fixtures, school schedules) that deploy automatically in markets where conditions are met. Most franchisees lack the bandwidth to run this in real time, which is exactly why the brand should.

One program that doesn't

Nationally managed social media for franchisee accounts does not work at scale. We have watched it tried at a dozen brands. Each location's account has too few followers, too local an audience, and too specific a content need for centralised management to beat benign neglect. Centralising the function is a strategic error, and we would make that call again.

The new row on the matrix, and the question your council will ask

Location-level AI visibility does not appear on anyone's centralize-or-localize matrix yet, and it will be raised at a franchise advisory council meeting in 2027 whether you have prepared for it or not. Our recommendation is to centralize the data layer and the tooling, and to localize the human judgment. The brand should own the listing management platform, the structured data standards, the menu feed, and the measurement, because those are single systems that benefit from one owner. The franchisee should own the review responses and the local photography, because those need a local voice and local knowledge.

Then name it explicitly in the allowed-use list before somebody asks. A franchisee who has paid a vendor for local listing management is going to want to know whether that is reimbursable, whether it conflicts with the brand program, and who is responsible when the two disagree. Having an answer ready is a small piece of work that buys a disproportionate amount of goodwill.

FIGURE 05 *What to centralize. What to localize.*

ACTIVITY	CENTRALIZE	LOCALIZE	WHY
National brand campaigns (CTV, out-of-home, audio)	<i>Yes</i>	No	Scale and consistency are only available at national level
Local digital playbook (Google local, Performance Max)	<i>Yes</i>	<i>Opt-in</i>	The brand manages, the franchisee contributes and benefits
Location records, structured data, and menu feeds	<i>Yes</i>	No	One system, one standard, and it now gates both organic and paid local performance
Review response and local photography	<i>Tooling</i>	<i>Yes</i>	The brand provides the platform, the franchisee supplies the local voice
Grand opening packages	<i>Yes</i>	No	Predictable trigger, standardised package, brand efficiency
Weather and event responses	<i>Yes</i>	No	Real-time triggers, and franchisees lack the capacity to execute
Local social media accounts	No	<i>Yes</i>	Local voice, hyperlocal content, low scale
Direct mail to trade area	<i>Template</i>	<i>Yes</i>	The brand provides creative, the franchisee triggers locally
Community sponsorships and events	No	<i>Yes</i>	Trust is built in person and the franchisee owns the relationship
Loyalty program	<i>Yes</i>	No	Single system, single dataset, brand-managed

Template and opt-in indicate a brand-managed framework with franchisee-level activation. Treat this matrix as the starting point for a conversation with your franchise advisory council, not as a verdict.

VIII.

—— PART EIGHT · A PRACTICAL SEQUENCE

The next 90 days.

A sequence for a CMO who has read this far
and is wondering where to start.

DAY 1-30

Measurement audit

DAY 30-60

Prompt panel

DAY 60-90

Franchisee dashboard

01 to 30 Audit your measurement stack.

Map every conversion action across every channel and find the gaps. Most brands discover that thirty to fifty percent of their reported conversions fire incorrectly or not at all. The audit costs nothing and reveals more about your program than any other single activity, which is why it was first in the first edition and is first again.

30 to 60 Run the prompt panel.

Twenty-five prompts, ten markets, three assistants, scored for appearance and for primary recommendation. It replaces the channel pilot we recommended last year because it is cheaper, faster, and produces a finding you can take to a board meeting rather than a media test you have to wait a quarter to read.

60 to 90 Build the franchisee dashboard, even roughly.

Use it once with your franchise advisory council and watch how the conversation changes. A rough version used once is worth more than a polished version built for nobody.

This sequence builds momentum, produces wins that compound, and changes the political dynamic around the ad fund. It also sets you up to fund the larger investments (first-party data infrastructure, feed and listing operations, server-side tracking) with credibility you have already earned.

The reallocation page

The first edition told you what to start and never told you what to stop, which is the part of this that actually requires courage. Three to five percent of a brand-level media budget is enough to fund everything recommended in this guide. Here is where we would find it.

FIGURE 06 *The 2027 reallocation: where three to five percent comes from, and what it buys.*

WHAT TO STOP OR REDUCE	WHY	WHAT IT FUNDS
<i>National linear television weight in your least concentrated markets</i>	You are paying to reach households outside any trade area, and CTV can buy the same reach addressably	Location records, structured data, and menu feed work
<i>Marketplace sponsored placement above your defensive cap</i>	Return-on-ad-spend on marketplace placement flatters itself by measuring the last click before an order that was already coming	One channel pilot with footfall measurement designed in advance
<i>Retainers on channels with no incrementality read</i>	If nobody can tell you what the channel produced incrementally after a year, that is the answer	The incrementality methodology itself, which then prices every other decision
<i>Print, and any legacy line nobody has defended in two budget cycles</i>	Not because print never works, but because it is almost never the thing being measured	The franchisee dashboard

These are the four places we most often find the money. The specific mix depends on your footprint and your current allocation, and the exercise is worth doing properly before the budget locks.

Three to five percent, found from things nobody has defended in two budget cycles, funds everything in this guide.

● A CORRECTION

What we got wrong.

We published the first edition with opinions in it, which means some of them did not hold. Here are the four worth naming.

01 We wrote a dollar floor for connected television measurement, and we should have written a design requirement. The first edition said footfall measurement becomes workable at roughly five thousand dollars a month at the brand level. That is true for a geographically concentrated brand and false for a brand spread across forty-seven markets, where the same spend will not reach significance in any single trade area. The right question is how many exposed households per trade area the test design needs, and whether your footprint can supply them. Chapter III is rewritten accordingly.

02 We treated first-party data as a prerequisite for AI. It is now a prerequisite for access. We told you that thin data would make your automated bidding less efficient. What has actually happened is that brands unable to supply clean server-side signal are being excluded from inventory and from newer campaign types altogether. We under-stated the stakes.

03 We did not see assistant-led discovery arriving this fast. In the first edition, AI in this category meant bidding, allocation, and creative production. Within a year it also meant distribution, and two of the eight chapters in this edition exist because of a shift we described as a bidding story.

04 We were right about content-recommendation networks, and we would say it again. The one call we are happy to reprint without amendment.

We will keep this page in every edition. A guide with opinions in it should be willing to show its record.

● ABOUT HAPPY COG

Nice people with interesting problems.

Happy Cog has built and run marketing programs for franchise and multi-location restaurant brands since 1999. Our team includes senior strategists, analysts, engineers, and creative directors who have worked across quick service, fast casual, casual dining, and emerging restaurant concepts.

We believe the best marketing work is plainspoken, evidence-driven, and respectful of the operational realities our clients face. We believe in people first.

If the questions in this field guide describe your situation, the next thirty minutes with us is probably worth your calendar.



Start a conversation.

If you want to talk through any of this against your own footprint, write to us. Tell us what your same-store number is doing and where you think the gap is, and we will tell you plainly what we would look at first.

HELLO@HAPPYCOG.COM →