Nuanced

WHITE PAPER | 2026 EDITION

DIGITAL MARKETING STRATEGY FOR URGENT CARE OPERATORS

Advertising cannot cause a cough.

A decision framework for Google Ads in urgent care

How to distinguish clicks from patients, reported conversions from incremental growth, and competitive pressure from profitable opportunity.

THE FOUR OPERATING POSTURES

STAND DOWN	DEFEND	CONTEST	ATTACK

PREPARED BY **Nuanced Technologies Inc.**

Healthcare software · e-commerce · digital marketing

THE CENTRAL THESIS

Advertising redistributes urgent-care demand. It does not create it.

Google Ads can redirect a patient who is already choosing where to seek care. The economic question is not whether an ad received a click or appeared in a conversion report. It is whether the ad changed the patient's destination - and whether that change created enough contribution margin to justify its cost.

THE OWNER'S QUESTION

How many patients did advertising genuinely win that the clinic would not otherwise have received?

01 **Attribution is not incrementality**
A dashboard credits the last visible interaction. It cannot tell whether the patient was already walking toward your door.

02 **Costs compound through the funnel**
Every leak between click, contact, booking, arrival and ad-won patient raises the true acquisition cost.

03 **Brand changes the result**
Established clinics intercept more habitual patients. New clinics are more likely to win genuinely incremental visits.

04 **Economics set the ceiling**
Contribution margin, repeat behavior and available capacity determine what an ad-won patient is worth.

WHAT OWNERS SHOULD DO NEXT

01

Measure the business: margin by service line, payer mix, repeat-visit rate, capacity and normal weekly demand.

02

Separate the traffic: branded from generic, contacts from arrivals, and attributed patients from ad-won patients.

03

Choose a posture: stand down, defend, contest or attack - then test the decision with a scheduled on/off experiment.

FIELD GUIDE

Nine chapters. One operating decision.

Read the argument in sequence, or use the final worksheet to identify the chapters that matter most to your clinic.

STAND DOWN
The auction is uneconomic or there is no meaningful choice to influence.

DEFEND
Protect existing brand demand and choose generic fights selectively.

CONTEST
Build awareness where pressure is limited and acquisition remains efficient.

ATTACK
Compete for high-intent demand when economics and capacity support growth.

All calculations in this guide are illustrative. Replace them with clinic-level data before making a budget decision.

01 · THE ONE TRUE THING

Demand is made by biology, not by budgets

People come to urgent care because of a fever at 9pm, a sliced thumb, a school physical due Friday, a bad flu season. Advertising creates none of that. No campaign has ever caused a pandemic, produced a cough, or driven a sick person around town in a bus.

What advertising *can* do is intercept a person who is already sick and already choosing, and tilt that choice toward one door instead of another. You are not buying patients. You are buying turns at an intersection patients were already driving through.

Hold onto this rule, because it applies at every scale - your whole market, your flu season, and your empty Tuesday afternoon. It will come back in Section 07 wearing different clothes. Two consequences arrive immediately:

- **If you're the only game in town, don't advertise.** There is no choice to tilt. The patient was coming to you anyway, and every ad dollar buys you your own customer. (One nuance below - the "only game" test is broader than you think.)
- **If competitors are advertising, you must respond - or knowingly decline to.** Respond by contesting the same intersections, or by being different in a way that doesn't enrich Google at all: reputation, tenure, speed, a doctor people know by name. Doing nothing while a competitor advertises is a slow leak - but a slow leak can be a rational choice, *provided you've priced it*. The subsection below does the pricing.

The nuance on "only game in town"

Your competition is not just the urgent care across the street. It's the ER, telehealth, the retail clinic in the pharmacy, and the family doctor with same-day slots. If patients in your zip code are choosing between you and *any* of those, a choice exists - and a small, cheap defense of your own name can be worth keeping even in a "monopoly." But it is defense, not growth.

Price the leak before you patch it

A leak has a monthly cost: patients lost × what a patient is worth (Section 07's lifetime math, honestly computed). The patch has a cost too: what defending actually spends. If the patch costs more than the leak - if the cost to win or hold a patient through ads exceeds their total lifetime value - then tolerating the leak is the rational posture. STAND DOWN stops being a monopoly privilege and becomes a pure economic verdict, available in any market.

There's a sharper version of this, worth savoring: auction prices are set by bidders, and bidders can be wrong. If clicks in your market have been bid up past what *any* clinic can recoup per patient, your competitor's advertising isn't a threat - it's them buying patients at a loss, with a dashboard flattering them exactly the way yours would flatter you. Letting an overpaying rival win the auction is letting them lose money. Competitor spend is pressure to be *priced*, never automatically matched.

One caution: the leak math only works with an honest patient value. A clinic that underestimates what a kept patient is worth will conclude every leak is fine, and quietly bleed. Price the leak with Section 07's numbers, not with pessimism.

How to read the numbers in this guide. Wherever we say **"% chance,"** read it as: out of 100 identical situations, how many go this way. "A 40% chance the patient books" means 40 book, 60 don't. Every % chance in this guide is a number you can estimate today and measure later - and half the value is writing down which is which.

DECISION CHECKPOINT · Q1

If every ad you run vanished tomorrow, how many of this month's patients would still have walked in? If your honest answer is "most of them," hold that thought - it becomes the central number of Sections 03 and 04.

The five-minute version of Google Ads

You bid to appear when someone in a chosen area types a chosen phrase. You pay when they click - whether or not they call, book, arrive, or ever become a patient. That's the whole machine. The strategy lives in *which phrases* and *which area*.

The three keyword families

- **Your own name.** "Lakeside Urgent Care." Cheap clicks, extremely high intent - and mostly people who already decided to come to you. Buying these is a fence, not a funnel - and understand who is on the other side of it: by definition, the audience for your own name is your *most loyal* segment. You are renting access to an asset you already own. The fence has exactly one job - blocking a competitor's ad from sitting above your listing - which makes it conditional: search your own name monthly. If no competitor ad appears, take the fence down; you're paying Google to stand between you and people already walking toward your door. Put it back up when someone actually attacks.
- **Generic intent.** "Urgent care near me," "walk-in clinic open now," "STD testing today." This is the real battlefield - the undecided patient, mid-choice. It's also where every competitor is bidding, so it's the most expensive place to be wrong about your numbers.
- **Competitor names.** Tempting, mostly futile. Google scores your ad's relevance to the search; your clinic is, by definition, not relevant to a search for someone else's clinic. You'll pay a punitive price per click to appear inconsistently in front of a person who typed your competitor's name because they wanted your competitor. Expect to be charged an arm and a leg to briefly annoy the loyal. There are edge cases (a competitor just closed, or has a reputation crisis) - but as a standing strategy it burns money.

Location is half the targeting

Urgent care is a drive-time business. A patient with a fever has a radius, not a loyalty. Your real market is roughly the 10-15 minute drive-shed around your door - and ads shown outside it are largely decoration. Any vendor who can't tell you your drive-shed hasn't started working yet.

The exception: the destination doctor

Primary care breaks the radius rule all the time - people who love *their* doctor will drive an hour past a dozen closer options. Occasionally an urgent care earns the same pull: a clinician with a decade of goodwill, a reputation that travels by word of mouth. If that's genuinely you, your drive-shed is wider than the map says, your brand gravity (Section 06) is higher, and more of your patients arrive by habit rather than by search.

But be honest, because this is the single most flattering thing an owner can believe about their own clinic. It's also checkable: pull patient home zip codes from your EMR. If a meaningful share of patients drives *past* a competitor to reach you, you have gravity - widen the targeting and lean on the brand. If nearly everyone comes from the nearest ten minutes, you're a convenience business like the rest of us, and the belief dies on contact with data.

DECISION CHECKPOINT · Q2

Stand in your parking lot and search "urgent care near me" on your phone. Who appears, in what order, and who is paying to be there? Do the same from a competitor's parking lot. That ten-minute exercise is your market map.

DECISION CHECKPOINT · Q3

Is your clinic actually special - the amazing doctor people cross town for? Not "we think so": what do the zip codes in your EMR say? Answer with data, then label it MEASURED.

Not every click is a patient - and costs compound at every layer

A click is not a conversion. A conversion is not a person at the front desk. And a person at the front desk is not necessarily a patient your ad *won*. The funnel leaks at every joint, and the same spend gets more expensive per unit at each one.

The last stage deserves its own vocabulary, because it's where the biggest lie in advertising lives. Of the patients your dashboard attributes to ads, most fall into one of two groups:

- **Intercepted patients** - they were already on their way to *your* clinic. The ad stepped in front of a person mid-journey, took the click, took the credit, and took your money. Nothing changed except your invoice - you paid Google, and then thanked your ads vendor for delivering a patient who was already yours.
- **Ad-won patients** - the ad genuinely changed the destination. These are the only patients advertising actually buys.

Every honest number in this guide is priced against ad-won patients only. The % chance that an "ad patient" was really intercepted is called *incrementality*, and Section 04 shows how to estimate and - more importantly - test it.

FIGURE 1

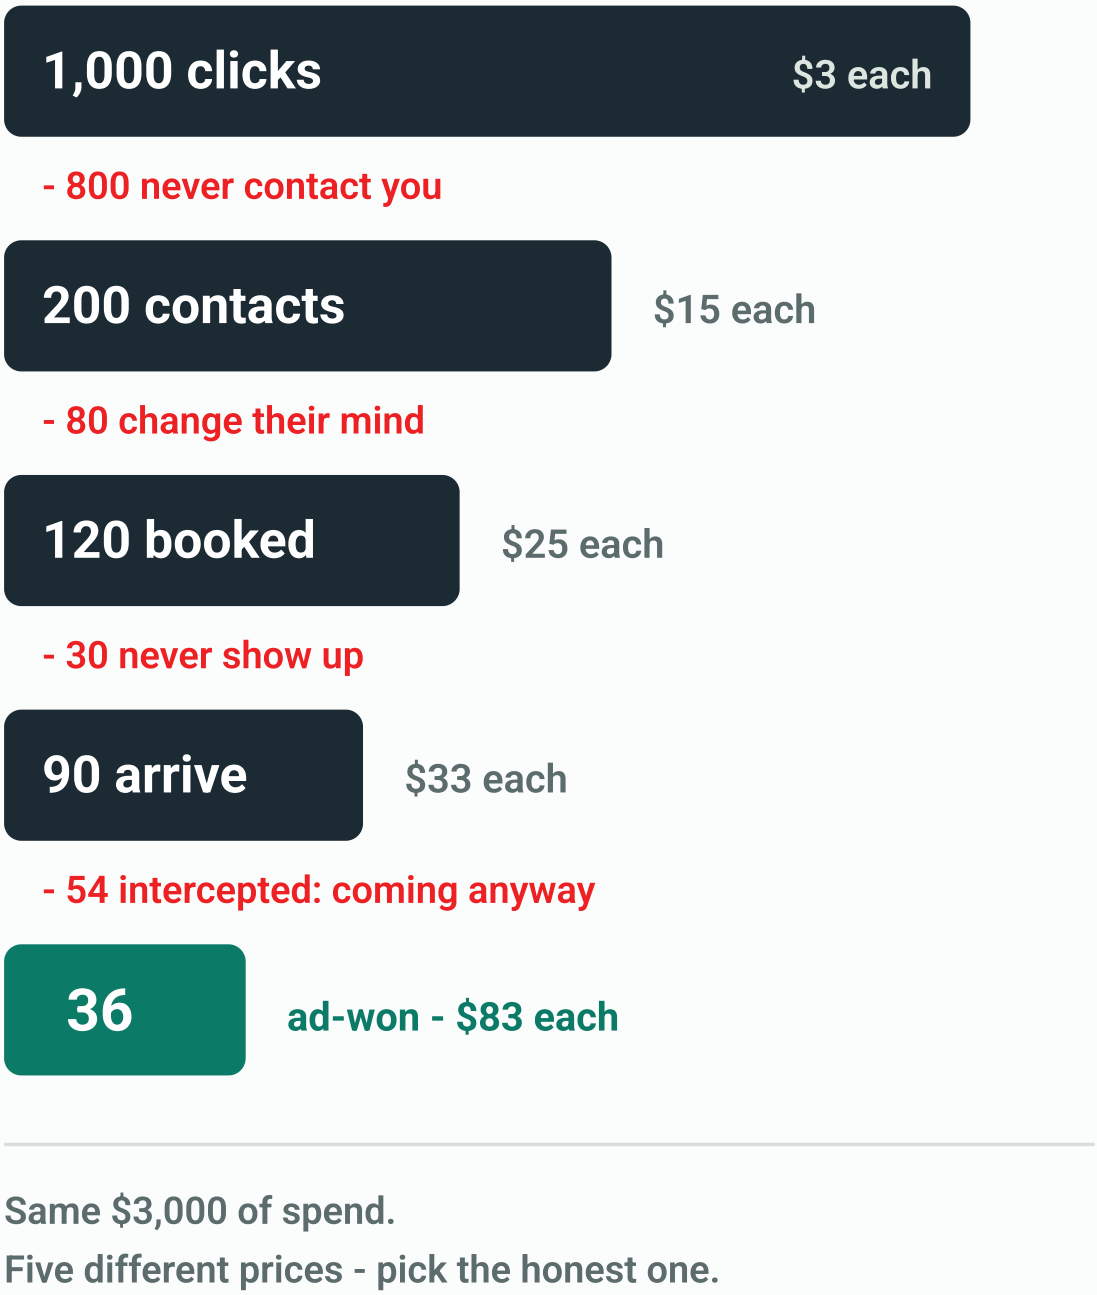

Illustrative numbers. Your leak rates will differ - the point is that they exist at every stage, and each one multiplies your real cost. The last cut is the cruelest: 54 of the 90 arrivals were already yours, and you paid $33 apiece to interrupt their walk to your own door.

The patient's actual path

Before you can price the funnel, trace the human walking through it. Every patient journey passes through the same forks - and ads can only touch some of them:

FIGURE 2

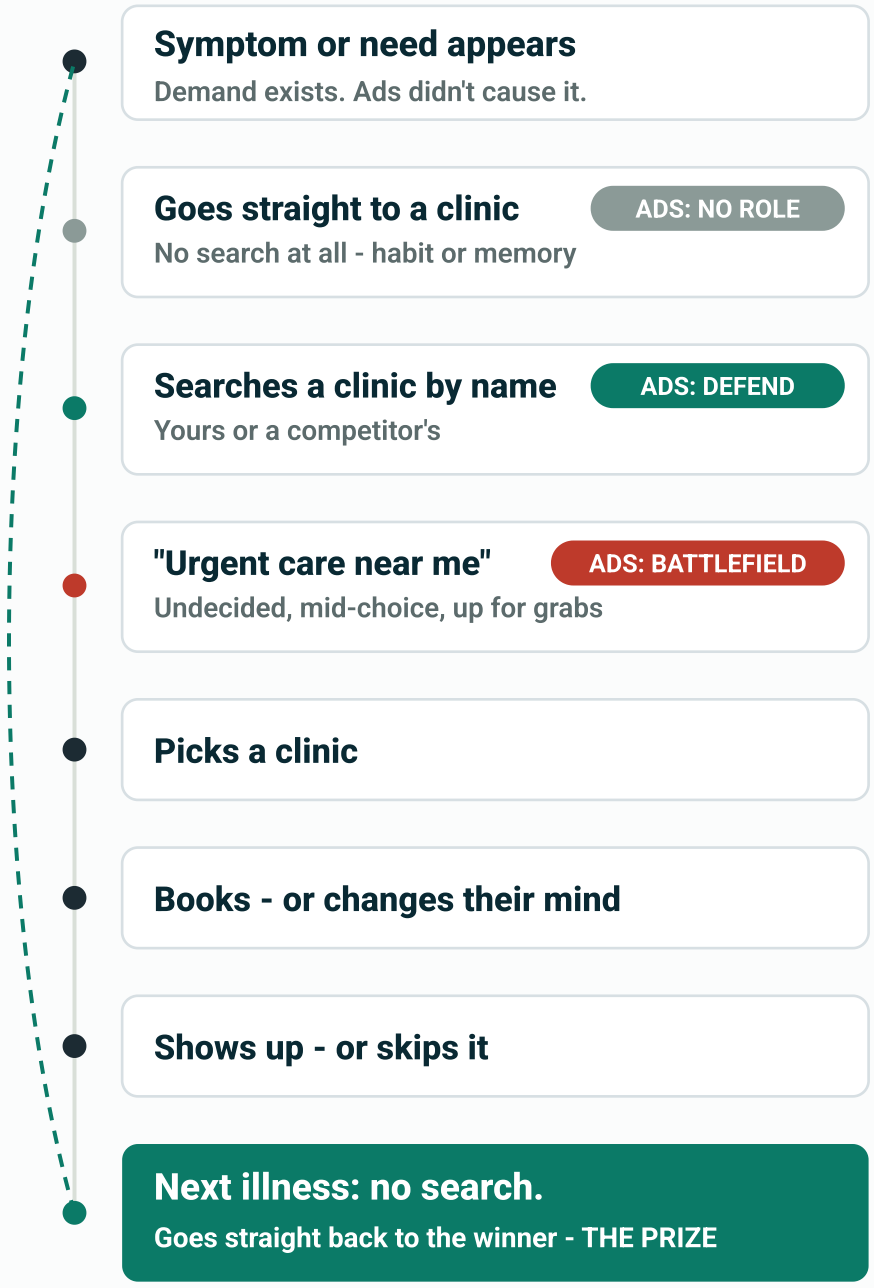

The dashed loop is the whole game - if your business has repeat visits to loop. A patient acquired once and treated well skips the search next time; Section 06 stress-tests how often there is a next time. Ads only operate at two forks in this path - everything else belongs to your front desk, your wait times, and your doctors.

DECISION CHECKPOINT · Q4

For each stage above, do you know your % chance - measured from your own data - or are you guessing? A guess is fine to start. Write it down anyway. A written guess can be corrected; a vibe cannot.

Intercepted vs. ad-won - and how to test which is which

Incrementality is the % chance that an ad-attributed patient was genuinely ad-won rather than intercepted. It is the single most important number in this entire subject, and most vendor reports silently assume it is 100%. It never is.

Meet the two patients your dashboard cannot tell apart. **Maria** has been to your clinic twice. Saturday morning her son spikes a fever; she searches your clinic's name to check the weekend hours, your ad sits above your own listing, and she taps it. **Dev** moved to town last month, wakes up with an ear that won't stop aching, searches "urgent care near me," and your ad beats the competitor's listing to his thumb. Two clicks, two arrivals, two identical rows in the report. But Maria was **intercepted** - the ad was a tollbooth built on the road that already led to your door. Dev was **ad-won** - the ad bent his path away from someone else's.

FIGURE 3

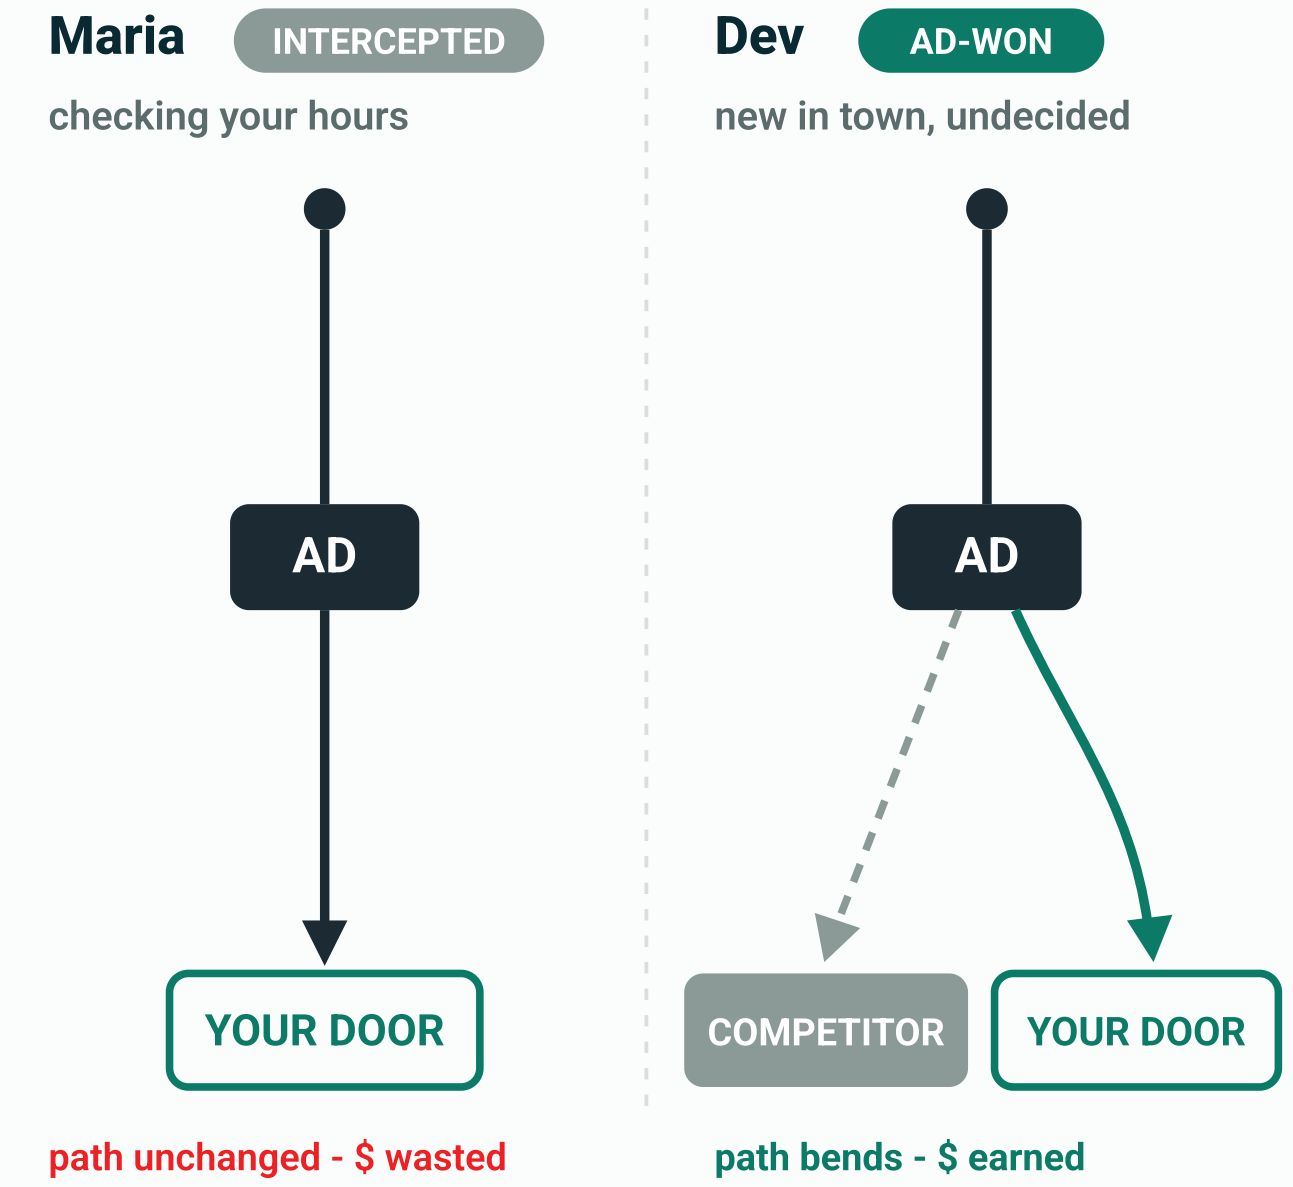

The dashboard records Maria and Dev identically. The ad only earns its cost when the path bends - the dashed line is where Dev was headed without it.

Maria is not rare. Think about who clicks your ad: people who searched your name (already coming), regulars who click the top result out of habit (already coming), neighbors for whom you're the obvious choice anyway (probably coming). The dashboard credits the ad for all of them. The real question is always: *compared to what would have happened anyway?*

```
TRUE cost per ad-won patient
  = spend ÷ (attributed patients × incrementality)

Worked example, one month:
  spend $3,000 · 100 attributed · incrementality 40%
  = $3,000 ÷ (100 × 0.40) = $100 per ad-won patient

What the dashboard says: $3,000 ÷ 100 = $30 ← fiction
```

FIGURE 4

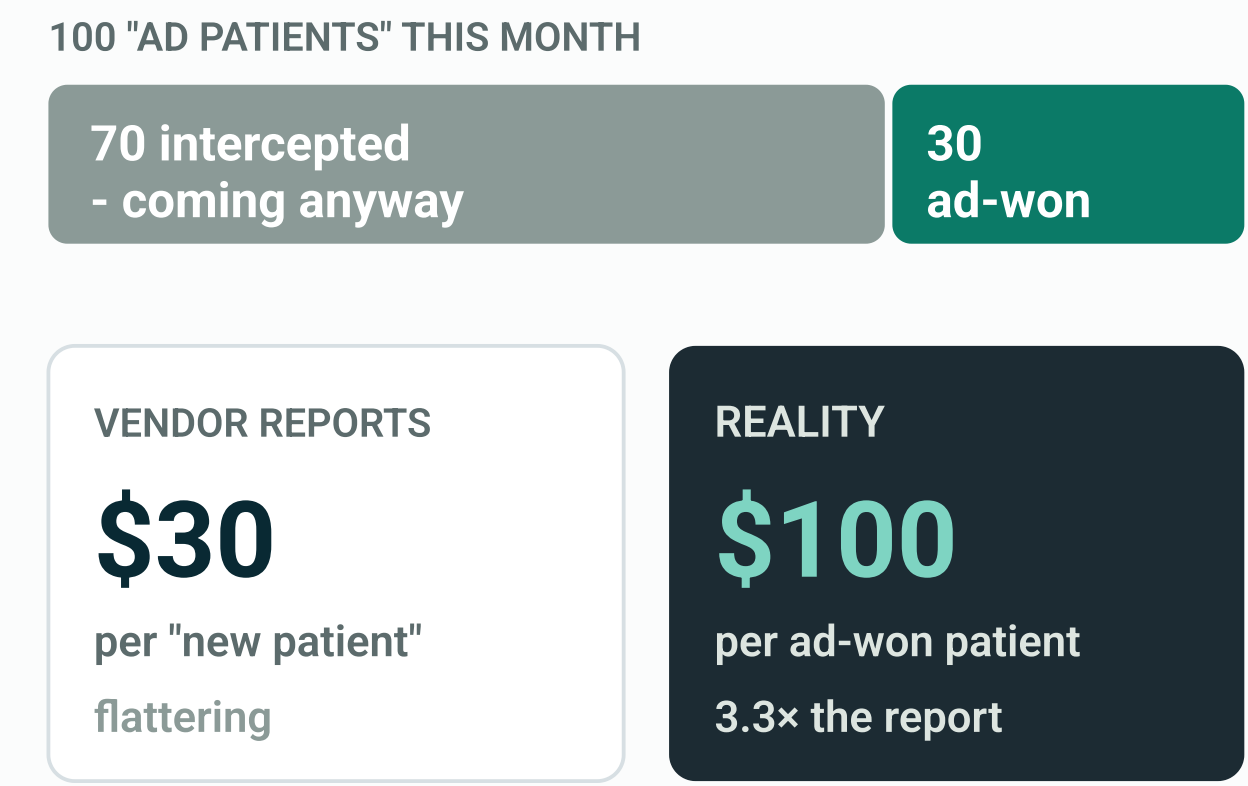

Same month, same spend, same dashboard. The only difference is asking "compared to what?"

Bracketing your number

You can't know incrementality perfectly, but you can bracket it honestly. Push it **down** (more interception) when: your clinic is old and known, your brand-name search volume is high, patients have few alternatives, your branded campaign is a big share of spend. Push it **up** when: you're new, unknown, in a crowded market, with spend concentrated on generic keywords. A decade-old clinic buying its own name might sit at 10-20%. A six-month-old clinic fighting for "urgent care near me" against five competitors might sit at 60-70%.

Don't estimate it. Test it.

The honest definition of incrementality is a counterfactual: two identical clinics in two identical towns, one advertising, one not - the difference in arrivals is what ads actually won. You can't run two towns. But you can run two time windows, and that's the whole test.

The gold standard is embarrassingly simple: **turn the ads off and count.** Run ads-on for a few weeks, ads-off for a few weeks, and measure the drop in patient counts - your front-desk counts, not the dashboard. The dip is what ads were actually winning you. Everything the dashboard claims beyond the dip is interception.

Three honesty rules for the test: run it in a stable season (a pause during a flu spike or a snowstorm measures weather, not ads); compare against a baseline (the same weeks last year); and accept that if the ads *do* work, the test costs some real revenue. Pay it - it's the cheapest audit you will ever buy, and it converts the most important BELIEF in your model into a MEASURED number.

One example of noise vs. signal

Before you trust any test, know the size of your normal wobble. Every clinic's weekly count swings for no reason at all - weather, school calendars, luck. Statisticians call the typical size of that swing the **standard deviation**; call it your wobble. A test can only detect a drop that's clearly bigger than the wobble.

```
Rules of thumb:
  · trust a drop only if it's ≈ 2× your wobble or more
  · averaging n weeks shrinks wobble by √n

CLINIC A - 400 visits/week, wobble ±25
  ads off → 360. Drop of 40 > 2 × 25? No - but close;
  3 weeks of averaging (wobble ±14) makes it clear-cut. Verdict in a month.

CLINIC B - 80 visits/week, wobble ±15
  ads off → 70. Drop of 10 < wobble 15 - an ordinary week.
  Needs ~9 weeks off (wobble shrinks to ±5) to see it. Slow, wide error bars.
```

The uncomfortable truth inside that example: small clinics may lack the volume to measure incrementality precisely. They must test longer, or accept an honest "somewhere between 20% and 60%" - which is still infinitely better than a dashboard's silent 100%. It also explains why vendors don't volunteer these tests: the honest answer is often "we can't tell yet," and that sentence doesn't sell retainers.

Which is exactly why it's the vendor filter. If your ads vendor has never proposed an on/off test - or resists one, or doesn't know what a standard deviation is - change vendors. You provide the patient counts; they design the windows. A vendor unwilling to measure whether their own work does anything is telling you what they think the answer is.

DECISION CHECKPOINT · Q5

Has your vendor ever proposed testing incrementality - or even said the word? Ask for an on/off test design, in writing, this week. The reaction is the audit.

Ads can choose patients - a little, and only in one direction

Advertising has a throttle for *attracting* types of patients, and a much weaker brake for *rejecting* them. Understanding the asymmetry is worth real money.

The throttle: campaigns select by service line

A campaign built around STD testing reaches a private, motivated, often self-pay patient who may spend hundreds of dollars on testing and treatment - and values discretion over price. A campaign around flu shots reaches a price-shopper for a commodity. Same platform, same clinic, wildly different economics. You choose the patient by choosing the promise.

Your campaign list should therefore read like a ranked list of your most profitable *advertisable* service lines, not a list of everything you do: lacerations, school and sports physicals, travel medicine, STD testing, IV fluids. Rank them by margin before your vendor ranks them by guesswork. (Why "advertisable" matters - and why some of your most valuable lines don't qualify - is the subsection after next.)

The flu shot: a war you're not even in

Flu shots deserve their own warning, because the problem isn't just thin margin - it's *who else is selling them*. CVS and Walmart price flu shots at effectively zero, often literally free with insurance, as loss leaders: the shot exists to pull a body into the store, where the profit is. When you bid on "flu shot near me," you're not competing with the clinic across town. You're entering a price war against companies with ten thousand locations whose strategy is to lose money on the product.

Could a loss leader ever make sense for a brand-new clinic trying to seed the habit loop? Only if the loop pays for it - so do the math instead of the vibes:

```
Line value = visit margin + (% chance of return × future margin)

Flu shot: $0 + (10% × $60) = $6 of value
Cost per ad-won flu shot at $4 clicks, 1-in-10 landing: ≈ $40

Verdict: paying $40 to acquire $6. Don't.
```

And the 10% is generous - pharmacy-style transactions build pharmacy-style loyalty, which is to say almost none. If you want a loss leader, pick one the behemoths aren't giving away.

Three kinds of searches - only one of them buys

Here's the trap in "advertise the valuable service lines": value to your P&L and advertisability are different things. What matters is the distance between the search and the purchase, and who's actually doing the buying. Every keyword falls into one of three baskets:

- **Now-intent.** "Urgent care near me," "stitches open now," "STD testing today." The searcher is the buyer, and the buying is happening within the hour. The whole funnel model of Section 03 applies. This is where ad money works.

- **Research now, buy later.** Immigration physicals are the canonical case. Enormously valuable visit - but the vast majority of searches are homework: "what is form I-693," "immigration physical requirements," "how much does it cost." You pay per click to educate a researcher, and when the need finally matures - weeks or months later, after a lawyer's timeline - can you be certain they come back to *your* clinic rather than starting a fresh search? You can't. You bought a click with no claim on the eventual visit. Note the line runs through *keywords*, not service lines: "civil surgeon near me, same day" is now-intent even inside a research-heavy category. Bid on the urgency phrases; starve the homework phrases.
- **B2B relationship sales.** Occupational medicine is the honest example of a line that's ridiculously difficult to advertise. It's valuable, recurring, contract-based revenue - and the buyer is an HR manager or operations director who chooses a clinic through sales calls, pricing sheets, and relationships. The people typing "DOT physical near me" are the *workers*, not the buyers. Do you want to pay per click for local truckers to read about physicals, in the hope their employer independently signs a contract with you? That's not a strategy; a sales call is. Some worker-side searches convert for one-off physicals, so a small campaign can exist - but the line grows by handshake, not by auction.

The weak brake: filtering leaks

- You can exclude phrases like "free clinic" - but the person seeking free care doesn't always type "free." They type "urgent care near me" like everyone else, click your ad, and you paid for the click.
- You cannot target by insurance. There is no setting for "commercial insurance only." If your worry is a payer whose reimbursement doesn't cover your cost to serve, ads cannot screen them out - the click happens before the insurance card does.
- Geography is a crude proxy for demographics, and an imperfect one. Use it, but don't expect surgical precision.

Reimbursement closes the loop

Cost to acquire is only half an equation. The other half is what the visit pays - and that swings enormously by service line and payer. A $100 true cost per ad-won patient is a bargain for a self-pay service line averaging $350 in revenue, and a slow bleed for visits reimbursing $85. You cannot judge the first number without knowing the second.

DECISION CHECKPOINT · Q6

Which three service lines do you actually want more of - and for each: what does the average visit pay, and which of the three baskets does its search traffic fall into? If you can't answer from your billing data today, that query is worth more than any campaign change.

DECISION CHECKPOINT · Q7

What share of your current visits reimburse below your cost to serve? Ads that grow your volume grow that share too, unless the campaigns are steered by service line.

Brand changes the math on both sides

Are you a national chain, or Dr. Jones who has sat in the same clinic for a decade? Is your competitor the other way around? These aren't soft questions - they move click-through rates, incrementality, and the entire defend-vs-attack split of your budget.

Two stories patients tell themselves

When "urgent care near me" shows a national brand next to a local independent, the listings do not perform equally - and which wins depends on your zip code. Some patients read national as clean, standardized, insurance-friendly, open Sundays. Others read it as revolving-door providers and upsells, and prefer the doctor who's known their family for ten years (the destination doctor of Section 02). Neither story is universally true. Your *local* read on which story wins is a genuine input to strategy - hand it to your vendor.

Sway and the loop: the two % chances brand controls

Strip branding down to arithmetic and two numbers fall out. Name them, estimate them, and half the strategy writes itself:

- **Sway factor** - the % chance an ad changes a patient's destination *this* visit. This is what you rent from Google, one auction at a time. It's also what competitor ads try to exert on *your* patients.
- **Loop factor** - the % chance that a patient you've served skips the search entirely next time and comes straight back. This is what you own. Nobody bills you for it.

A strong brand moves both numbers in your favor at once: your patients are harder to sway away (competitor ads bounce off), and more of them never re-enter the funnel at all (you stop paying tolls). Which gives you the sentence that should hang over this whole section: **advertising is a toll you pay every visit until your brand is strong enough to stop paying it.**

This is also where lifetime value really lives. A patient's three-year value is mostly a function of the loop factor - and the loop factor is mostly a function of brand and experience: wait times, the front desk, whether the doctor remembered them. So when you ask "was that $100 ad-won patient worth it?", you're really asking: is this a *permanent win* or a *rented visit*? A clinic with a 60% loop factor bought a stream of visits; a clinic with a 15% loop factor bought one transaction and gets to bid for the same human again next year. All of it is measurable: repeat rate in your EMR, share of arrivals with no search, monthly searches for your name and whether the count grows.

The loop is a hypothesis, not a law

One honest caveat before you bank on the loop: it assumes urgent care is a repeat business, and for part of your patient base it isn't. A family with three young kids is a genuine loop asset - strep tests, ear infections, sports physicals, several visits a year. A healthy 28-year-old getting stitches may not need urgent care again for three years - and across a gap that long, loyalty decays. They move, they forget, they simply search again as if you'd never met. For that segment the loop factor is effectively zero, and an ad buys a transaction, full stop.

So the loop factor is really the % chance a patient returns *before the memory fades* - episode frequency matters as much as satisfaction. One EMR query settles which business you're in: the distribution of visits per unique patient over 24 months. If your median patient visits once, you are, for lifetime-value purposes, a transaction business. That's survivable - but it changes every downstream number: Gate 2 in Section 07 shrinks toward single-visit margin, acquisition ceilings drop, and the whole posture should cool.

It also gives brand its second job. Advertising pays the toll per episode; **brand is the bridge between episodes** - the thing that survives a two-year silence so the loop can still close across it. Signage, community presence, a doctor's name people repeat at dinner: their function in this model is to stretch the memory horizon.

What your position implies

- **Local independent vs. a chain:** the chain bids on generic terms with a budget you can't outspend. Your play is a cheap, airtight fence around your own name, plus differentiation the chain cannot copy - tenure, continuity, the doctor's actual name. Head-on generic bidding wars against a national budget are the losing move.
- **Chain or franchise location:** your name already does work - people search it directly. That's an asset, and a trap: branded campaigns will look spectacular in reports precisely because those patients were intercepted, not won. Incrementality discounting cuts hardest on the strongest brands.
- **Tenure is an incrementality proxy.** A decade-old clinic has a large habitual base - ads mostly intercept it. A new clinic has almost no base - most ad-driven patients are genuinely ad-won. The same $3,000 buys mostly fiction at one clinic and mostly patients at the other.
- **Your competitor's brand matters as much as yours.** Against another solo doc, nobody searches either name much - generics are the battlefield. Against a national name, their branded searches are high-volume but Google will price interception punitively (Section 02). Don't fund that fight.

FIGURE 5

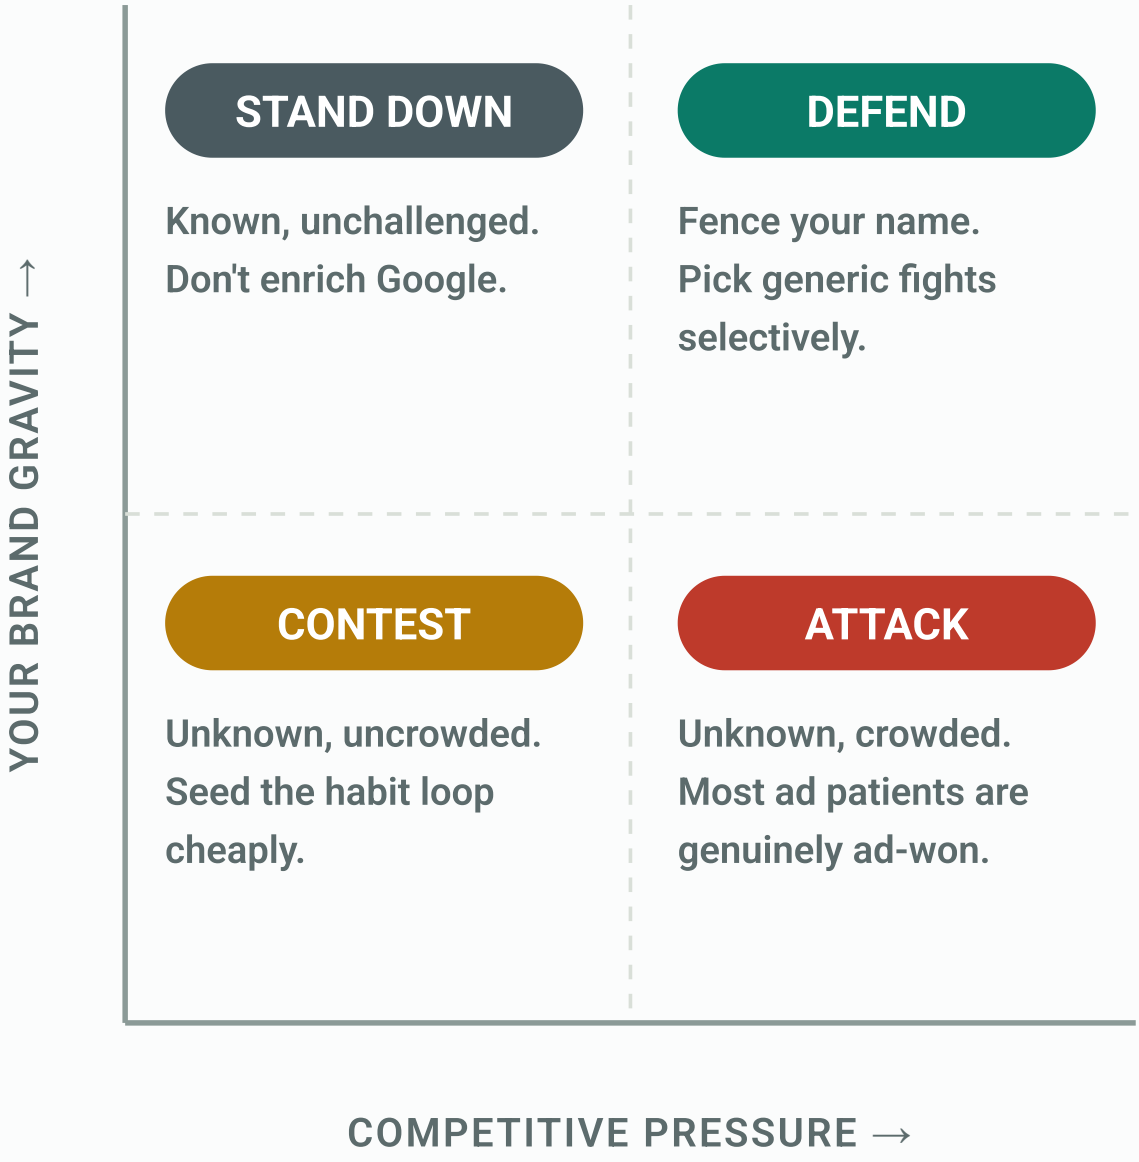

Where you sit on this map is set by two things you already know: how many rivals are advertising, and how much pull your name has. The postures return in Section 09, worked through on a real chart.

Your worldview is a hypothesis

"People here prefer a local doctor" is a belief until tested. Some of it is checkable: review sentiment, repeat-visit rates in your EMR, monthly search volume for your name versus your competitor's (your vendor can pull this in minutes). Tell your vendor your worldview - but mark each parameter BELIEF or MEASURED, and revisit the beliefs quarterly. A model with labeled guesses beats no model; an unlabeled guess masquerading as data beats nothing at ruining a budget.

CHAPTER CHECKPOINTS

DECISION CHECKPOINT · Q8

Does your name pull patients toward you - or is it just a name on a sign? Rough test: how many people per month search for your clinic by name, and is that number growing?

DECISION CHECKPOINT · Q9

When you win a patient, do you keep them - and how often is there anything to keep? What's the % chance a first-timer returns within 12 months, and what's the median number of visits per unique patient over 24 months - measured from your EMR, not remembered from your gut? Those two numbers are your loop factor and its reality check; together they decide whether ads buy assets or transactions.

DECISION CHECKPOINT · Q10

Which of your opinions about "local vs. national" in your market is measured, and which is believed? Label them before your vendor builds a strategy on top of them.

Three gates every ad dollar must pass

Gate 1 - Contribution margin, by service line

Not revenue: *margin*. What a visit pays minus the variable cost of serving it (supplies, labs, per-visit staffing). Averaged across your real payer mix, per major category. This is the ceiling that acquisition cost lives under - and it's why "we got 90 patients from ads" is not yet a sentence with meaning.

Gate 2 - Lifetime value and the habit loop

Section 06 named the loop factor; here's what it's worth in dollars. A patient acquired once and kept skips the search next time - you bought all their future visits with one click:

```
Patient value ≈ margin now + (loop factor × visits/yr × margin × years)

Example: $60 + (50% × 2 × $60 × 3) = $240, not $60
Same patient at a 15% loop factor: $60 + $54 = $114
```

Your EMR knows your real loop factor - the share of last year's first-time patients who came back within 12 months. Pull it. It can double or quadruple what an ad-won patient is worth, and it's the number your front desk and wait times control, not your vendor. And pair it with Section 06's reality check: the formula silently assumes "visits per year" is a real number. A clinic whose median patient visits once ever has a patient value near single-visit margin - and every acquisition ceiling in this guide drops accordingly.

Gate 3 - Capacity, with a trap inside

Urgent care is a fixed-cost machine: rent, staff, and equipment are paid whether rooms are full or empty. If you're running at capacity, ads are pure waste - you're paying to lengthen a line that already exists, and worsening the waits that feed your reviews. If rooms sit empty, even a modest-margin patient contributes to fixed costs you're paying anyway. So far, so simple.

Now the trap. "Empty rooms" is the strongest honest argument for advertising - *and* the easiest one to misuse. Empty rooms only argue for ads if there is demand at that hour going to somebody. If your 2pm Tuesday lull is empty because nobody in town is seeking care at 2pm on a Tuesday - every clinic's lobby is quiet - then there is nothing to redirect. Ads shown into a dead afternoon buy a handful of clicks from researchers and intercepted regulars at a rising unit cost, and your ad-won count barely moves. This is Section 01's rule again, at a smaller scale: **the thesis is fractal.** Ads redistribute demand and never create it - across your market, across your season, and across your Tuesday afternoon. Before spending into a slow daypart, ask the only question that matters: are patients walking into *someone's* clinic at that hour? If yes, compete for them. If no, no budget fixes an empty town.

Seasonality: bring out the bucket when it rains

Demand follows the calendar - flu season, allergy spring, summer stitches, back-to-school physicals. Since ads only redirect demand that exists, spend should follow the demand curve: heavier into flu season (as long as Gate 3 leaves room - peak season is exactly when you might be full), and scaled back to embers in the summer trough. You cannot catch rain that isn't falling.

The trough is doubly expensive, and the reason is subtle: even if the price per click falls in the slow season, the thin remaining demand skews toward habitual patients - people who already know where they're going. Your % of ad-won arrivals drops, so the cost per *ad-won* patient rises just as the dashboard looks cheapest. A vendor spending the same amount every month has scheduled your budget for their convenience, not your rain.

FIGURE 6

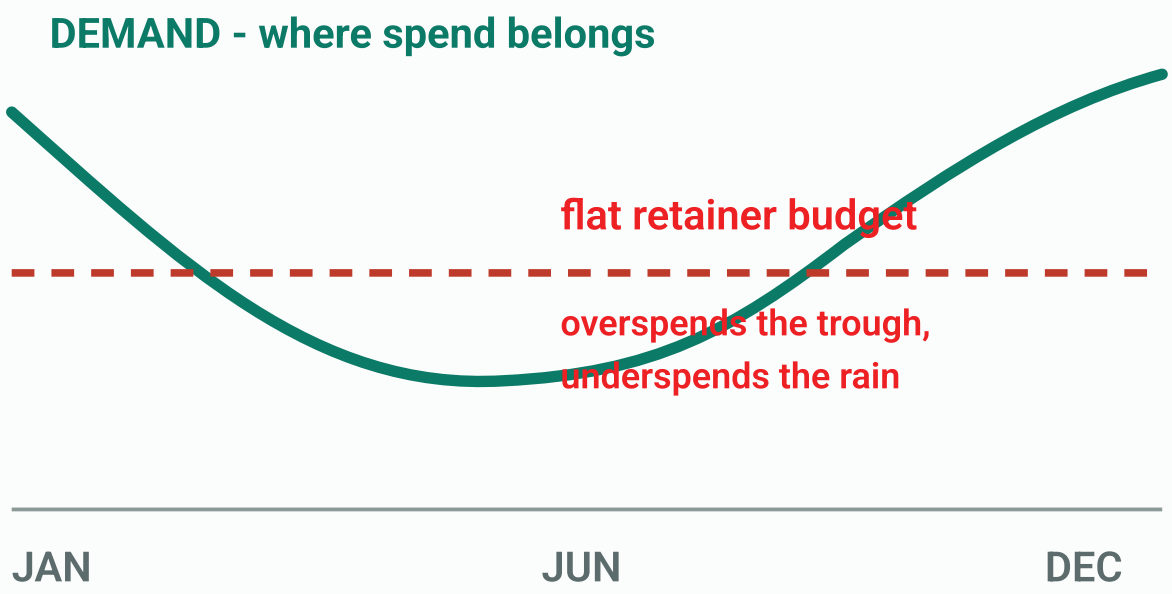

Bring out the bucket when it rains; put it away in the dry months. The red line is a budget scheduled for the vendor's convenience.

DECISION CHECKPOINT · Q11

What is a new patient worth over three years - not one visit? Take your loop factor from Q9 and run the Gate 2 math with your own margin.

DECISION CHECKPOINT · Q12

When are your rooms empty - and during those hours, are patients walking into anyone else's clinic? Empty rooms plus busy competitors is the case for ads. Empty rooms plus an empty town is just an afternoon.

What to hand them, what to demand, what they cannot do

A Google Ads vendor operates a machine. They cannot know your margins, your payer mix, your empty hours, or your neighborhood's feelings about national chains - unless you tell them. The single biggest upgrade to any campaign is the owner showing up with a model.

Hand them (your side of the table)

- Your drive-shed - the real 10-15 minute radius, or the wider one your zip-code data earned in Section 02.
- Service lines ranked by contribution margin, each tagged with its search basket: now-intent, research, or B2B.
- Your payer mix, and which payers you'd rather not grow.
- Capacity by daypart and season - when you want volume, when you're full, and which lulls are dead-town hours no budget can fill.
- Your brand worldview from Section 06, with each parameter labeled *measured* or *belief*.
- Your funnel % chances, your loop factor, and weekly patient counts for incrementality testing.

Demand back (their side of the table)

- **Branded vs. generic split in every report.** One blended number is the oldest trick in the book: cheap branded clicks from intercepted patients, averaged in to flatter everything else. Insist on seeing the two families separately, every month, no exceptions.
- **An on/off incrementality test, designed and scheduled.** Not an incrementality *assumption* - a test. You provide the patient counts; they design the windows around your stable season. If they've never run one, or hedge about "losing momentum," or can't explain what a standard deviation is and why it sets the test's length - that's not a negotiation, that's your cue to switch vendors.
- The search-terms report - the actual phrases real people typed before clicking. Read it monthly; it's where the research-basket and B2B clicks hide in plain sight.
- A negative-keyword list that grows every month. If it's static, nobody is steering.
- Cost per *arrived* patient, not per click or per call - which requires them to accept your front-desk data as the scoreboard.
- Spend that follows your demand curve and your capacity calendar.

What no vendor can do

- Create demand. (Section 01. It bears repeating exactly as many times as a vendor implies otherwise.)
- See inside your EMR. Only your front desk can match an arrival to a click, a payer, and a first-visit flag. Closing that loop is *your* job - a simple "how did you hear about us?" plus a weekly tally is enough to start. No vendor can do it for you, and any vendor claiming complete attribution without your data is describing a guess.
- Know a clicker's insurance before the click.
- Make competitor-name bidding cheap. Google's pricing sees to that.
- Guarantee new patients. They can guarantee clicks. The distance between those two words is Sections 03 and 04.

Red flags, collected

- Never addresses incrementality - no test proposed, no interception discount applied, no curiosity about which arrivals were coming anyway. Remember the incentive: every intercepted patient becomes a conversion in the report, the report justifies the retainer, and the retainer funds more interception. The test threatens the report card, so the demand for one must come from you. This is the master red flag; a vendor uninterested in whether their work does anything will exhibit all the others eventually.
- No deeper interest in your business: never asks about margins, payer mix, capacity, or seasonality. They're running a machine, not a strategy.
- One blended ROAS number, campaigns not broken out.
- "Wins" concentrated in branded campaigns.
- The same spend every month of the year.
- Reluctance to share search-terms reports or the account itself (you should own the account; they should be a manager on it).
- Any promise that sounds like demand creation: "we'll grow your market."
- "Conversions" reported without a definition you've agreed to.

> **DECISION CHECKPOINT · Q13**
>
> For last month: can your front desk tell you which arrivals came from an ad, and how many of those were first-time patients? If not, fix that before touching the budget - every other number in this guide depends on it.

09 · YOUR CHART

The intake form for your own business

You made a patient fill out a form before treating them. Same rule applies to your marketing. Print this page, fill it in honestly - guesses labeled as guesses - and read your posture the way the worked example below reads its own.

1 · URGENT CARE COMPETITORS WITHIN A 15-MINUTE DRIVE

☐ None ☐ 1-2 ☐ 3 or more

2 · ARE THEY ADVERTISING ON GOOGLE? (THE Q2 PARKING-LOT TEST)

☐ No ads showing ☐ Some ☐ Heavily

3 · YEARS THIS CLINIC (OR ITS LEAD CLINICIAN) HAS SERVED THIS COMMUNITY

..... years

4 · YOU ARE

☐ Independent ☐ Regional group ☐ National brand

5 · YOUR MAIN COMPETITOR IS

☐ Nobody, really ☐ Independent ☐ National chain

6 · TYPICAL CAPACITY UTILIZATION

☐ Under 60% ☐ 60-85% ☐ Over 85% - usually full

7 · AVERAGE CONTRIBUTION MARGIN PER VISIT

$ ..... BELIEF / MEASURED - circle one

8 · SHARE OF VISITS THAT ARE COMMERCIAL INSURANCE OR SELF-PAY

..... % BELIEF / MEASURED - circle one

9 · % CHANCE A TYPICAL "AD PATIENT" WAS INTERCEPTED (COMING ANYWAY)

☐ Low, under 30% - we're new ☐ 30-60% ☐ Over 60% - everyone knows us

10 · LOOP FACTOR - % CHANCE A FIRST-TIME PATIENT RETURNS WITHIN 12 MONTHS

..... % BELIEF / MEASURED - circle one

One clinic, charted

Here is the whole model run once, on a made-up clinic, so you can see how a filled chart turns into a posture.

CASE STUDY · MAPLE & 5TH URGENT CARE

Competitors within 15 min: **4, two advertising**

Tenure: **2 years** · Independent · Main rival: **national chain**

Capacity: ≈ **60%** · Margin/visit: **$55** · Commercial/self-pay: **50%**

Intercepted % chance: **low (~25%)** - too new to be a habit yet

Loop factor: **45%** BELIEF - EMR query pending

Reading the chart: a young clinic has almost no habitual base to intercept, so incrementality runs high - around 75% of ad-driven arrivals are genuinely ad-won, and the funnel math prices them near `$3,000 ÷ (100 × 0.75) ≈ $40` each against a Gate-2 patient value of `$55 + (45% × 2 × $55 × 3) ≈ $200`. Rooms are empty during hours when competitors' lobbies aren't, four rivals are already bidding, and half the payer mix reimburses well. Every arrow points the same way:

STAND DOWN	DEFEND	CONTEST	**ATTACK**

Posture: ATTACK - a real budget on now-intent generic keywords, shaped to flu season; a cheap fence on its own name; not one dollar bidding on the chain's name; and an on/off test booked for the first stable month to convert that 25% from belief to measurement.

Now age the same clinic ten years. Tenure 12 years, everyone knows the name, intercepted % chance now *high* - and the identical market reads **DEFEND**: fence the name, contest generics only where margins justify it, and discount every dashboard "win" hard. Same town, same competitors, opposite posture. The chart is the difference.

The numbers to bring to your vendor

Every blank is something covered above - and every blank you can fill makes your vendor measurably better at their job.

- Average revenue per visit, top 3 service lines: ..... / ..... / .....
- Search basket per line (now-intent / research / B2B): .....
- Average contribution margin per visit: .....
- Payer mix (commercial / medicare-medicaid / self-pay): ..... / ..... / .....
- % chance a contact books: ..... `measured / belief?`
- % chance a booked patient shows up: .....
- Loop factor - % chance a first-timer returns within 12 months: .....
- Median visits per unique patient over 24 months: .....
- Incrementality: last on/off test date and result: .....
- Weekly visit count and its normal wobble (standard deviation): ..... ± .....
- Monthly searches for our name vs. main competitor's name: ..... vs .....
- Hours / seasons we want volume - and dead-town hours we don't: .....

Nuanced

ABOUT THE PUBLISHER

Technology advice, grounded in the business.

Nuanced Technologies Inc. provides healthcare software development, e-commerce and digital marketing services for small and mid-sized organizations. Our work connects technical execution to the operating numbers that determine whether an investment actually performs.

Healthcare software

.NET, Blazor and SQL Server application development for healthcare organizations.

Online marketing

Google Ads, social advertising and measurement built around business outcomes.

Business-first advice

A direct, no-BS approach that treats the client relationship as a working partnership.

ABOUT THE FRAMEWORK

A decision model, not a universal prescription.

The examples in this paper are illustrative and are designed to make assumptions visible. Actual advertising performance depends on market conditions, clinical economics, payer mix, capacity, tracking quality and execution. Validate each parameter with clinic-level data and professional advice appropriate to your organization.

Nuanced Technologies Inc. | nuanced.ca | 1-833-NUANCED

THE BOTTOM LINE

Google sells clicks. Your job is to know what a click is worth before you buy a thousand of them.

The clinic that knows its numbers buys patients. The clinic that does not buys hope - at auction, against neighbors doing the same.

Healthcare software, e-commerce and digital marketing

42 Carden Street · Guelph, Ontario · N1H 3A2

nuanced.ca

1-833-NUANCED · 1-833-682-6233