




— A WORKING GUIDE FOR PAID MEDIA IN AI ASSISTANTS

TIEBREAK

Only pay when the AI can't decide.

How we decide where ChatGPT, Google AI Mode and Perplexity ad budget goes, and where it doesn't. The original framework, the reasoning behind it, and the operating playbook we now run on client accounts.

THE METRIC

Hedge rate

RUNS PER QUESTION

20 minimum

THE BUDGET RULE

60/25/15/0

Hassan Tahir

Founder & Lead PPC Strategist, Global Performance Propel Company
Google Partner & Meta Business Partner

globalppc.org

A short note, then the contents

I wrote the first version of this framework in a week when three clients asked me the same question: "Should we be in ChatGPT ads?" I didn't have a good answer, and the answer I eventually landed on wasn't "yes" or "no". It was "only on some questions, and here's how to find them." This document is that answer, written out properly, with the parts I've added since we started running it for real.

Hassan Tahir, September 2026

If you read nothing else. An AI assistant writes its recommendation before it shows you an ad, so an ad only changes a purchase when the assistant hasn't made up its mind. TIEBREAK measures that indecision (we call it the hedge rate) and then scores each buyer question by how winnable it actually is. Budget goes to the questions the AI can't settle, where you are already on its shortlist. Nowhere else.

PART 1 · THE FRAMEWORK

03 Why buying intent stops working inside an assistant

04 The three states a question can be in

05 Hedge rate, the settledness map and the fork rule

06 Running the audit, and when to leave the framework alone

07 What we're basing this on

PART 2 · THE OPERATING PLAYBOOK

08 The eight-step loop

09 Why one signal wasn't enough

10 Scoring a question, and the five tiers

11 The decision flow we hand to buyers

12 Mapping the fork tree

13 What actually tips a tie, and the Tiebreaker Card

14 Splitting the budget, and arbitrage between assistants

15 Proving it worked when the dashboard can't

16 Category notes and the mistakes we keep seeing

17 Where it sits next to Google, Meta and GEO · the monthly rhythm

18 Next steps

Part 1 is the argument. Read it once. Part 2 is the manual; keep it open while you run the audit. The charts in Part 2 use example numbers to show the method, and every threshold is where we start, not where you have to end up.

— PART 1

Why buying intent stops working inside an assistant

For ten years the job was to find the buyer with intent and get in front of them. In an answer engine the buyer has intent and it doesn't matter, because someone else has already answered.

Take a search for "espresso machine under £600". On Google that query shows ten links and the buyer picks. Bidding on it is sensible because the decision is still open when your ad loads. Ask the same thing in ChatGPT and something different happens: the model reads the question, weighs what it knows, writes "get the Sage Bambino", and *then* the sponsored card appears underneath. The buyer has already been told what to buy by the thing they trust. Your ad is a stranger disagreeing.

Three things about the channel make this worse, not better.

- ◆ **The order is fixed.** OpenAI has said plainly that ads don't influence answers. The recommendation is written first and the ad is attached after. There is no auction position that puts you ahead of the verdict.
- ◆ **You're buying a context, not a person.** Matching is contextual, reporting is aggregate. You can't out-target a rival with audience data the way you can on Meta. The only lever left is what the model says.
- ◆ **The money is real.** ChatGPT reaches most US generative-AI users, and self-serve entry sits at a \$50,000 minimum. Budget is pouring into a channel whose deciding variable doesn't appear in any dashboard.

WHERE THE DECISION HAPPENS

Search engine

Query

Your ad · one of ten options

Nine other options

Buyer decides here · ad in play

AI assistant

Question

Model decides here · verdict written

Answer: "Buy the X"

Your ad · after the verdict

The ad is visible either way. It only wins if the verdict isn't settled yet.

So the uncomfortable bit: the high-intent commercial questions that everyone is queueing up to buy in ChatGPT are, more often than not, the ones where the model has already decided. That's the leak. The rest of this guide is a way to find the questions where it hasn't.

The whole thing in one line

Only pay when the AI can't decide.

— PART 1

The three states a question can be in

Ask the same buyer question twenty times in fresh sessions and you'll see the assistant land in one of three places. Each one has a single instruction, and the instruction for the first and third is the one most advertisers find hardest to follow.

THREE STATES, THREE INSTRUCTIONS



Settled against you

Don't bid yet · unsettle first



Open

Bid hard · one fact tips it



Settled for you

Spend zero · monitor

STATE	WHAT YOU SEE OVER 20 RUNS	WHAT TO DO
Settled against you	One competitor named on 70% or more of runs.	DON'T BID YET Publish a counter-fact the model can source, get it cited, then re-measure. Bid only once the hedge rate starts climbing.
Open	No brand above 50%, or a shortlist / "it depends" on 40%+ of runs.	BID HARD The model is missing one fact. A card that supplies it completes the answer instead of arguing with it.
Settled for you	You named on 70% or more of runs.	SPEND ZERO You're recommended for free. Re-run monthly; a rising hedge rate means a rival has published something the model likes.

Why "settled against you" is the one people argue with

It's the exact high-intent question you'd love to own, and every instinct from search says go and win it. But an ad placed beneath a confident recommendation for someone else isn't competing on equal terms. Engagement is poor, and the impression quietly teaches the buyer that you're the one the AI didn't pick. Change the facts before you change the budget: publish the head-to-head, the spec they can't match, the price condition under which you win. Get it onto pages the model reads. Then ask again.

And why "settled for you" gets bought anyway


Because it feels like defending your patch. It isn't. The model is defending it for you, for nothing. Your job there is surveillance, not spend.

— PART 1

Hedge rate, the settledness map and the fork rule


The **hedge rate** is just the share of runs where the assistant won't name a single winner. Ask "best project tool for a 12-person team" twenty times: fourteen outright picks and six shortlists is a 30% hedge rate; three picks and seventeen shortlists is 85%. Put hedge rate on one axis and purchase value on the other and you have what we call the **settledness map**. Budget lives in the top-right corner.

HEDGE RATE BY BUYER QUESTION (EXAMPLE)



Grey: settled, spend nothing. Blue: borderline, worth a test.
Green: open, bid.

THE SETTLEDNESS MAP



Top-right is where one impression can change a real purchase.

The fork rule: the decision happens on question two

Here's something we didn't expect when we started logging conversations. In search the decision happens on the query. In chat it happens on the follow-up. Almost nobody buys off the first answer to "best espresso machine". They narrow: "under £600", "for a beginner", "for a small kitchen". Every one of those constraints resets the tie, because the model has thinner data the more specific you get. A question that's settled at the top level is frequently wide open two turns later.

So two rules. Measure hedge rate on the follow-ups, not just the opening question, because the follow-ups are more open and closer to the purchase. And write your creative for the fork: the ad that wins has already answered "but is it any good for a beginner?" before the buyer types it.

Why twenty runs, fresh sessions and two assistants

Below twenty, one odd answer swings the number by five points. Past thirty you learn very little more. Memory has to be off, otherwise you're measuring your own chat history. And ChatGPT, Gemini and Perplexity read different sources and settle different questions, so one assistant's verdict is not the market's.

— PART 1

Running the audit, and when to leave the framework alone

The audit, in seven steps

- 1 **List the questions.** Everything a buyer might ask between "I have a problem" and "I bought it", forks included. Twenty to forty is typical. Search query reports and sales-call notes are the best raw material.
- 2 **Run each one twenty times.** Fresh session, memory off. ChatGPT plus at least one of Gemini AI Mode or Perplexity. Log the named winner, or "hedge".
- 3 **Work out hedge rate and share.** Hedge rate is runs with no winner. Share is how often each brand gets named.
- 4 **Attach a value.** Order value or deal size per question. Plot hedge against value.
- 5 **Move the money.** Off settled questions, onto open ones. Write for the fork and lead with the provable fact.
- 6 **Unsettle where you're losing.** Publish the counter-fact, get it cited, re-measure. No bidding until the hedge rate moves.
- 7 **Do it again in 30 days.** Models update and competitors publish. Nothing here is permanent.

It's worth running when

- ◆ You're spending, or about to, inside ChatGPT, AI Mode, Perplexity or Copilot.
- ◆ There are four or more credible options in your category. That's where open questions live.
- ◆ A single sale is worth thousands. Home energy, private dental, B2B software, premium ecommerce.
- ◆ You already do GEO or AI-visibility work and want the paid half of that.

Leave it alone when

- ◆ Every question is already settled for you. The audit will tell you to spend nothing, and it'll be right.
- ◆ Your category is off-limits. OpenAI shows nothing near health, mental-health or political queries.
- ◆ Your budget is under the \$50,000 minimum. Run the audit anyway for the intelligence; put the media into Google Ads.
- ◆ You want it to replace Google or Meta. It won't. Search still rewards intent and social still rewards creative. Run them on their own logic.

PLATFORM	THE AD UNIT	WHAT CHANGES
ChatGPT	Sponsored card beneath the answer; feed-powered shopping cards.	Audit with memory off. Put the tie-breaking spec in the feed itself.
Google AI Mode	Conversational Discovery Ads and Highlighted Answers, inside the answer.	Open questions are worth more here. Cross-check with your Search terms report.
Perplexity	Sponsored follow-up questions.	The fork rule sold as a product. Buy the follow-up buyers actually ask.
Copilot	Performance Max conversational placements.	Least control. Use the audit to decide whether to opt in at all.

The honest caveat

The mechanism, answer variance as a confidence signal, is real and you can check it in an afternoon. The budget rule built on it is new and has only been proven on our accounts. Run it on ten questions in your own category before you move serious money.

— PART 1

What we're basing this on

Three things: what the platforms have shipped, how language models behave when you ask them the same thing repeatedly, and what the first wave of advertisers told the trade press. Here's each, so you can check the reasoning rather than take it on trust.

What the platforms shipped

- **LATE 2025**
OpenAI starts pitching sponsored placements. First real ad channel inside a generative chatbot.
- **16 JAN 2026**
Ads announced for US Free and Go tiers. Nothing near health, mental-health or political queries.
- **9 FEB 2026**
ChatGPT ads go live. Early reports of a \$60 CPM. Target, Ford and Adobe in the first wave.
- **APR 2026**
Ads manager launches; pilot minimum drops from \$200k to \$50k. Six formats live or piloting across ChatGPT, Perplexity, AI Mode and Copilot.
- **MAY 2026**
First DSP access. A penetration study finds sponsored units in roughly a quarter of responses overall and about half of US ones. Google Marketing Live unveils AI Mode ad formats; AI Mode passes a billion monthly users.
- **JUL 2026**
Full "Advertise in ChatGPT" manager, with Best Buy, Lowe's and VistaPrint as launch cases.
- **2027–28**
Chatbot ads projected under \$1bn for 2026, around 3% of AI ad spend, with growth weighted to 2027–28. Don't expect confidence signals in a dashboard before then.

How models actually answer

A language model samples from a probability distribution. A question with one dominant answer returns that answer nearly every time; a question with several close candidates returns different picks on different runs, or hedges with a shortlist. That variance isn't noise. It's a readout of how settled the recommendation is. TIEBREAK does nothing cleverer than turn the readout into a number.

Two more properties matter. Models resolve close calls on facts they can source, so a verifiable claim moves a recommendation where "market-leading" doesn't. And persistent memory now shapes answers per user, which is the whole reason the audit runs in fresh sessions on more than one assistant.

What early advertisers said

Trade coverage through spring 2026 agrees on one thing: measurement was thin. Advertisers got impressions and not much else, which OpenAI has been patching with click pricing and a LiveRamp partnership. The creative lesson is equally consistent: units that open with a discount code die, because the text around them is an answer and a discount reads as noise.


Both findings point the same way

If the dashboard can't tell you whether an ad changed the outcome, the question to answer before spending is: was there an outcome to change? That's the question this framework exists to answer.

PART 2

The eight-step loop

Part 1 stops the biggest leak. It doesn't run an account, because it treats every torn question as equally torn and every month like the last. The playbook fixes that. Each letter now has a job, an output and someone who owns it.



LETTER	WHAT YOU DO	WHAT COMES OUT	WHO
T • Test	Run every buyer question, forks included, twenty times in clean sessions on two or more assistants.	Raw answer log	Analyst
I • Index	Turn the log into hedge rate, tie width and shortlist presence.	Signal table	Analyst
E • Evaluate	Attach purchase value; rate the fact gap.	Winnability Score	Strategist
B • Bucket	Sort every question into one of five tiers.	Settledness map	Strategist
R • Reallocate	Apply 60 / 25 / 15 / 0. Zero out settled questions in both directions.	Budget plan	Media buyer
E • Engineer	One Tiebreaker Card per fork, leading with the tipping fact.	Creative set	Copywriter
A • Attack	Publish and seed counter-facts on high-value questions you're losing. No bidding yet.	Content brief	Content / GEO
K • Keep watch	Re-run at day 30. Log every tier change. Feed it back into T.	Delta report	Analyst

On a small team all four owners are the same person. The steps don't change; the hand-offs just disappear.


PART 2

Why one signal wasn't enough

The first time we ran the audit for a client with a crowded category, the hedge rate told us to bid on almost everything. It was technically right, the model was undecided, and it would have wasted most of the budget, because the model was undecided between five brands and ours wasn't one of them. Hedge rate says the model is torn. It doesn't say whether it's torn in a way you can use. So we added three more things to the log.

SIGNAL	WHAT IT MEASURES	HOW WE READ IT
Hedge rate	Share of runs with no single winner.	Under 30% settled · 30–50% borderline · over 50% open.
Tie width	How many brands get named when the model hedges.	Two is a coin toss and easy to tip. Five or more is a crowd and hard to.
Shortlist presence	Share of hedged runs where you are one of the names.	Over 60% you're in the tie · 20–60% adjacent · under 20% absent, and no ad fixes that on its own.
Fact gap	Whether a specific, checkable claim exists that would decide it for you.	Published: 1.0 · exists but unpublished: 0.6 · doesn't exist: 0.2.

SAME HEDGE RATE, VERY DIFFERENT OPPORTUNITY



Example data. All five questions hedge around 70% of the time. Two are prime targets. Hedge rate alone would have bought all five.

The one that matters most is shortlist presence, and it's the one people skip. An ad on an answer surface doesn't introduce you to the buyer; the model does that. If the model never names you when it hedges, your card is asking the buyer to consider a brand the thing they trust has already filtered out. That's a visibility problem, and it belongs to your GEO work, not your media budget. We treat absent questions as a hard stop for paid spend no matter what the hedge rate says.

PART 2

Scoring a question, and the five tiers


Four signals and a value need to collapse into one number so questions can be ranked, not just labelled. We kept the maths deliberately plain. It's for comparing questions inside your own category, not for precision.

Winnability = Hedge rate × Presence × Fact gap × Width factor × log10(Value)
 Width factor: 2 brands = 1.0 · 3 = 0.8 · 4 = 0.6 · 5+ = 0.4 · Value = order value or first-year contract

QUESTION	HEDGE	PRESENCE	FACT GAP	WIDTH	VALUE (LOG)	SCORE
Espresso, under £600	0.70	0.92	1.0	1.0	£520 (2.72)	1.75
Heat pump installer, 4-bed	0.70	0.75	1.0	0.8	£11,000 (4.04)	1.70
Project tool, 12 seats	0.70	0.55	0.6	0.6	£4,300 (3.63)	0.50
CRM for agencies	0.70	0.35	0.6	0.4	£9,600 (3.98)	0.23
Invisalign, London	0.70	0.08	1.0	0.6	£3,800 (3.58)	0.12

Notice the espresso question edges out the heat pump one on a twentieth of the value. It's a two-way tie we're nearly always in and the deciding fact is already public. Value matters; winnability matters more.

WINNABILITY SCORE BY QUESTION




TIER	DEFINITION	SHARE	WHAT TO DO
1 · In the tie	Open, presence over 60%, fact gap 0.6 or better.	60%	Bid at max. Creative leads with the tipping fact.
2 · Adjacent	Open, presence 20–60%.	25%	Bid, with GEO running alongside to lift presence. Watch cost per flip.
3 · Unsettle	Settled against you, high value, a counter-fact exists.	15%	Content and seeding, not ads. Bid once hedge rate passes 40%.
4 · Absent	Open, presence under 20%.	0%	No ads. Visibility problem. Send to the GEO backlog.
5 · Held	Settled for you.	0%	Watch it. Flag if hedge rate climbs five points in a month.

PART 2

The decision flow we hand to buyers

This is the sheet a media buyer works from for every question in the audit. The point is to take judgement out of the day-to-day. The only arguments left are about where the thresholds sit, and we have those once a quarter, not every Tuesday.

PER-QUESTION DECISION FLOW




These thresholds are where we've landed after a year of running it. Tighten them in crowded categories. Loosen them where you hold a clear fact advantage nobody else can match.

— PART 2

Mapping the fork tree

Buyers don't purchase off an opening question. They add a constraint, then another, and the model re-decides at every step. So we map it as a tree, with the four signals recorded at each node. The pattern you'll usually find: the opening question is settled, the first fork is borderline, and the second fork, the one nearest the purchase, is wide open and cheap to tip.

EXAMPLE FORK TREE · HOME HEAT PUMPS



Example data. The two green nodes are where one sponsored card decides a five-figure sale. The red node hedges hard but we're not on the shortlist, so it goes to the content team rather than the media buyer.

Three things we've learned doing this

Go two levels deep, minimum. Level one is nearly always a category constraint: size, budget, use case. Level two is logistics or trust: timing, location, who handles the admin. Level two is where the money is.

Weight nodes by traffic, not just score. Use search query reports and sales-call transcripts to estimate how many buyers actually reach each fork. A perfect score on a fork nobody asks is a footnote.


Write for the parent, deploy on the child. The card that wins at "installer near Leeds" is the one that has already answered "installed before winter". Buyers carry their previous constraint into the next question with them.

— PART 2

What actually tips a tie, and the Tiebreaker Card

Language models settle close calls on claims they can attribute and check. Across our audits the same hierarchy keeps showing up, and it's roughly the reverse of what most ad copy leads with.

WHICH CLAIMS MOVE A HEDGED ANSWER



Indexed to the strongest type = 100. Directional, from our UK audits. The bottom rung is still the one most advertisers open with.

The attribution test

Before a fact goes into a card, ask one question: if the model wanted to cite this, could it? A number with a source behind it (a lab, a standards body, a spec sheet, a price page with conditions) passes. An adjective doesn't. Facts that pass do two jobs: they tip the ad, and they're the raw material for unsettling a rival's held question.

What we've learned writing cards

- ◆ **One card per fork, never one per product.**
- ◆ **The FACT line carries its source.** Sample size, standard, test publisher.
- ◆ **FIT isn't optional.** Saying who you're not for is the quickest way to get read as an answer.
- ◆ **No price-off, no urgency.** Conditional price facts go in FACT, stated flatly.
- ◆ **Match the answer's register.** Plain, second person, no exclamation marks.

The Tiebreaker Card

A sponsored unit on an answer surface gets read as a continuation of the answer. Anything that reads as an interruption (a discount, a slogan, a brand-first headline) breaks that and gets skipped. The card is four lines that carry on where the answer left off.

TIEBREAKER CARD · STRUCTURE, WITH AN EXAMPLE

1 · FORK	Name the constraint the buyer is about to add. "Need it installed before winter?"
2 · FACT	The one claim that resolves the tie, source visible. "Average 11 working days from survey to commissioning across 340 installs in 2026, MCS-certified."
3 · FIT	Who this is right for, and who it isn't. "Best for 3 to 5-bed homes on mains gas. Not for flats or off-grid."
4 · NEXT	One low-friction next step, phrased as the buyer's next question. "See what your home qualifies for (2 minutes)."

— PART 2

Splitting the budget, and arbitrage between assistants

Once every question has a tier, allocation is arithmetic. The split we start from puts most of the money where one impression can change a verdict, keeps a working share on questions you're close to winning, ring-fences a little for offence, and forces the amount spent on settled questions to zero.

DEFAULT SPLIT BY TIER

60% · Tier 1 · In the tie

Paid placements, max bids

25% · Tier 2

Adjacent · bid + GEO

15%

Tier 3 · unsettle

0% · Tiers 4 and 5 · Absent and Held · no paid spend

Tier 3 money goes on content, seeding and monitoring, not media. It becomes media when the question opens.

SITUATION	ADJUSTMENT	WHY
Not many Tier 1 questions	60 → 40, Tier 2 → 40	You can't saturate a small set. Over-bidding Tier 1 just raises your own cost per flip.
A big rival holds your best question	Tier 3 → 25	Opening one high-value settled question beats marginal Tier 2 gains.
You hold most questions	Tier 1 → 30, hold 30 back	Dry powder, so you can respond in days if a held question starts to hedge.
Minimums exceed sensible spend	Pause paid, keep the audit	Run Google Ads and organic AI visibility off the map until the channel fits your size.

Arbitrage between assistants

ChatGPT, AI Mode, Perplexity and Copilot read different sources and settle different questions. The same buyer question can be held by a rival on one and wide open on another. It's the closest thing this channel has to an underpriced auction, and you only see it if you audit more than one.

PATTERN	WHAT IT LOOKS LIKE	THE MOVE
Split verdict	Settled for you on one assistant, open on another.	Spend on the open one. Borrow the settled one's cited sources as your counter-fact elsewhere.
Source lag	Open because one assistant hasn't picked up a rival's new claim yet.	Bid now; the window closes as sources sync. Get the counter-fact ready in parallel.
Format lean	One hedges with shortlists, another with "it depends".	Shortlist hedges want a FACT-led card; "it depends" wants a FIT-led one.
Region drift	Open in UK sessions, settled in US sessions.	Geo-limit the spend. One country's audit isn't global truth.


PART 2

Proving it worked when the dashboard can't

Ad platforms inside assistants report impressions, clicks and not a lot else.

TIEBREAK gives you an outcome the platforms can't: whether the model's own verdict moved. We call the test the **Settledness Delta**, and it costs nothing but time.

THE SETTLEDNESS DELTA · A 30-DAY HOLDOUT



- 1 Baseline.** Before any spend, record the four signals for every Tier 1 and Tier 2 question. That's your day-0 map.
- 2 Hold some out.** Pick a third of the Tier 1 questions at random and don't advertise on them.
- 3 Run 30 days.** Advertise on the rest as normal.
- 4 Re-measure.** Day 30, run the full audit again. Compare shortlist presence and named-winner share on advertised questions against the held-out ones.
- 5 Read the gap.** If advertised questions moved toward "settled for you" and controls didn't, the spend changed verdicts, not just clicks. If both moved, that's your content working, not your media.

This doesn't replace revenue attribution; it comes before it. In a channel this thin on reporting, knowing that spend moved the recommendation is the best leading indicator you'll get.

A caution on cause and effect

Models update on their own schedule, and the platforms say sponsored placements don't touch the answers directly. What ads change is buyer behaviour: clicks, visits, reviews, mentions, which eventually become sources. Expect the delta to lag spend by weeks, and treat one month's movement as a reason to keep testing rather than proof.

METRIC

Flip cost

- ◆ Spend to move a question from open to settled for you.
- ◆ The AI-era cost per acquisition.

METRIC

Unsettle time

- ◆ Days for a counter-fact to lift a rival's hedge rate.
- ◆ The 2027 version of "time to rank".

METRIC

Δ Settledness

- ◆ Verdict movement, advertised minus control.
- ◆ The lift number the dashboard can't give you.

— PART 2

Category notes, and the mistakes we keep seeing

The framework doesn't change between categories. The forks, the facts that tip and the sensible thresholds do.

HIGH-TICKET HOME SERVICES

Heat pumps, solar, windows, roofing

- ◆ Forks are logistics: timing, area, grant handling, disruption.
- ◆ What tips it: certified install counts, time-to-commission, accreditations.
- ◆ Bid from hedge 40%; deal value carries lower presence.

B2B SOFTWARE

Tools bought by teams

- ◆ Forks are fit: team size, integrations, compliance, migration.
- ◆ What tips it: named integrations, uptime, security standards.
- ◆ Tie widths run wide here; weight presence over hedge rate.

PREMIUM ECOMMERCE

Considered products, £300 to £3,000

- ◆ Forks are constraints: budget, size, skill level, use case.
- ◆ What tips it: spec differentials, independent test scores.
- ◆ Your feed is the FACT line; keep specs complete and current.

REGULATED SERVICES

Dental, clinics, finance

- ◆ Many placements restricted; audit for visibility, not media.
- ◆ What tips it: regulator registrations, published outcomes.
- ◆ Use the map to steer Google Ads and GEO instead.

Nine ways it goes wrong

- 1 Auditing with memory on.** You're measuring your own history, not the model.
- 2 Stopping at the opening question.** The forks are where the purchases happen.
- 3 Treating hedge rate as the whole story.** Absent plus high hedge is a content brief, not an ad buy.
- 4 Bidding on held questions "to be safe".** That's just margin leaving the building.
- 5 Leading creative with the brand.** Fork and fact go first. Brand is line five.
- 6 Using superlatives as facts.** Bottom of the ladder. Swap for something with a source.
- 7 Auditing one assistant.** You'll miss the arbitrage entirely.
- 8 Running the audit once.** June's settled question is September's open one.
- 9 Judging it on clicks.** Run the Settledness Delta or you'll never know whether the spend did anything.

— PART 2

Where it sits next to Google, Meta and GEO, and the monthly rhythm

TIEBREAK doesn't replace search, social or organic AI visibility. It sits in the middle and tells each one what to do about the questions the assistant is deciding.

THE SETTLEDNESS MAP AS THE SHARED BRAIN



One audit, four instructions. Search keeps buying intent, social keeps buying attention, GEO keeps changing what the model says, and AI ads only pay where the model is torn.

The 30-day rhythm

- DAY 1–3**
Test and Index. Full re-run on every audited assistant. Update the signal table.
- DAY 4**
Evaluate and Bucket. Recompute scores. Log every tier change and why it moved.
- DAY 5**
Reallocate. Apply the split. Pause anything that dropped to Tier 4 or 5; promote anything that rose into Tier 1.
- DAY 6–8**
Engineer. Refresh cards for new forks. Retire any card whose fact a rival has now matched.
- DAY 9–25**
Attack. Ship the counter-fact content for Tier 3. Seed it where the models read.
- DAY 26–30**
Keep watch. Spot-check held questions. Run the Settledness Delta. Pull next month's new questions from sales calls.

— NEXT STEPS

Stop paying for answers the AI has already given.

Run the audit on ten questions in your own category before you move real money. If hedge rate and shortlist presence predict where your ads move outcomes, you'll have a case study nobody else has yet.

STEP 1

List 20–40 buyer questions

STEP 2

Run each 20× on 2 assistants

STEP 3

Score, bucket, reallocate

A note on the numbers

Thresholds, weights, tipping-power figures and budget shares are Global PPC operating defaults, drawn from audits on UK accounts. They're where we start, not universal constants. Charts marked as example data use illustrative figures to show the method; run the audit on your own category before allocating budget.

© 2026 Global Performance Propel Company. Free to share with attribution. Edition 1, September 2026.