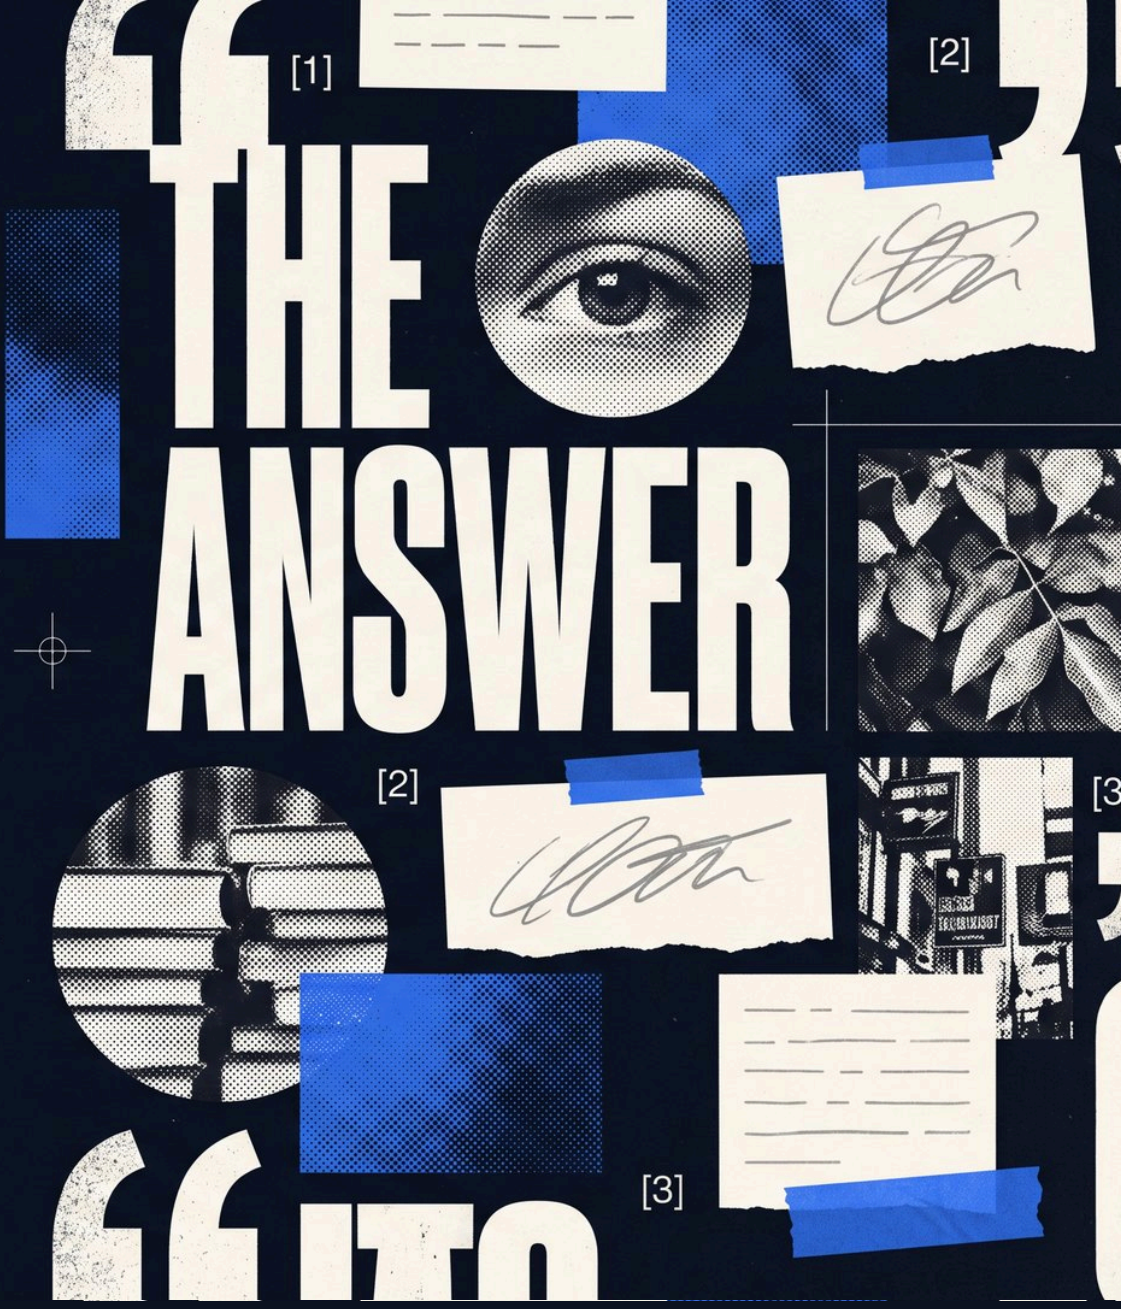


GEO PLAYBOOK · v1.0

GEO Playbook

General Edition

Follow it in any industry and earn your place in AI answers.

Canlah AI

Canlah AI · Singapore · A measured method for AI search · September 2026

Canlah AI

GEO Playbook

General Edition

Follow it in any industry and earn your place in AI answers.

Version	v1.0 · September 2026
Author	Canlah AI (Haoyang Pang)
Licence	Content CC BY 4.0: share, adapt and use commercially; credit "Canlah AI · GEO Playbook" and link the original page
Archive	DOI 10.5281/zenodo.23005022 (Zenodo)
Data	DOI 10.5281/zenodo.23005020 · canlah.ai/data/ai-citations-sg-2026/

This book does not promise rankings. Every rule states its evidence and sample size, or is marked as not measured. The compliance guidance is not legal advice.
Cite as: Canlah AI (2026). GEO Playbook · General Edition, v1.0. https://doi.org/10.5281/zenodo.23005022

The web edition has one page per chapter and is kept up to date; it is the reference version:
https://canlah.ai/playbook/
Scan for the web edition

Contents

Chapter 0 **Start here: how to read this book, and how tightly your advertising is regulated** 7

0.1 The book in one sentence, and the 90-day reading order 8

0.2 Decide which side you are on first 12

0.3 Which page types are limited for you, and what to build instead 19

0.4 The three labels for compliance sentences, and the stop-and-escalate rule 25

Chapter 1 **How AI picks its answers: what buyers ask, and which pages ChatGPT and Google's AI cite** 29

1.1 How buyers ask, and what AI reads 30

1.2 Two legs: ChatGPT looks for the source, AI Mode for second-hand summaries 35

1.3 Three gates and three paths 37

1.4 Three kinds of things not to do 42

Chapter 2 **Let AI read your website: check first, then fix the settings** 49

2.1 The five gates at a glance, and gate 0: access logs (read-only) 50

2.2 Gate 1 read-only audit: what to check in nosnippet, robots and the WAF 54

2.3 Gate 2: four identities, fetched live (read-only) 55

2.4 What each crawler is for, and judging the door leg by leg (read-only) 57

2.5 Gate 3, read-only: connect the two free AI reports 63

2.6 Fixing the door: nosnippet, the four-step robots merge, the WAF allowlist 64

2.7 Gates 3 and 4: indexing paths, and JSON-LD sealed once 68

2.8 The launch gate for every page 70

2.9 Door-layer troubleshooting and this chapter's acceptance checks 71

Chapter 3 **Get AI to recognise your business: the same facts everywhere** 75

3.1 Why this comes before writing pages · The two-hour checklist 76

3.2 The fact table and the organisation facts page 79

3.3 Person pages and six-trace alignment 85

3.4 Third-party credential tiers, Wikidata and the five business profiles 92

3.5 The brand six questions, and what to do when you find an error 96

3.6 The positioning-statement myth, and the exits from this chapter 100

Chapter 4 **Which questions to go after: pick 30 buyer questions and record how AI answers them today** 105

4.1 The full target-picking process and the twelve pre-start questions 106

4.2 The question pool: where the 30 questions come from, how they are balanced, how they are signed 110

4.3 Coarse screen and piles, full rounds, and the brand-six baseline 113

4.4 The web control leg and the frozen baseline (the book's only full spec) 117

4.5 Buyer types, annual opportunity value and the three checks 122

4.6 Four states, two denominators and the abstain line 125

4.7 Actions by state, the ten question types and this quarter's slots 129

4.8 Two shortcuts and this chapter's checklist 135

Chapter 5 **Writing pages (1): the 46 page types AI cites, and the rules for every page** 139

5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side 140

5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers 149

5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions 159

5.4 The 30-day writing order and the three mechanical gates 168

5.5 The page unit: one intent cluster, one page 175

5.6 Two legs, two kinds of sentence 179

5.7 Seven rules that hold for every page type 183

5.8 The conclusion block, the density gate and the source gate 186

5.9 Content form, FAQ, Chinese pages and other-language pages 191

5.10 Shared parts for all page types: six checks for numbers, no public price, billing units, dates and schema 194

Chapter 5 **Writing pages (2): the 9 most-cited page types, block by block** 198

5.11 Single-service price page (pt01) 199

5.12 Price guide page (1): why it is always the first page (pt02) 201

5.13 Price guide page (2): blueprint, column headers, sentences (pt02) 204

5.14 Official pricing page (pt03) 206

5.15 Comparison page (pt04) 208

5.16 List page (pt05) 212

5.17 One question, one page (pt06) 214

5.18 Eligibility and process page (pt07) 217

5.19 Remedy and second-opinion page (pt08) 219

5.20 Entity anchor pages: three subtypes, and the product detail page (pt09) 223

Chapter 5 Writing pages (3): quick cards for 33 more page types 227

5.21 Tier-B price types: subsidy limits, parameter basis, category lists, calculators, time-limited promotions 228

5.22 Tier-B selection and reputation types: third-party reviews, sentiment and forums, review aggregates, self-built reputation, verdict-first 236

5.23 Tier-B questions and rules types (1): regulatory obligations, definitions 245

5.24 Tier-B questions and rules types (2): procedure steps, preparation lists 249

5.25 Tier-B questions and rules types (3): schedules, collected FAQ, policy hubs, change notices, misconceptions 252

5.26 Tier-B entity and product facts types (1): branches, verification pages, directory listings 259

5.27 Tier-B entity and product facts types (2): compliance proof, terms, how-to docs and integrations 265

5.28 Tier-B primary sources (1): legislation, regulator guidance, official replies, device documents, case law 272

5.29 Tier-B primary sources (2): statistics sources, registers 281

Chapter 5 Writing pages (4): 4 unverified page types, and checks before you publish 287

5.30 The four tier-C types: unverified page types 288

5.31 Handing over: five steps per page, the red line on lying, and self-checks 292

5.32 Master table of common mistakes 298

Chapter 5 Writing pages (extra): producing hundreds of pages, from studying a competitor to launching in batches 301

5.D1 The seven steps at a glance, and what to tell clients 302

5.D2 Pick the right study target: the competitor that wins on on-site form 303

5.D3 Take the architecture apart until you can write from it 305

5.D4 Correct with measurement: the competitor sets the skeleton, measurement sets the details 307

5.D5 Topic selection: turn every head and long-tail keyword into a list with one main query per page 309

5.D6 The production pipeline: three agent passes per article, plus automated lint gates 311

5.D7 Six traps we hit: build the safeguards in before you start 314

5.D8 Source audit: trace every number in data reports to its source 315

5.D9 Launch in batches: from launch day to the keep-or-cut decision on day 45 317

5.D10 Add new keywords every month: new keywords come in, a person reviews before anything goes out 318

Chapter 6 Beyond your website: get onto the 20 pages AI already cites 320

6.1 Where the effort goes, and two numbers that must not sit side by side 321

6.2 The send gate, paid listings and the citation-slot table 323

6.3 Six dispositions, bylined contributions and outreach letters 325

6.4 Three source groups, and the Group B actions 329

6.5 Reviews: how to ask, how to reply, how to fill the empty slot 331

6.6 Annual rankings and off-site asset decay 334

6.7 Long videos narrated by the named expert 337

6.8 Off-site verdicts, common mistakes and rituals to cut 341

Chapter 7 **Monthly retest: how to tell whether it's working and what to do next month** 345

7.1 What a retest produces, and the rulers 346

7.2 The noise band, the page-level signal and the ten monthly steps 351

7.3 Not-moved triage and next month's three points 355

7.4 Holding position and the 90-day settlement 364

7.5 Ruler discipline: change the ruler and nothing is comparable 369

Chapter 8 **Using this in a new industry: four steps to work out your approach** 374

8.1 Four steps for a new industry: ten questions → side check → ten axes → decision table 375

8.2 Five steps to decide the side, and three criteria 378

8.3 The ten axes (1): page types, containers, anchors, regulation, question shapes 384

8.4 The ten axes (2): decision-makers, capacity, attribution, half-life, language 392

8.5 The decision table and rulings 396

8.6 Seven tactics: how regulated and unregulated sides differ, and why each moves rankings 400

8.7 Cheap tactics on the unregulated side, and substitutes on the regulated sides 406

8.8 Unregulated ≠ easy: admission criteria J0–J7 408

8.9 Cross-border: whose law applies 412

Appendix A **The evidence: the numbers and sources behind every conclusion** 416

A.1 Evidence for mechanics, the door and identity 417

A.2 Evidence for picking targets, writing pages, off-site and retests 425

A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types 432

A.4 Page-type measurements (2): tier B and C samples, and the AI Mode retest list 439

A.5 List of sources 446

Appendix B **Templates: robots.txt, wording, outreach letters and tables** 448

B.1 Door templates 449

B.2 Approved wording and scripts 452

B.3 Outreach emails and review replies 457

B.4 Headers for work orders, ledgers and monthly reports 463

B.5 Blueprints for unverified page types (for industry editions to include) 467

Appendix C **Five industries side by side** 473

C.1 Ten-axis values for the five industries covered, and three things for week one 474

Appendix D **How to write numbers, and the glossary** 487

D.1 Number discipline: how to label numbers, and what stays internal 488

D.2 Glossary 492

Chapter 0

Start here: how to read this book, and how tightly your advertising is regulated

The 90-day reading order, how strict your industry's rules are, and which pages you can't build

0.1 The book in one sentence, and the 90-day reading order

0.2 Decide which side you are on first

0.3 Which page types are limited for you, and what to build instead

0.4 The three labels for compliance sentences, and the stop-and-escalate rule

0.1 The book in one sentence, and the 90-day reading order

What you'll do in this section: Memorise the book in one sentence, and read through the book following the 90-day roadmap. When you're done, you'll have a schedule: which section to read on which day, and which days are read-only with no changes allowed; and if you can only do three things, which three they are and what to measure before you start.

The whole book compressed into one sentence:

> **Write your facts in clean, visible HTML, let the crawlers in, make machines sure which business you are, then get onto the third-party pages AI already reads.**

For why this sentence holds (AI fetches pages live, reads only visible text, and the main battlefield is not your own website), see → General Edition 1.1 How buyers ask, and what AI reads.

The book uses only this one week-numbering system: D1–D6 are the audit and baseline freeze; W1 counts from the day you fix the door (the split day); the D90 settlement falls around week 13. The writing calendar in → General Edition 5.4 The 30-day writing order and the three mechanical gates and the settlement in → General Edition 7.4 Holding position and the 90-day settlement both run on this same week numbering.

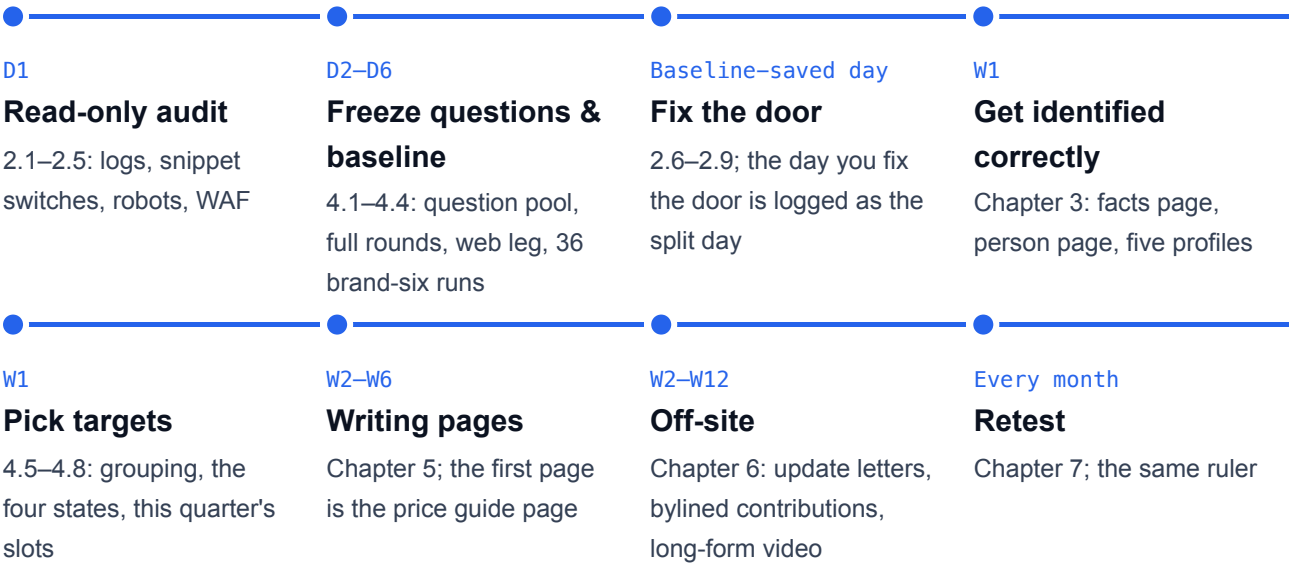


Figure: read the book in this order across 90 days; the first two rows are read-only, no changes

Four things the figure cannot show:

- **Why the baseline comes before you fix the door**: If you freeze the baseline after the door is already fixed, what you freeze is a shelf already contaminated by your own changes, and before-and-after become impossible to compare. The cost is only a few days' delay in fixing the door, in

exchange for an honest set of before-and-after numbers on D90. "Baseline saved" means all three are on disk: the baseline Top 20, the noise band measured from the same-week retest, and the 36 brand-six runs (spec in → General Edition 4.4 The web control leg and the frozen baseline (the book's only full spec)).

- **On D1, besides the read-only audit, there is only one more thing to do: get access.** Access logs need ops or the hosting provider to grant permission, which is often slower than the audit itself, so ask for it on day one.
- **The build order cannot be swapped**: with the door closed, everything downstream is multiplied by 0. Door and identity are ×0/1 switches; writing pages and off-site are the only variables that can add points month by month (see → General Edition 1.3 Three gates and three paths). The one exception is the product feed: it does not go through the door, so it can run in parallel with fixing the door. Build order and measurement order are two different things — the ruler's acceptance check comes last, and the freeze comes first.
- **Writing pages and off-site both start in W2**: the vast majority of the pages AI cites are not on your own website, so off-site needs the same order of magnitude of effort as changing your site (evidence in Appendix A.1). Where the sequencing of effort conflicts with this figure, the figure wins.

If you can only do three things

Prerequisite
Side check and baseline
0.2 side check; the 48-hour baseline in 4.8, read-only until saved

Thing 1
Door and identity
2.6–2.9 fix the door; 3.1–3.3 get identified correctly

Thing 2
Price guide page
One page for all products, in English and in Chinese; 5.12–5.13

Thing 3
Institutional sources and video
6.3–6.4 institutional sources; 6.7 your own long-form video

Every month
Retest
7.1–7.2; only counts as moved once it's outside the noise band

Figure: the order for the three things, with measurement before you start

What decides the order: **evidence strength × how reachable it is within 90 days × whether it is a ×0 multiplier switch**.

Thing	Why it's this one	How you know it's done
Prerequisite · Side check and baseline	Skip the side check and the pages you write may break the rules for your side; skip freezing the baseline and you will have no honest before-and-after numbers when you're done	The first line of the work order states the side and the switches; the 48-hour baseline is saved before the door is fixed, and the noise band is filled in during week 1
Thing 1 · Door and identity	A ×0/1 multiplier switch: if crawlers cannot get in, 20 pages written is still 0; if AI identifies the wrong business, the credit goes to a namesake instead. Skip this and the other two are wasted	Access logs show real OAI-SearchBot hits returning 200; your price figures can be found in the HTML from Bing's Live Test; ChatGPT can quote the price line word for word when given the URL; factual errors in the brand six questions = 0
Thing 2 · The all-products price guide page (pt02)	This page type's citation density is counted by page type, not by who owns the domain, so a page you build yourself benefits just the same; the price guide pages being cited locally right now are already supplied by practitioners themselves; for Chinese-language queries the field is mostly open, the only battlefield where you can go from 0 seats to some within 90 days, and all it takes to enter is one landing page (evidence in Appendix A.2)	The 2–3 unique fact strings logged for this page are found, by string search, in the actual wording of this month's saved answers
Thing 3 · Institutional sources and your own long-form video	Official registers, association profiles, manufacturer dealer locators, and government or statutory practitioner registers are the most certain sources: they don't depend on a third-party editor's approval, they aren't rewritten every year the way rankings are, and their share is growing; long-form video is the highest-return item per unit of effort — cover only process, equipment and cost breakdown, never outcomes or testimonials	On the list and cited X/20 rises by 1 each month; new landing pages start appearing in the Search Console generative AI report impressions for the Google leg

Two additions: on the strictly regulated side, Thing 2 is written differently: no peer price ranges, and the official benchmark gets its own section, with your own price never appearing inside it (see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers). For Thing 3, 60% of off-site effort goes to this group of sources; every thick page gets one 8–15 minute long-form video, with the full transcript and that page's URL in the description, timestamped by H2 — spec in → General Edition 6.7 Long videos narrated by the named expert.

Three other common candidates	Why they don't make the cut
Writing schema page by page, filling in sameAs	Live fetch doesn't read JSON-LD; it only enters Google's Knowledge Graph indirectly
Paid listings and directory submissions	Paid placements and advertorials are only 0.3% of AI citations
Piling on more articles	With the door closed, page count is multiplied by 0; with it open, 5 exclusive facts beat 5 educational articles

Figure: three other common candidates, and why they don't make the cut

The first row doesn't mean schema is useless: set the template up once when the site is built, and it doesn't belong on the per-page checklist (see → General Edition 2.7 Gates 3 and 4: indexing paths, and JSON-LD sealed once). Samples and sources for all three rows are in Appendix A.1.

Which section to turn to now

Your situation	Go to
You haven't decided whether to do this	4.8, shortcut 1: half a day to judge whether it's worth starting; for an industry this book doesn't cover, read 8.8 first
You've decided, and it's day one	2.1–2.5, read-only, change not a single word
The door is open, but AI gets your information wrong, or confuses you with a namesake	Chapter 3
You want to know "do I actually have any seats"	4.1–4.4: get a ruler first, then ask whether it has gone up
You have a ruler, but don't know which questions to run	4.5–4.8
You're about to start writing pages	Chapter 5; read 5.1–5.3 first
You're about to send letters or shoot video	Chapter 6
A month has passed and the numbers haven't moved	The not-moved triage in 7.3; start at layer 1, no skipping layers
You're in a regulated industry, before writing any paragraph	0.2, 0.4, 5.2–5.3
You want to quote a number publicly	Appendices D.1 and B.2

Your situation	Go to
You're switching to an industry this book doesn't cover	Chapter 8, then back to 0.2 to recheck your side

Two more sentences to remember

> **AI reads visible HTML.** Prices and facts must be written on the page in clean `<table>` / `<dl>` markup; schema is only a copy.
>
> **A number without a noise band is not a number.** Same batch of questions, same engine, same number of passes, retested three times in the same week: only a change larger than the range (max − min) counts as movement.

0.2 Decide which side you are on first

What you'll do in this section: Answer four questions to work out whether you're on the strictly regulated side, the lightly regulated side, or the unregulated side; then check off the other five add-on switches; where you're unsure, start on the stricter side, and write the conclusion into the first line of the *First work order*.

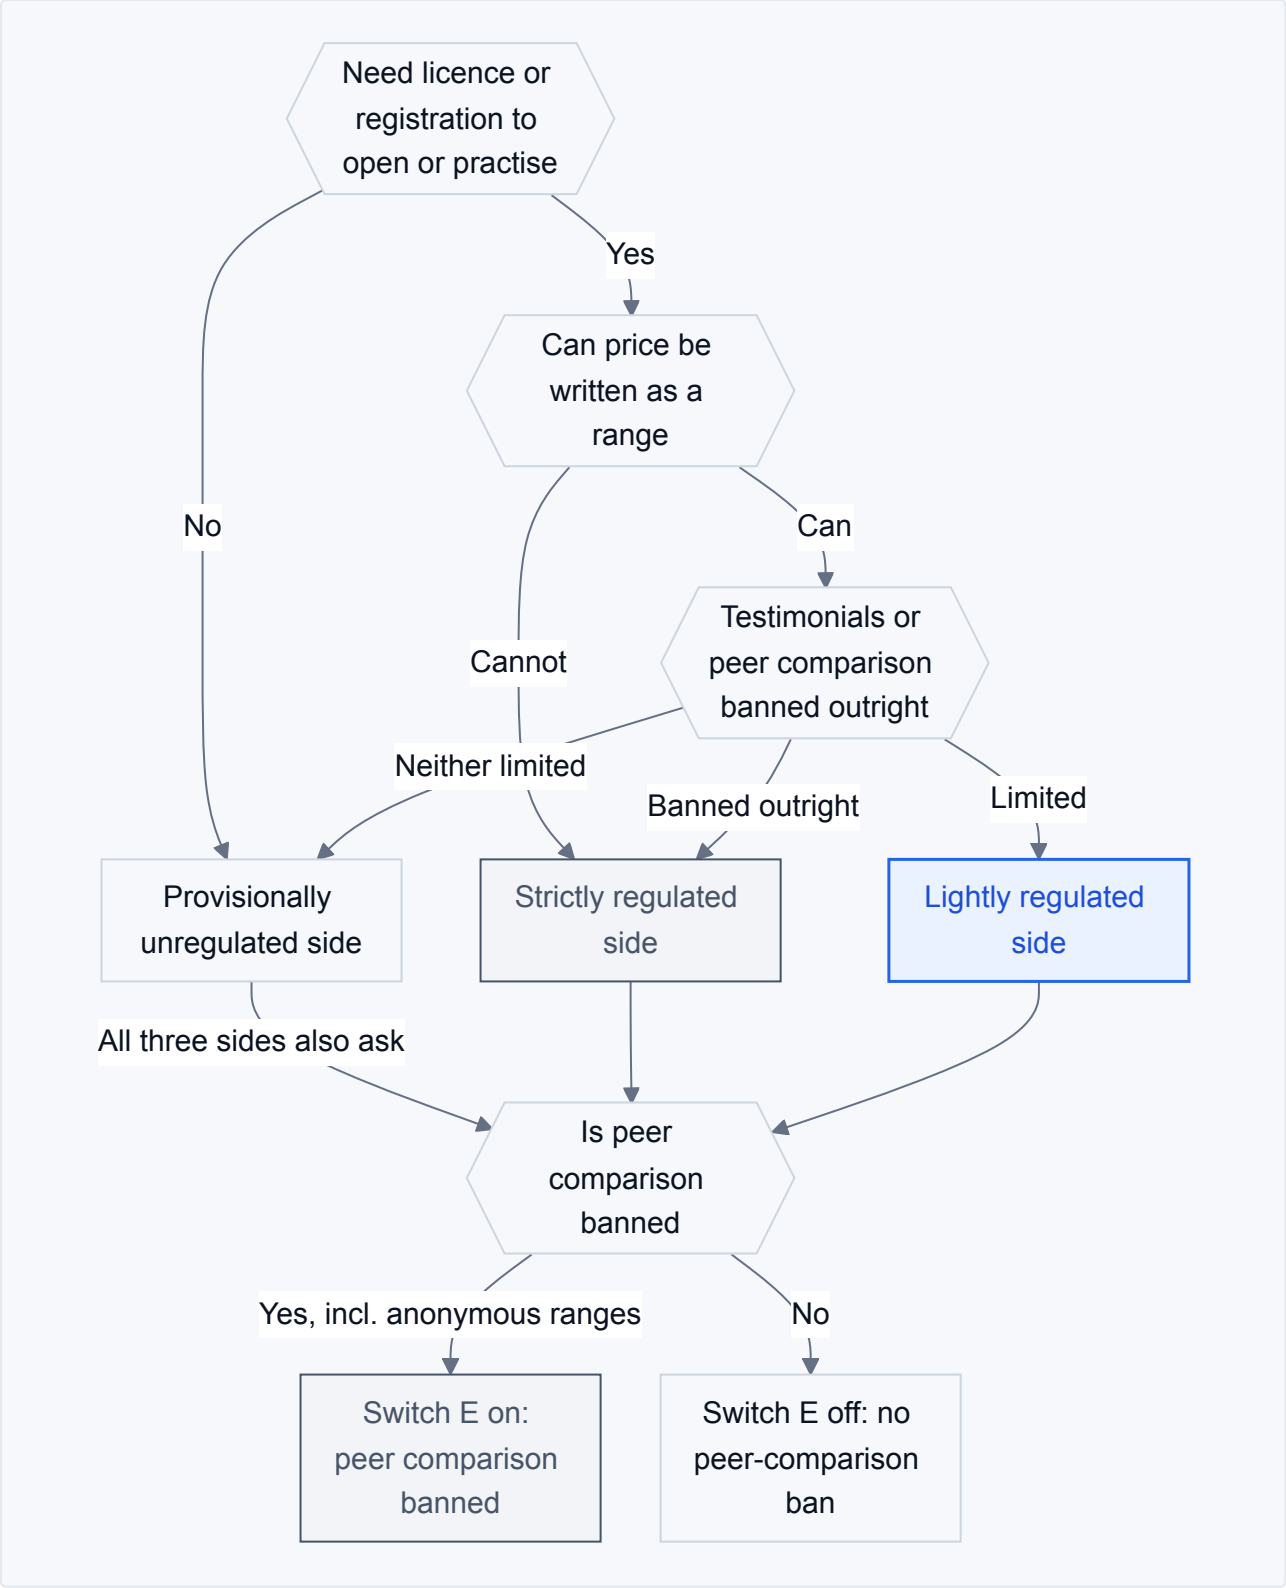


Figure: four questions decide which side you're on; question 4 alone decides switch E

For question 1, count it "Yes" whether the business itself needs a licence to open, or the individuals doing the work need a personal professional registration; where the individual needs registration, an extra layer of professional-conduct rules applies. Answering "No" to question 1 only makes you provisionally unregulated: you still have to run the three reverse checks in → General Edition 8.2 Five steps to decide the side, and three criteria (does the copy make any claims about health, efficacy, earnings,

outcomes or employment; does the audience include minors; is the buyer overseas) — any hit moves you up one level. And "unregulated" does not mean "no rules" — see → General Edition 8.8 Unregulated ≠ easy: admission criteria J0–J7. In question 3, "limited" means testimonials can be written but with conditions, or only part of comparison is banned; "neither limited" means there is no rule specifically governing advertising and only the general consumer-protection baseline applies; even if you need a licence, you still land on the unregulated side. In question 4, "peer comparison" includes both anonymous market ranges that don't name anyone and ranges that cite a third-party source.

Where you land somewhere in the middle and can't decide, start on the stricter side — never carry a "TBD" into the work order. The reason is that the costs are asymmetric: misjudge yourself as lightly regulated, and the pages you write may trigger penalties and joint liability; misjudge yourself as strictly regulated, and it only costs you a few tactics, which you can fix later. These four questions are the field version; the full method (checking the regulator's documents, checking product law, the three reverse checks) is also in 8.2.

> **Example** (medical aesthetics) A clinic holds a healthcare services licence, so question 1 is "yes"; prices can only be fixed in tiers, never written as a range, so question 2 already puts it on the strictly regulated side and question 3 doesn't need answering; question 4 is still asked: Statute text bans peer comparison, so switch E (peer comparison banned) is on.

Once the four questions have settled your side, there are five more switches to check off separately — switches don't change which side you land on, they just add one more hard gate on top of that side:

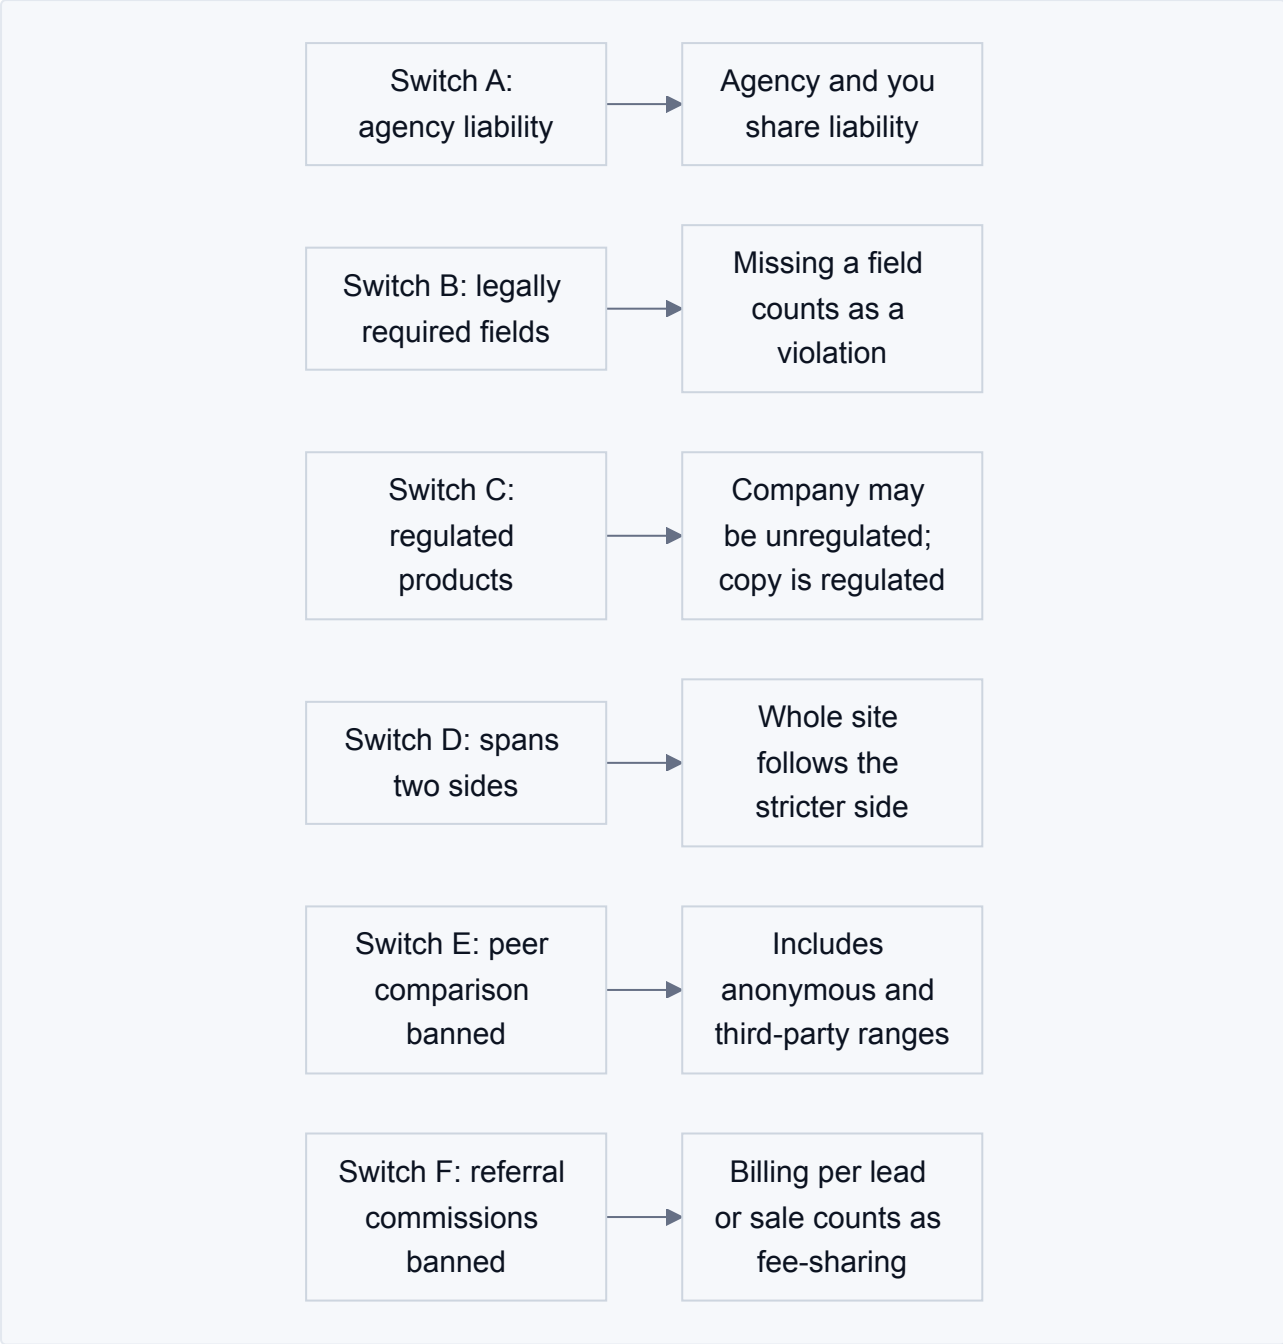


Figure: what hard gate each of the six add-on switches adds

Three of these are the easiest to miss. **Switch A (agency liability)** doesn't depend on whether you run the ads yourself — as soon as someone else writes or publishes on your behalf, if the material has a problem, the agency and you share liability; your own responsibility isn't reduced at all, and every outreach letter or piece of material someone else sends on your behalf counts as if you sent it yourself. **Switch C (regulated products)** can still hit even after question 1 has already put you on the unregulated side: the company itself may need no licence, but if what it sells needs to be registered, licensed or notified, that's a hit — "company unregulated, copy regulated" are two separate judgements, decided by product category, not by company; passing on the company side doesn't mean the copy passes too. **Switch D (spans two sides)** is judged the simplest way: when one business runs both a regulated and an unregulated line, you don't split the judgement by page or subdomain — the whole site follows the stricter side, and you cannot mix strict-side wording on one page with light-side wording

on another. The full criteria for switch B (legally required fields) and switch F (referral commissions banned) are in → General Edition 8.2 Five steps to decide the side, and three criteria; for now, remember the conclusions: when switch B is on, fields like the registration number and the registered name aren't a nice-to-have — missing them is a violation, and they're also the cheapest disambiguation anchor for the identity gate; the regulated side gets one for free, the unregulated side doesn't and has to build its own, see → General Edition 8.6 Seven tactics: how regulated and unregulated sides differ, and why each moves rankings; when switch F is on, any referral channel billed per lead, per sale or by commission is treated as fee-sharing for clients on your part, and cannot be used.

> **Example** (e-commerce) An online store selling skincare, supplements and massage devices needs no industry licence for the company itself, but its ad copy is regulated — switch C (regulated products) is on, judged category by category.

The strictly regulated side has the most switches on, and they shut things down the hardest. Which shelves get closed entirely:

What a regulated business cannot do		The shelf it shuts
Buy or self-build a Best/Top list spot	→	The whole list-page shelf is closed; redirect the effort elsewhere
Write price as "from / as low as / original price"	→	The main page type for price questions — its wording is locked down by regulation
Build a peer comparison page	→	The comparison page, the single page type with the highest citation density
Actively request reviews	→	Reviews can only be observed passively, never used as a metric
Testimonials, before-and-after photos, laudatory terms	→	Review, reputation and verdict-first page types

Figure: what a regulated business cannot do, and which shelf each one shuts

Deciding your side isn't a legal question, it's a rankings question: every row above closes one of the shelves with the largest share of AI citations (share evidence in Appendices A.1 and A.2). These five rows aren't the whole list — the full list is in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers and → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions, and every closed item comes with a specific replacement — it is never left blank.

Once those shelves are closed, persuasiveness on the strictly regulated side can only come from checkable facts: a fixed final price and its tiers, what's included and what's excepted, the registration

number and a link to the official register, the real model numbers of equipment and materials, procedure duration and number of sessions, the term and scope of the follow-up and replacement policy, and operating figures stated with a clear basis and no adjectives (list in → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions). At the organisation level there is almost no room to boast, but a practitioner is allowed to publish their own qualifications, specialisms, practice arrangements and contact details — so on the strictly regulated side, the person page isn't a bonus, it's the main line (see → General Edition 3.3 Person pages and six-trace alignment).

Deciding your side doesn't change the process itself — it only changes how six steps within that process are written:

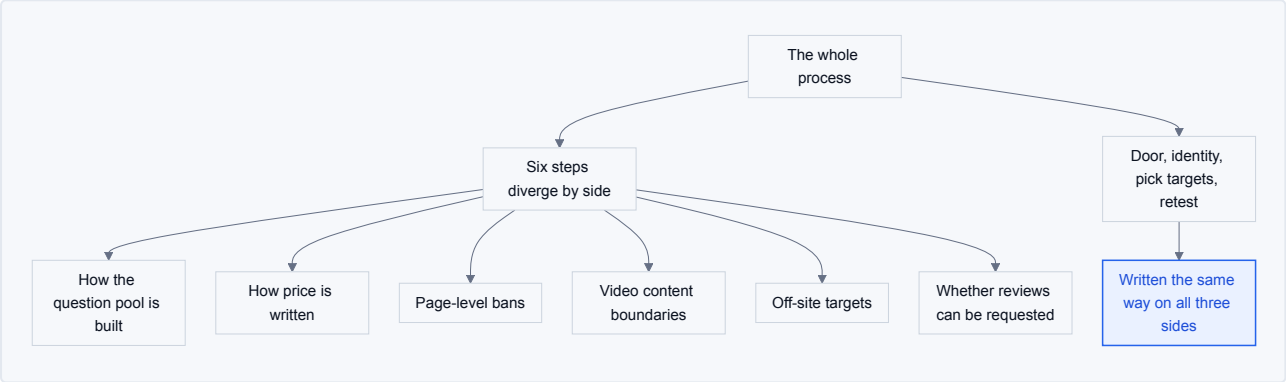


Figure: the three sides diverge in only six steps; everything else is shared

Deciding your side doesn't change the door or identity layers; it only decides which tactics you can use, and what replaces the ones that are taken away. How each of the six steps is written, by side:

Step	Strictly regulated side	Lightly regulated side	Unregulated side
Question pool	Price questions must be phrased so a "fixed final price" can answer them	No restriction	No restriction
Price wording	Fixed, itemised prices by tier, one item one price, with the tax basis locked in writing; never write prefixes like "from" or "as low as", never write a range (handle case-by-case variation with tiers plus a defining condition for each tier); never write a percentage change, an original price, a discount, a bundle, a time limit or instalments	May write a range and a billing method, but must state the scope, what's excluded and any third-party fees	May write a range and a billing method, but must state the scope and what's excluded; mandatory fees are included in the stated price

Step	Strictly regulated side	Lightly regulated side	Unregulated side
Page-level bans	Never write laudatory terms, true or not; never write outcomes, success rates or recovery time; never post before-and-after photos; never reprint testimonials or reviews (including screenshots and star ratings); never compare against peers, named or unnamed; never write "free / complimentary / $0"; never cite a third-party directory that calls itself "government-approved"; never write a statement you cannot prove	Laudatory terms must be provable; testimonials may be used but with conditions; do not compare against peer businesses, only compare plans and delivery formats (Conservative line (not statute text))	Laudatory terms may be written as long as they're provable; case studies and testimonials may be posted; peer comparison pages may be built (same basis, with sources and a check date)
Video content boundaries	Cover only process steps and duration/number of sessions, equipment model numbers, cost breakdown and what's included or excepted, and when the procedure isn't suitable; never cover outcomes, never show before-and-after images, never feature a customer, never mention peers	This side's page-level bans apply the same way; otherwise no restriction	No restriction
Off-site targets	Never buy or go after a spot on a Best/Top list; disclosing that it's paid doesn't excuse it either; if you are already listed on one of these pages, ask the publisher to remove you, and don't send a correction letter. The main effort goes to official registers, association profiles, manufacturer locators, and educational bylined contributions	Compete for editorial slots, send update letters and correction letters; if switch F (referral commissions banned) is on, don't use referral platforms billed per lead, per sale or by commission	Go after Best/Top lists: directories and lists are both normal targets; for a paid slot, first confirm the page is already being cited
Reviews	Never actively request reviews, under any circumstances; observe passively only — never use as a work metric or an acceptance measure	May actively request: no incentives, no ghostwriting, no cherry-picking, no dictating the content, and never confirm a customer's identity externally. Reprinting a single review on a page follows the testimonial rule (written consent, with name, relationship and year attached); for law firms, follow the industry edition	May actively request: no incentives, no ghostwriting, no cherry-picking, no dictating the content; may display on the page (as visible text, stating sample size, period and source, without altering the wording), but never post a self-assigned rating

The price guide page rule "never place peer data side by side" also follows the side: on the strictly regulated side and wherever switch E (peer comparison banned) is on, never place peer data alongside your own — on the strictly regulated side, even the official benchmark gets its own section, with your own price never appearing inside it; on the lightly regulated side (where switch E is off), the market range can only be an independent, sourced fact section, never in the same table as your own price (Conservative line (not statute text)), and comparison covers only plans and delivery formats, never businesses (Conservative line (not statute text)); the unregulated side can build peer comparisons, on the same basis, with sources and a check date (see → General Edition 5.15 Comparison page (pt04)). Full wording for each step: for the question pool and price, see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers; for page-level bans and reviews, see → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions; for video, see → General Edition 6.7 Long videos narrated by the named expert; for off-site, see → General Edition 6.2 The send gate, paid listings and the citation-slot table.

0.3 Which page types are limited for you, and what to build instead

What you'll do in this section: First look at the 15 page types with restrictions; where the strictly regulated side can't build one, swap in the legal page type for the same intent, then, if switch E (peer comparison banned) is on, adjust again using the second figure; every other page type may be self-built on all three sides, under the conditions in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers and → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions; finally, check whether you need to read the matching industry edition.

Not built on the strictly regulated side	Build this type instead, for the same intent
Official pricing page (pt03)	Single-service price page, fixed price by tier (pt01)
List page (pt05)	Buyer's selection framework, list criteria without ranking (pt43)
Third-party single-business review (pt24)	Organisation facts page, checkable facts only (pt09)
Self-built reputation and credentials page (pt29)	Organisation facts page plus person profile page (pt09)
Verdict-first page (pt41)	Entity anchor page, persuade with facts (pt09)
Time-limited promotion page (pt42)	Standing fixed-price page (pt01)
Review aggregate page (pt28)	Being listed on review platforms and directories, with information matching the website (pt33)

Figure: page types the strictly regulated side doesn't build, each with a substitute for the same intent

The seven pairs above are the "don't build" page types the strictly regulated side runs into most often; the category list page (pt37) in the table below is also not built — build the organisation's own service catalogue instead; two more types aren't built on the strictly regulated side, outside the table: the case-law page (pt36, whose equivalent is pt32) and the market observation page (pt46, also not scheduled on the lightly regulated side), see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. If the strictly regulated side wants onto the list shelf, the only legal route right now is the criteria-based selection guide (pt43): list the criteria, don't rank, don't name (not measured; 65% confidence). Full three-side comparison for all 15 types:

Page type	Strictly regulated side	Lightly regulated side	Unregulated side	Switch E on (peer comparison banned)
Official pricing page (pt03)	Not built; build pt01 instead	May build after rewriting, as an all-in price page	May self-build	Not built; build pt01 instead

Page type	Strictly regulated side	Lightly regulated side	Unregulated side	Switch E on (peer comparison banned)
Comparison page (pt04)	May build after rewriting, comparing plans only, not businesses	May build after rewriting, comparing only plans and delivery formats, not businesses (Conservative line (not statute text))	May self-build	May build after rewriting, comparing only plans and procedures, not businesses
List page (pt05)	Not built (at the page-type level, Conservative line (not statute text); ranking is itself comparison, which runs into the Statute text ban on comparison — a sign-off cannot produce a publishable version)	Not built (Conservative line (not statute text); may be loosened with the client's written reasons)	May self-build; if you rank yourself first, disclose it	Not built
One question, one page (pt06)	May self-build; for symptom-type questions, build pt08 instead	May self-build	May self-build	May build after rewriting; for cost questions, do not write a peer range
Remedy and second-opinion page (pt08)	May self-build; the safety-first order may not be reordered	May self-build	May self-build	May build after rewriting; do not criticise the previous provider
Organisation facts page (pt09)	May self-build; never write superlatives, testimonials or star ratings	May self-build	May self-build	May build after rewriting; do not state your position relative to peers
Subsidy and limit rules page (pt17)	Official version: downstream citation only; your own version may be built after rewriting: the subsidy stays in an independent fact section, not in the same table as your own price and not adjacent to it, with no case-by-case confirmation tier	Official version: downstream citation only; your own version may be self-built	Official version: downstream citation only; your own version may be self-built	Not applicable

Page type	Strictly regulated side	Lightly regulated side	Unregulated side	Switch E on (peer comparison banned)
Third-party single-business review (pt24)	Not built (Conservative line (not statute text))	Not built (Conservative line (not statute text))	May self-build	Not built
Review aggregate page (pt28)	Not built (a rating is itself a testimonial, Statute text)	May build after rewriting, with testimonials attaching name, relationship and year	May self-build	Not built
Self-built reputation and credentials page (pt29)	Not built (testimonials and star ratings, Statute text)	May build after rewriting, testimonials must attach name, relationship and year	May self-build	May build after rewriting, stating only provable expertise, never implying standing relative to peers
Category list page (pt37)	Not built; build the organisation's own service catalogue instead	May build after rewriting, as the organisation's own catalogue, with names matching the registration	May self-build	Not built; build the organisation's own service catalogue instead
Calculator page (pt38)	May build after rewriting, without outputting a price commitment	May self-build	May self-build	May build after rewriting, stating "subject to formal engagement", not comparing price with peers
Verdict-first page (pt41)	Not built (Conservative line (not statute text); a verdict is itself a comparison)	May build after rewriting, as a scenario comparison table only, never writing "best" or "best value"	May self-build	Not built

Page type	Strictly regulated side	Lightly regulated side	Unregulated side	Switch E on (peer comparison banned)
Time-limited promotion page (pt42)	Not built (Statute text)	May self-build; if switch B (legally required fields) is on, write out the full validity period and conditions	May self-build	Not applicable
Buyer's selection framework (pt43)	May build after rewriting, removing the business-by-business comparison and keeping only the evaluation criteria	May self-build	May self-build	May build after rewriting, removing the comparison block

That the lightly regulated side doesn't self-build lists or reviews is, in the example industry, an internal conservative policy, not a written ban — hence labelled "Conservative line (not statute text)": the client's compliance officer can loosen it by giving written reasons (see → General Edition 0.4 The three labels for compliance sentences, and the stop-and-escalate rule). A business governed only by consumer-protection law, with no advertising-specific rule, falls on the unregulated side, and may self-build lists: the ranking needs a published method and checkable basis, and if you rank yourself first, disclose it.

Seven more types don't vary by side: the regulatory obligations and penalties page (pt10), the legislation text page (pt15), the register / approved list page (pt20), official replies and speech records (pt23), sentiment, forum and news pages (pt25), device and product regulatory documents (pt32), and the third-party directory listing (pt33). None of the three sides can self-build these; all you can do is make sure you're listed, fill in your own row completely, or build a downstream restatement page. For device and product regulatory documents (pt32), your own page's side-effects section must also match them point by point. Apart from pt36 and pt46 mentioned above, every other page type may be self-built on all three sides, under the conditions in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers and → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.

When switch E (peer comparison banned) is on, every block that lets the reader compare you with peers has to change:

Allowed before switch E (peer comparison banned) is on		Change to, once it's on
Price page gives the market range first, then your own price	→	Delete the market range; replace with an official-basis sentence
Comparison page compares peers' checkable, objective items	→	Compare only plans and procedures, not businesses
Self-build a list, disclosing if you rank yourself first	→	Don't self-build; go through a third party's normal review process
Verdict-first page gives a verdict	→	Don't build it; a verdict is itself a comparison
Selection framework page includes a business-by-business comparison	→	Remove the comparison block, keep only the evaluation criteria
Reputation page states awards and rankings	→	State only your own item, without implying standing over others

Figure: once switch E is on, every block that lets readers compare you with peers must change

The first row is the easiest to change incompletely: deleting the words "within the market range" doesn't count as done — the structure itself, "market range first, then your own price", is what's banned. In the example industry this splits into two tiers: comparison that names, or lets a reader identify, a specific business is Statute text, and even a sign-off cannot loosen it; an anonymous market range that names no one is Conservative line (not statute text) — but as long as the sample and basis can't be verified, a separate rule explicitly rules it out too, so in practice it can almost never be loosened. For how to write an official-basis sentence, see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers.

Questions like "who's best" are the easiest to assign to the wrong page type. Not every "who's best" question goes to a list page:

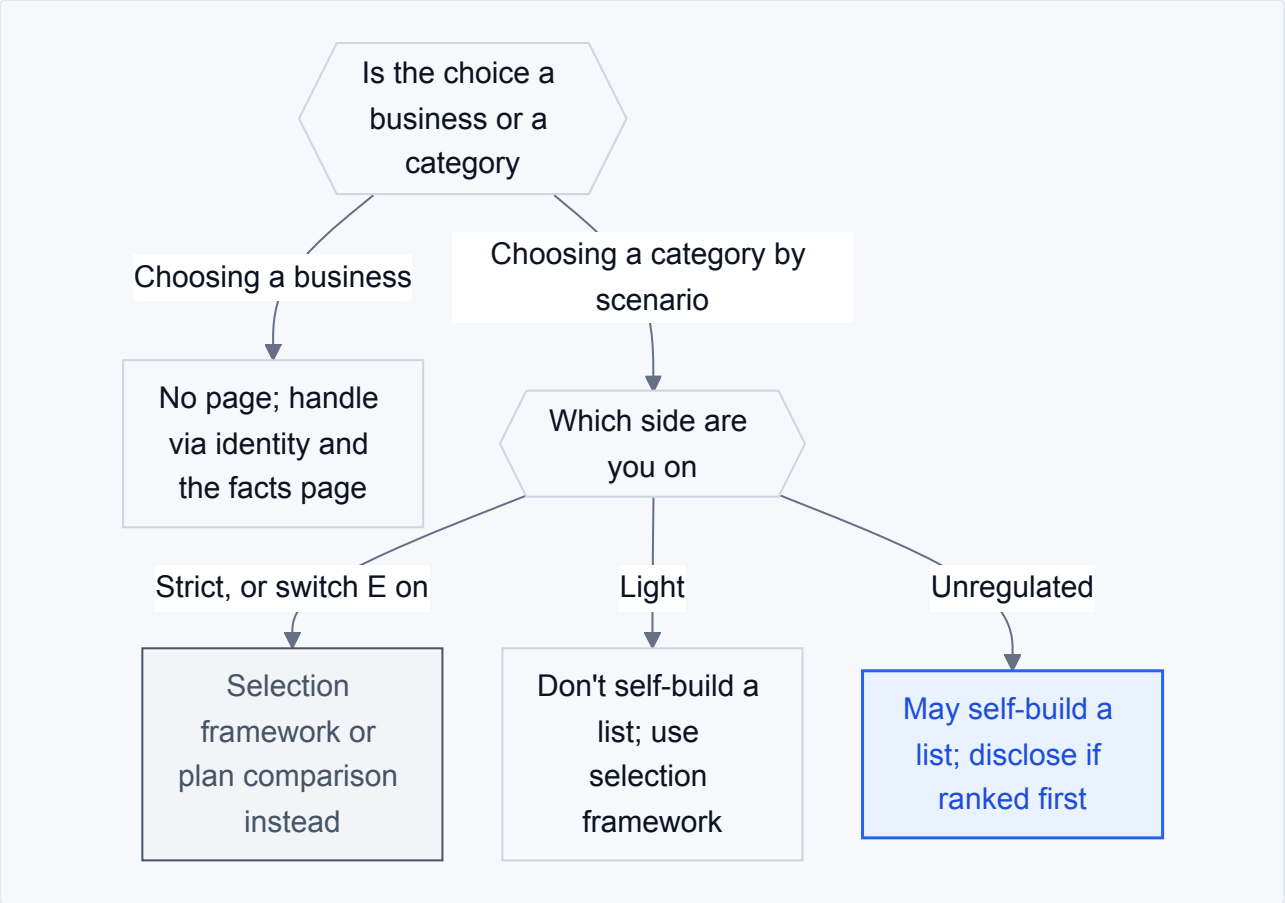


Figure: for "who's best", first split business choice from category choice, then route by side

No page is written for the "choosing a business" branch, because when buyers ask "who's best", neither leg cites a "best" article written by the business itself. The lightly regulated branch is Conservative line (not statute text); loosening it needs the client's compliance officer to give written reasons.

Industry editions published so far: *Dental & Aesthetics Edition* — for dental clinics and aesthetic clinics; *Law Firms Edition* — for Singapore law practices serving individuals and small businesses.

0.4 The three labels for compliance sentences, and the stop-and-escalate rule

What you'll do in this section: Understand what the three labels mean in the General Edition — each one states how firmly this pattern is backed in the example industry, not the law of your own industry; when you hit one of three situations, stop and escalate to the client's compliance officer for a written sign-off; remember a sign-off can only loosen situations ① and ②.

Items labelled [Compliance] or [Safety] are not judged by "does it move rankings": write them to the regulation and to personal safety, and where that conflicts with what's measured, follow the regulation — you may not loosen a rule just because "the cited pages write it that way too".

The labels in the General Edition state "how firmly this pattern is backed in the example industry" — not the law of your own industry. For your own industry, look up the actual rule text following → General Edition 8.2 Five steps to decide the side, and three criteria; until you find it, label it "Original text not obtained" across the board, and never present it to the client as statute; for how to write it in the meantime, follow the "when unsure, start on the stricter side" rule in → General Edition 0.2 Decide which side you are on first. That way you never pass off the example industry's rule as your own industry's law, and you never loosen the wording just because you "can't find it yet" either.

In three situations you must stop and escalate to the client's compliance officer for a written sign-off:

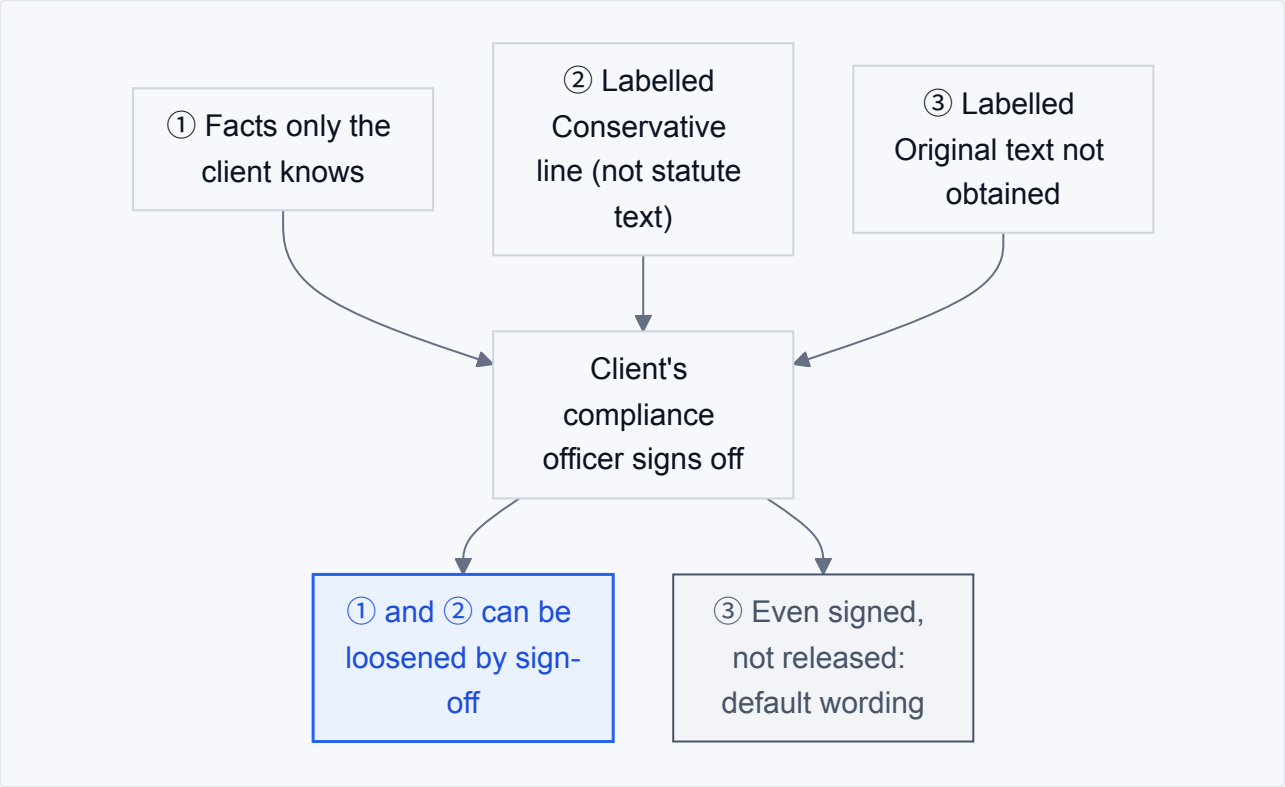


Figure: all three stop-and-escalate situations go to sign-off, but a sign-off can only loosen the first two

① Facts only the client knows: whether they've ever bought a review or ranking package, whether they have documentation on file for a claimed specialism, which country the certifying body is registered in, whether this price is the final, tax-inclusive price — only the client knows things like this; we can't judge them, so ask before writing. ② A judgement the handbook labels "Conservative line (not statute text)": this is an internal conservative policy, not a statutory ban; if the client wants to loosen it, have them give written reasons, and we write to what's signed off and keep it on file. ③ An item the handbook labels "Original text not obtained": a sign-off cannot supply the missing original text — **even with a sign-off, it is not released: write the handbook's default wording, and do not use it as a basis for judgement, assuming neither that the rule exists nor that it does not**; **never estimate an official figure, and never carry over an old value** — if you can't get the original table, don't write a number; write a pointer sentence such as "per the official schedule, see <link>", and update it once you actually obtain the original text. **A written sign-off resolves only these three stop-and-escalate situations,**

and it can only loosen ① and ②; it never turns something banned as Statute text into something publishable.

A sign-off is not a release from liability for us either: it is the judgement of the client's compliance officer, and it doesn't remove the joint liability of the party writing and publishing on the client's behalf (when switch A (agency liability) is on). Before writing, there are also five stop conditions: identity is unclear, the price cannot be made a final price, a number fails the six checks, the client hasn't authorised it, or the comparison bases differ — hit any one of these and escalate to the compliance officer; never fill in a number yourself. See → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers.

How a writer handles each of the three labels once they have it:

Label	How the writer handles it
Statute text	Write it as given; the ban has no exemption, and a sign-off cannot loosen it either
Conservative line (not statute text)	Write it as given by default; only the client's written reasons can loosen it
Original text not obtained	Not released even with a sign-off; write the default wording

Figure: the three labels, and how a writer handles each one

Why we're the ones writing these compliance rules: not because we know their industry better than the client — on the business and on the practice, they're the expert, not us. It's because a compliance rule decides **the layout's own cells**: which column can or can't carry a price range, which fields must be recorded next to each other, whether a promotion cell exists on the page at all — we're the ones who draw the cells, and if we draw one wrong the whole table has to be redone, usually after the body text has already been written to the wrong structure. The party writing and publishing on the client's behalf shares liability with the client — when switch A (agency liability) is on, that liability can't be avoided. So every rule only states "this is what this cell's wording is limited to"; it never gives legal advice; in the industry edition, every rule carries a rule number, the original sentence and a URL — a paraphrase doesn't count, and if it doesn't match, flag it; every ban must come with a matching "rewrite this cell as" entry, and a draft that deletes something and leaves the cell blank is sent back for insufficient density — a page with no exclusive checkable facts gives AI no reason to cite it.

Where a label comes from is never the writer's own call:

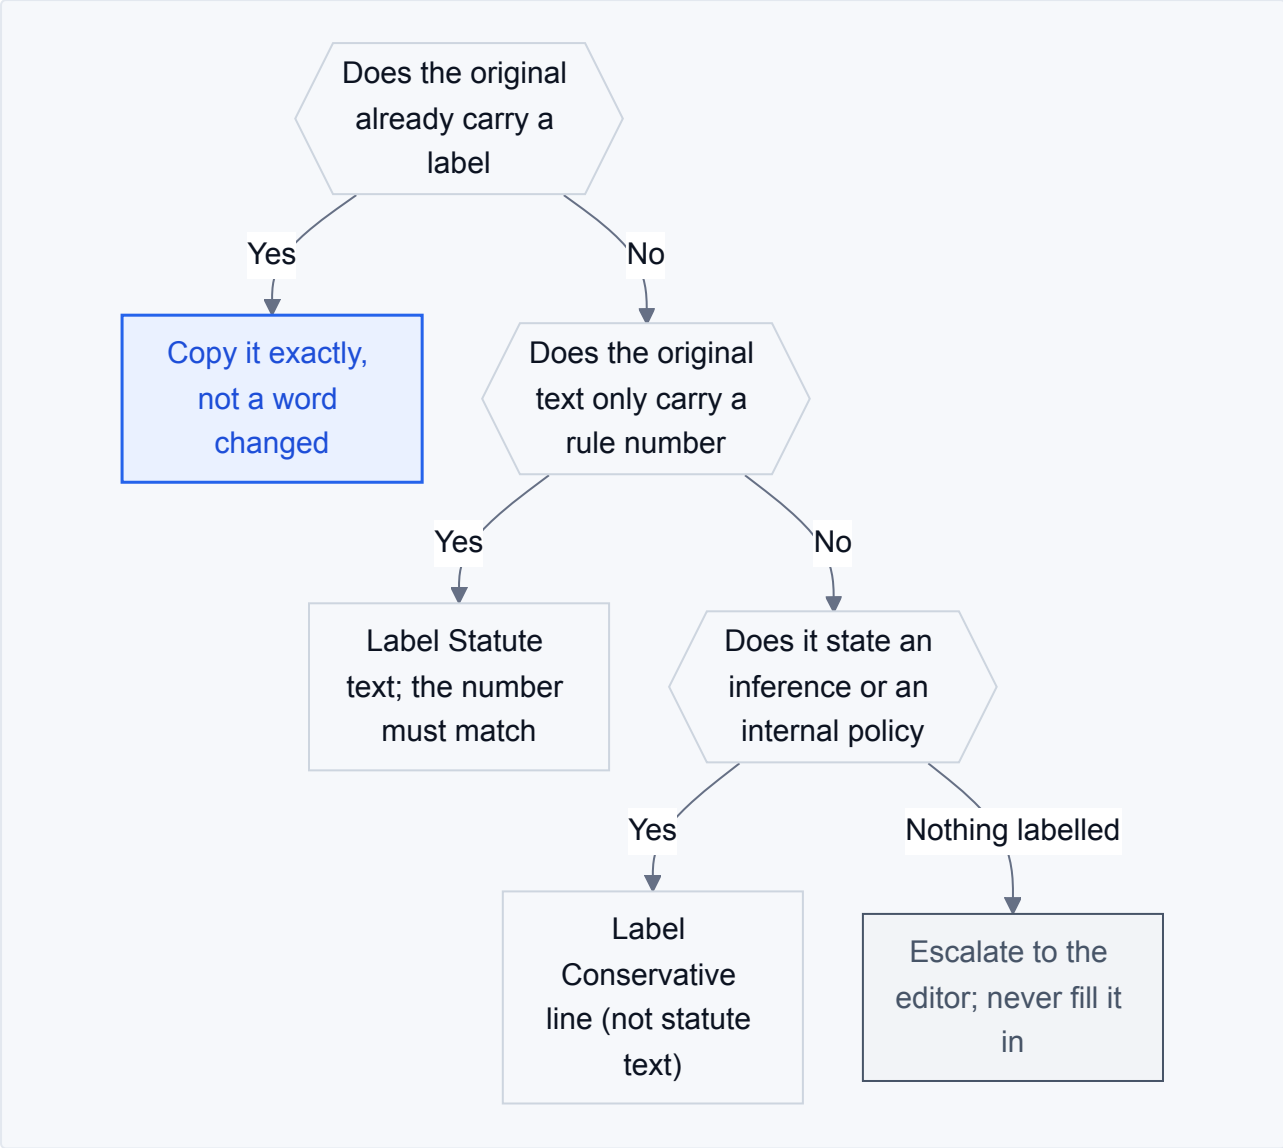


Figure: where a compliance sentence's label comes from; a writer may never add "Statute text" themselves

In the second box, "the number must match" means this rule number has to match the original text in the industry edition. Every compliance sentence in the book must ultimately carry a label, but the label can only come from the original text: a large number of original sentences only carry a rule number, or carry nothing at all, and the moment a writer fills the gap themselves to complete the cell, it's easiest for an inference to get written up as "Statute text". This is the step with the highest risk of lying. **Any ban with no original rule sentence behind it is labelled Conservative line (not statute text) across the board; a writer may never add "Statute text" themselves.**

For how to explain these three labels to a client, see → General Edition B.2 Approved wording and scripts.

Chapter 1

How AI picks its answers: what buyers ask, and which pages ChatGPT and Google's AI cite

AI reads pages live, the two engines cite different pages, and three things decide whether you get named

1.1 How buyers ask, and what AI reads

1.2 Two legs: ChatGPT looks for the source, AI Mode for second-hand summaries

1.3 Three gates and three paths

1.4 Three kinds of things not to do

1.1 How buyers ask, and what AI reads

What you'll do in this section: understand where the paragraph AI gives the buyer comes from. AI fetches a batch of pages live, reads only the text visible on those pages and does not execute JavaScript. After reading this, you'll know why the later chapters fix the door first, and why the main battlefield is not your website.

> **Example** (dental) Buyers no longer scroll through ten blue links. They open ChatGPT and ask directly: "新加坡哪家做隐形矫正好？" (which clinic in Singapore is best for clear aligner treatment?) AI replies with a paragraph that names a few clinics.

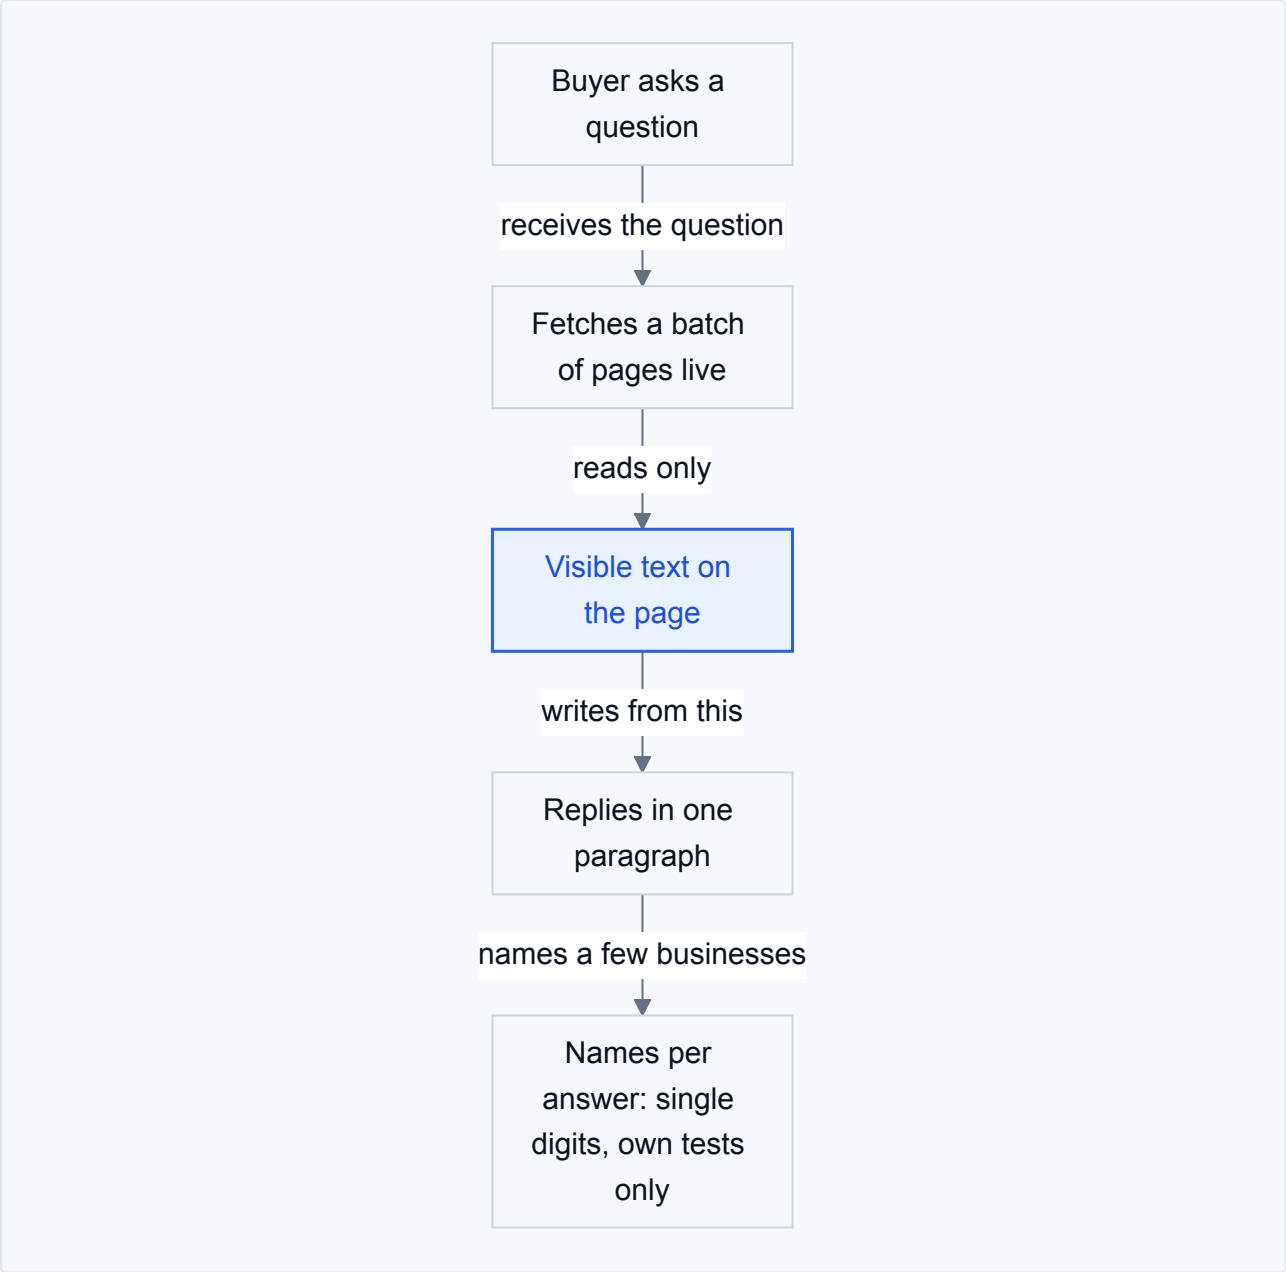


Figure: AI does not answer from memory. It fetches pages live and reads only visible text; the businesses the answer names make up its total names

How many names an answer holds is not something someone else's industry report can tell you. Trust only the number you count when you measure your own set of questions. On the same set of questions, some answers name 3 businesses, some name 8, and some name none at all (they only explain the basics and tell you to "consult a professional").

After the buyer has a name

① **See the names**
The answer names a few businesses; about 1% of visits click the links in the summary (US sample)

② **Search again**
They skip the links and search the name directly

③ **Go to the website**
Check price, process, credentials

④ **Check profiles**
Reviews, directories, business profiles

⑤ **Get in touch**
This is where the sale happens

Figure: Buyers rarely click the links in the answer. Once they have a name they go back and check it themselves, so pick "who to buy from" questions, not educational questions

- Educational questions often produce answers that name no business at all. Even if the answer cites you, the buyer has no name to go back and search for.
- Steps ②–④ affect the conversion rate, not exposure: a wrong address, price or business status stated by AI is what the buyer runs into at these steps. This is one reason identity comes before writing pages; see → General Edition 3.1 Why this comes before writing pages · The two-hour checklist.
- The sample and source for "about 1%" are in → General Edition A.1 Evidence for mechanics, the door and identity.

Where citations come from: the main battlefield is not your website

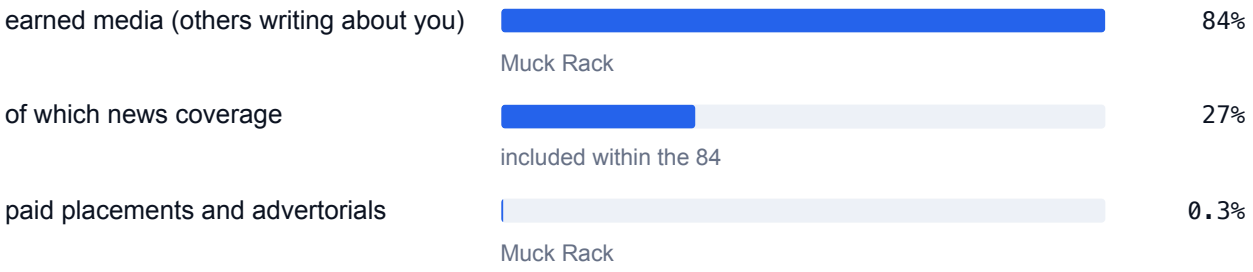


Figure: Most AI citations come from pages other people write about you; paid placements and advertorials are close to zero

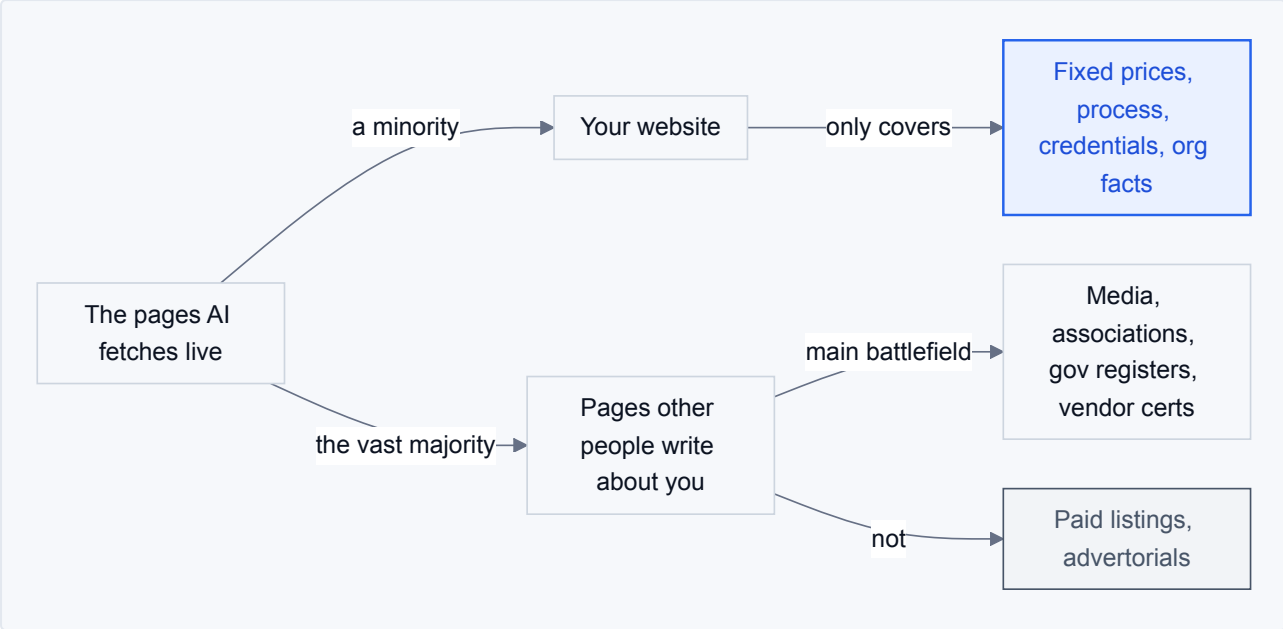


Figure: Your website only covers the cells third parties cannot fill; the main off-site battlefield is media and institutions, not paid listings

- Besides covering these cells, your website has one more job: supplying original facts that other people can copy.
- Even when the buyer searches your brand name, owned content accounts for only 23% of AI citations (Omniscient; its denominator is different from the bar chart above, so the two numbers must not be compared side by side).
- So when you record the shelf, what you record is not "am I there" but **which containers currently hold the seats for this question**. For how to split off-site effort, see → General Edition 6.1 Where the effort goes, and two numbers that must not sit side by side.
- The samples and how the numbers were counted in both studies are in → General Edition A.1 Evidence for mechanics, the door and identity.

What AI reads: visible HTML, no JavaScript

Do not execute JS (zero evidence of execution in a large sample)	Render pages (exceptions)
GPTBot, OAI-SearchBot, ChatGPT-User	Gemini: fully rendered through Google's infrastructure
ClaudeBot, Claude-User	Applebot: a browser-style crawler
PerplexityBot, Bytespider	—
Meta-ExternalAgent	—
Price written only in JS: not a single character comes through	Price written only in JS: it comes through

Figure: Mainstream AI crawlers do not execute JavaScript; only two exceptions render pages

- **This exception is useful: if the Google leg has seats and the ChatGPT leg has zero, check CSR first, not your choice of questions.** For how to check, see → General Edition 2.9 Door-layer troubleshooting and this chapter's acceptance checks.
- Structured data that is not visible on the page is not read either. When the five main systems (ChatGPT / Claude / Perplexity / Gemini / AI Mode) fetch live, not one of them reads JSON-LD; hidden Microdata and RDFa are ignored too, and every test price was taken from visible HTML. This only shows that "schema does nothing on the live-fetch path". It does not follow that "schema is useless". For where schema belongs and how far to build it out, see → General Edition 2.7 Gates 3 and 4: indexing paths, and JSON-LD sealed once.
- The samples and dates of both tests, and how many prices each system retrieved, are in → General Edition A.1 Evidence for mechanics, the door and identity.

Three things later chapters take as given

1. Before answering, AI fetches pages live and reads only visible HTML: it does not read JSON-LD, and mainstream retrieval crawlers do not execute JavaScript (the exceptions are Gemini and Applebot).
2. Most citations come from pages other people write about you (earned media, 84%); paid placements and advertorials make up 0.3%. Your website only covers the cells third parties cannot fill.
3. The order is door × identity × shelf (1.3), and progress is measured only with the three rulers plus the noise band from same-week retests (1.3, → General Edition 7.1 What a retest produces, and the rulers). The other premises (citation density by page type, the correlation coefficient for video, the

share of price questions that trigger AI Overviews, annual opportunity value) are given in the section where each is used.

1.2 Two legs: ChatGPT looks for the source, AI Mode for second-hand summaries

What you'll do in this section: tell apart which pages each leg looks for and which ruler measures each one. After reading this, you'll have one more habit: before you write any sentence, ask "which leg is this sentence for?"

ChatGPT leg	AI Mode leg
Looks for the source: government pages, official pricing pages, product pages	Looks for second-hand summaries: peer articles, lists, videos, Google cards
Takes the source facts and writes its own conclusion	Cites pages other people have already summarised
Feeds on own sentences: you as the subject, exact amount and unit	Feeds on market sentences: range + variables + source and date

Figure: The two legs do not look for the same kind of page, and they want different kinds of sentences

In our measurements, the direction was the same in all five industries. For what each leg cites and in what share, industry by industry, see → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types.

> **Example** (B2B SaaS) Official pricing pages make up about 53% of ChatGPT's citations, plus about 7 citations of help centres, and then ChatGPT writes its own conclusion. For AI Mode, YouTube, Reddit and LinkedIn together make up 40%, and the rest are peers' "vs" pages and lists.

- The last row of the figure is the direct takeaway for writing pages. ChatGPT also uses market sentences when it writes an overview, so write both kinds of sentence on every page, each in its own place; the spec is in → General Edition 5.6 Two legs, two kinds of sentence.
- **In regulated industries, market sentences are written differently, and some may not be written at all.** First check the quick reference for your side in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers. Do not copy sentence patterns straight from this figure.
- Local questions draw on business-profile sources, and IndexNow is not a switch for the ChatGPT leg either; both points are covered in → General Edition 2.7 Gates 3 and 4: indexing paths, and JSON-LD

sealed once.

Three rulers

The figure above answers "which pages to write, which sentences to write"; the one below answers "what to measure with". In the page-type tests, the AI Mode leg is used only to see which kinds of page it cites. It is not used as a monthly ruler.

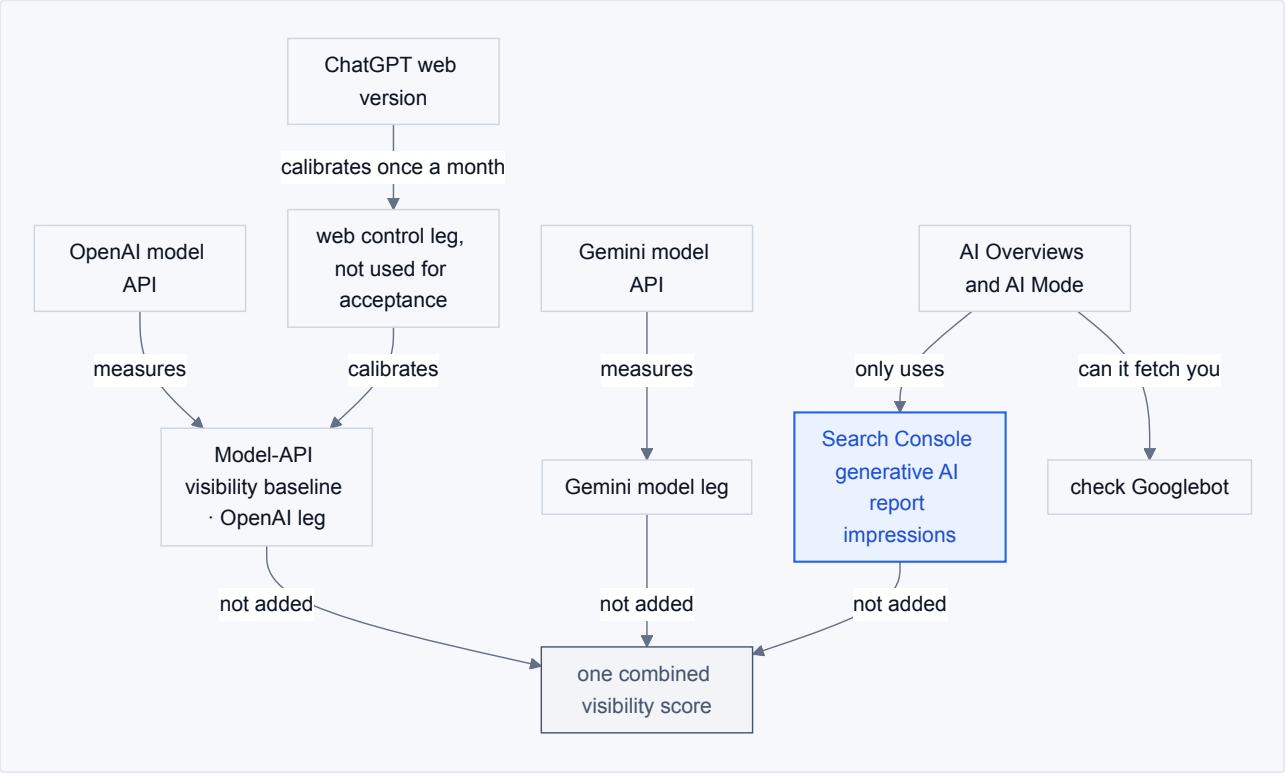


Figure: Three systems, three rulers: each measures its own thing, and they are never added together

- **In reports, the ruler legs may only be called this**: "Model-API visibility baseline · OpenAI leg" and "Gemini model leg", printed in the footer of every report. **Never write "what users see in ChatGPT" or "real user visibility".** Citation distributions on the API and on the web version show a measured, systematic skew. It is not noise, and running more rounds does not smooth it out. For how large the skew is and how to run the web control leg, see → General Edition 4.4 The web control leg and the frozen baseline (the book's only full spec).
- **The naming rule is not about wording**: get the name wrong, and a whole quarter's off-site effort gets aimed at the API leg's top 20, while buyers may be seeing a different set of pages.
- **Gemini API ≠ AI Overviews ≠ AI Mode.** They are three different products, with different fetch paths, different ways of generating answers and different source pools. Passing off Gemini model leg numbers as visibility in AI Overviews means measuring A and using it to sign off B.
- The Search Console ruler has impressions only, with no times named and no position. It is already included in total impressions, so it must not be added to total impressions.

- For why you check `Googlebot` rather than `Google-Extended`, see → General Edition 2.4 What each crawler is for, and judging the door leg by leg (read-only).

1.3 Three gates and three paths

What you'll do in this section: memorise the order of door × identity × shelf and the acceptance ruler for each stage, and know that only three paths can move your ranking. After reading this, you'll be able to tag every action on your plate ①, ② or ③ and delete whatever you cannot tag. You'll also be able to judge for yourself whether you are winning or losing this month.

Three gates: the first two are switches, only the third adds points

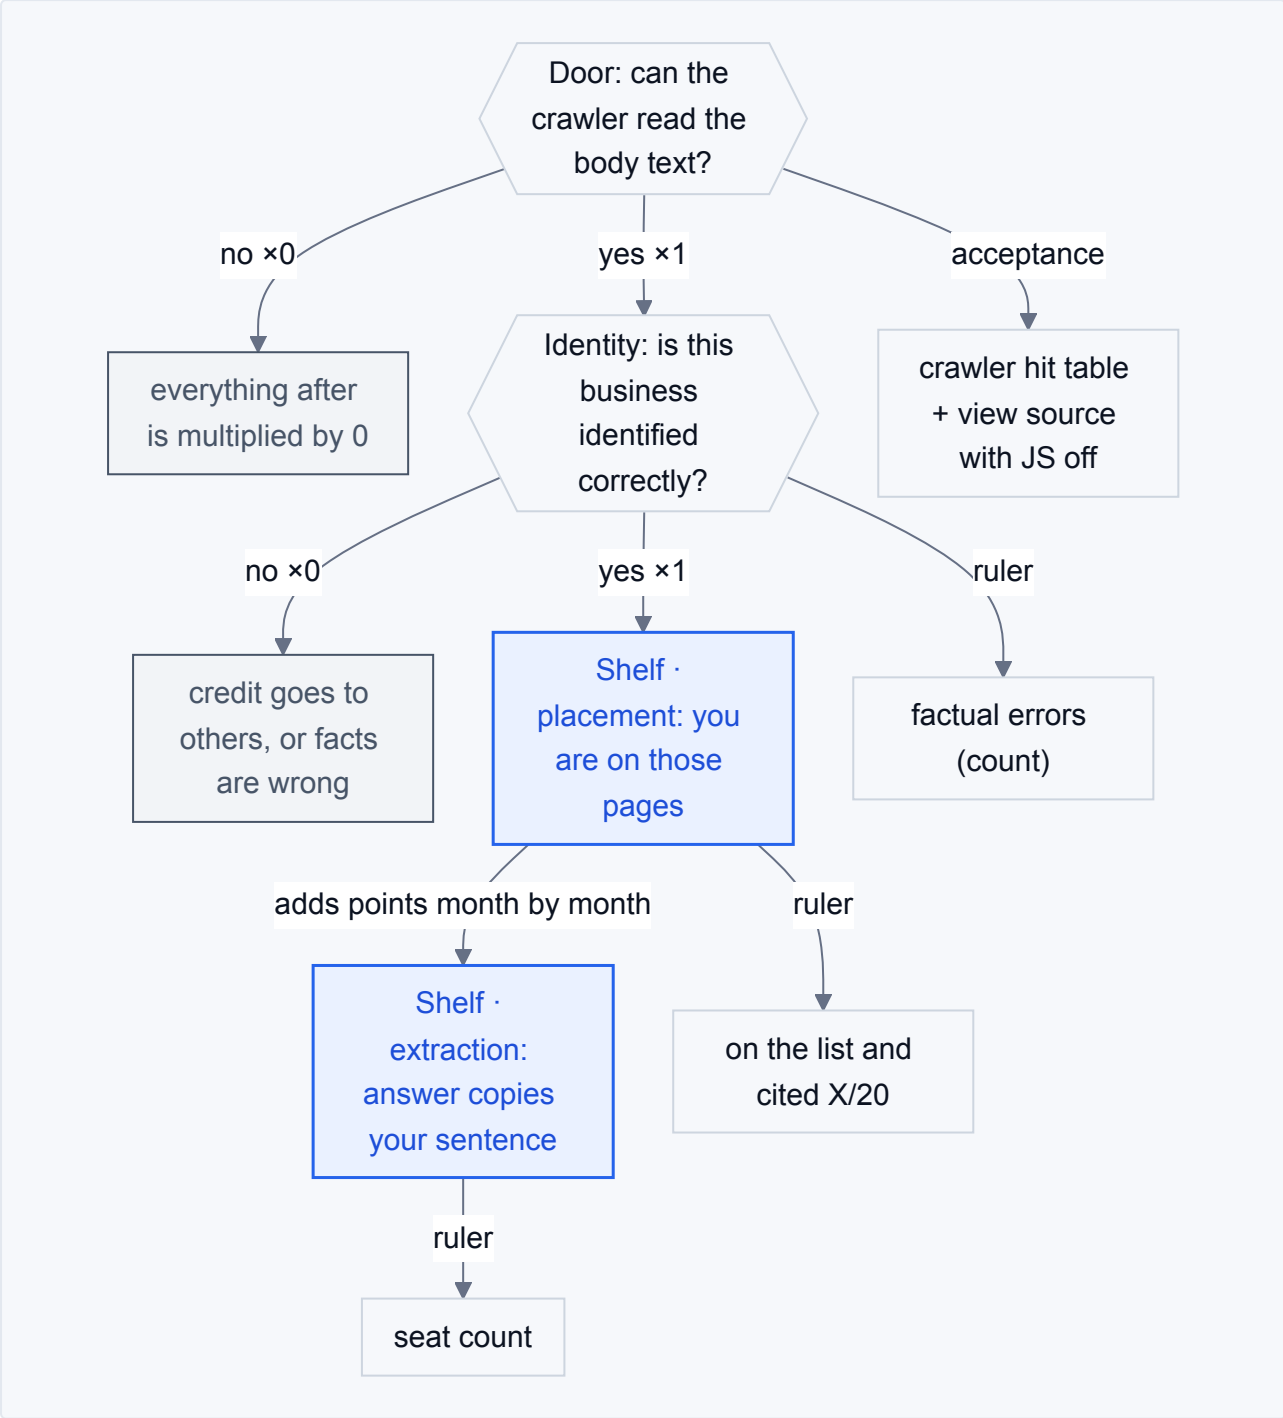


Figure: The door and identity are ×0/1 switches; only the shelf adds points month by month. Each stage has its own acceptance ruler

- **Why the order cannot be swapped**: the typical result of reversing it is spending three months on the shelf (writing pages, sending outreach letters), then finding on day 91 that the door was closed the whole time. Everything produced in the first 90 days goes back to zero and has to be redone.
- **Why the shelf is split into two stages**: when you are listed as a source but not written into the answer, adding pages does not help; you need to change the sentences. Without the "extraction" stage, the people doing the work will just keep adding pages.

> **Example** (e-commerce) A category page with numbers but no sentences most often gets stuck at the "extraction" stage: you are on the page, but the sentence the answer copies into its body is not yours.

What counts as a win on each of the three rulers

Ruler	What it measures	What counts as a win
Seat count	How many times you are named on the frozen questions	The increase is larger than your own measured noise band, and the direction holds for two months running
On the list and cited X/20	Of the 20 third-party pages AI most often cites for this set of questions, the number that include your name and were actually cited this month	The baseline is usually 0–2; +1 a month, ≥ +3 at 90 days
Factual errors (count)	How many things AI gets wrong about your address, price, business status and credentials	Down to 0, with before-and-after screenshots for each one

- The three rows sit side by side: never merge them and never add them together. The full definitions and how to run them each month are in → General Edition 7.1 What a retest produces, and the rulers; how to measure the noise band is in → General Edition 7.2 The noise band, the page-level signal and the ten monthly steps.
- **If you watch only one number, watch the factual error count.** AI copies factual errors from the source pages it retrieves, so when you fix the source, both legs change with it. Of all the rulers it is the least sensitive to the difference between the API leg and the web leg, and the only one in the whole set that is not contaminated by the sampling-leg skew.

Only three paths move your ranking

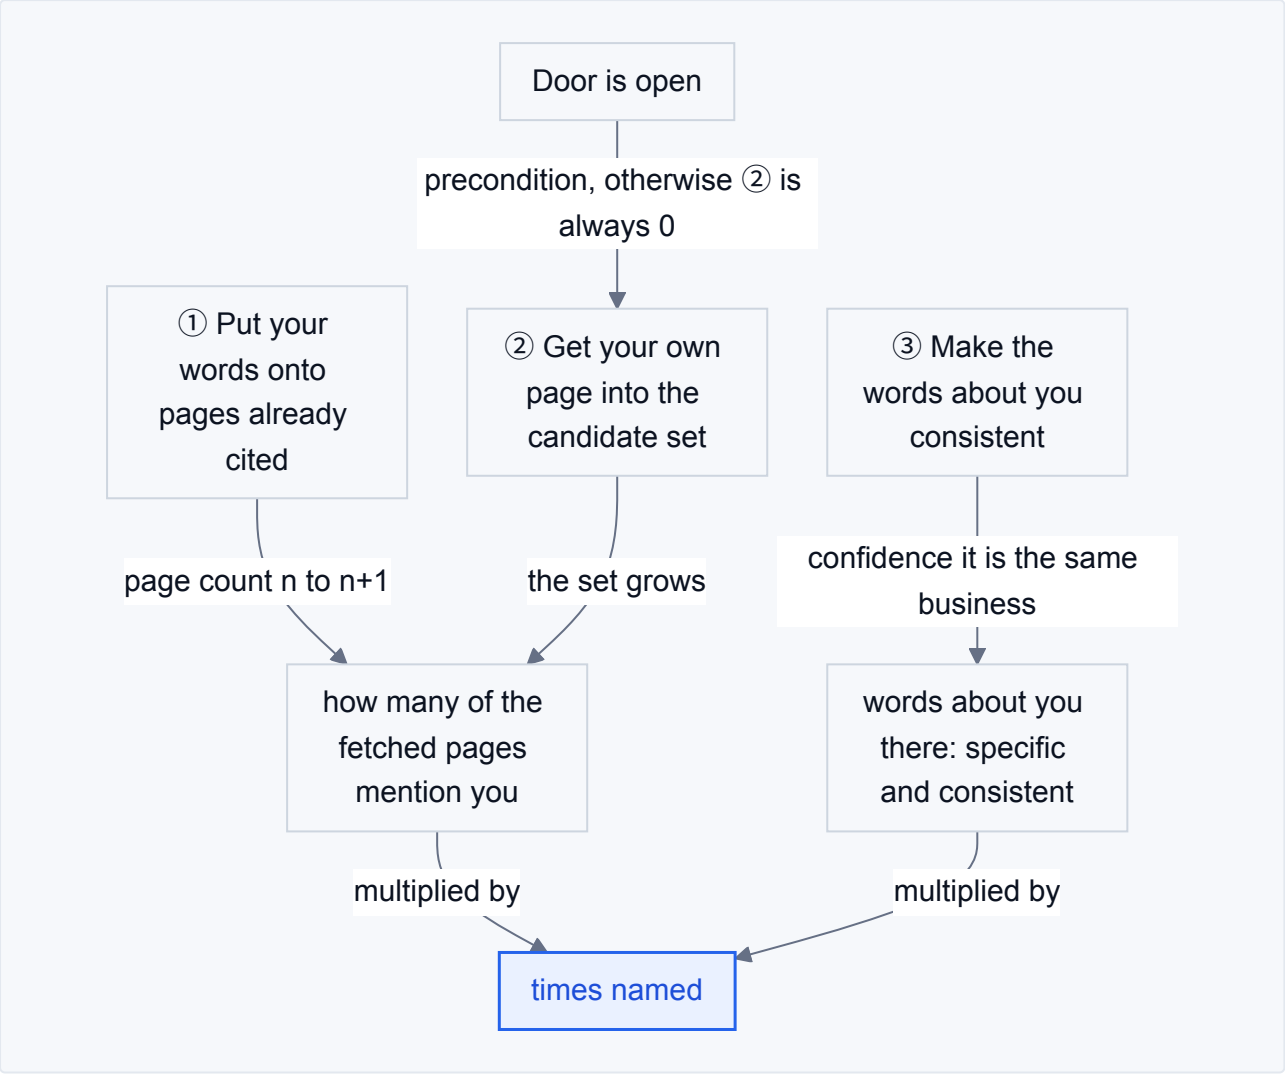


Figure: Times named is the product of two factors, so only three paths can move it; there is no fourth

Path	What the action looks like	Which stage it lands in
①	Update letters, correction letters, bylined contributions, listings, platform profiles, official registers	Shelf · placement. The candidate set stays the same; **the shortest causal chain**
②	Door + high-density page type + checkable numbers + sources + a visible update date	Door + shelf · placement. The whole page that joins the set is about you
③	Unique spelling, `/facts`, person pages, six-trace word-for-word alignment, corrections	Identity. Adds no pages; raises the confidence that "these pages are about the same business"

- The shelf · extraction stage does not need another path. It depends on how the sentence is written: checkable numbers, no adjectives, a source.

- The feed channel that sends product data to the shopping shelf does not go through the door, so count it as a variant of ②. Flag this difference separately. If you don't, you get one of two mistakes: "the door is broken, so the feed must be useless too", or "the feed is on, so there's no need to fix the door". For how to align product pages with the feed, see → General Edition 5.20 Entity anchor pages: three subtypes, and the product detail page (pt09).

How to tell you are losing



Figure: Pages on the list can rise while you are still losing. Check cited first, then whether seats have cleared the noise band

- **This is what typical losing looks like**: three months have passed, 8 pages have been written, 40 outreach letters have gone out off-site, and pages on the list have risen from 9/20 to 13/20. Every

table is filled in, but the seat count wobbles inside the noise band, and nobody can say whether it has moved at all, let alone why it hasn't.

- "The door is not really open": you think crawlers are let through, but the WAF is actually blocking them, or the price exists only in JS. "The ruler measured wrong": using one sampling gap as the threshold for a different sampling gap, adding passes partway through, or describing an API-leg reading as "what users see".
- **The invisible way to lose** is harder to spot. If you watch only "on the list", the whole loop can close out normally while rankings do not move at all, and after three months not one number can answer "did AI actually cite the pages we got our words onto, or not?". For how to triage when nothing has moved, see → General Edition 7.3 Not-moved triage and next month's three points.

Two disciplines

1. **Tagging discipline**: every action must be tagged ①, ② or ③. **Delete any action you cannot tag**, however much it looks like marketing work. Writing proposals, building competitor analysis tables, rewriting the brand positioning statement, posting social media image posts: none of these can be tagged, so none of them belong in this playbook.
2. **Confidence discipline**: low-confidence items still go into the playbook, with "**not measured**" at the end of the line; never leave something out just because there is no evidence. But the mechanism must be spelled out. An action you cannot tag ①, ② or ③ is not "not measured"; it "**does not hold**". Never mix the two up.

1.4 Three kinds of things not to do

What you'll do in this section: check your current plan against the three figures below. Delete the actions that do not move rankings, replace the ones that backfire, and fix the wrong criteria that would leave you unable to find the cause. The third kind is the most expensive: it has you put another three months into the wrong direction.

Things that do not move rankings

Do not do		Do instead
Write llms.txt	→	Do not write it
Write JSON-LD page by page and check it word for word	→	Set up a template once when building the site; leave it off the per-page checklist
Treat sameAs as the main entity lever	→	Use a visible-HTML dl with one field per line, plus clickable links
Stack FAQPage schema, add self-rated scores	→	Question-form H2s + a short Q&A at the foot of the page
Chase view counts, buy views, do flashy edits	→	Produce reference material that can be cited
Make paid listings, advertorials and directory submissions the main line	→	earned media + institutions, government, associations
Keep a six-category ledger of citation sources	→	Assign tasks by the eight page-type categories
Notarise screenshots and run permutation tests during the baseline period	→	Save the raw answers to disk + the noise band
Put pure symptom questions into the question pool	→	Use the symptom + option + price three-part form
Build a separate page for every phrasing of the question	→	Fold the variants into H2s on the same page

Figure: Ten things that do not move rankings, each with a replacement to do instead

- **llms.txt**: take it out of the prediction model and citation frequency is predicted more accurately, not less; it adds noise, not signal. Google's AI optimisation guidance also says outright that AI Overviews / AI Mode do not need it. It is the biggest fake GEO lever of 2026.
- **Page-by-page JSON-LD**: after JSON-LD was added, AI Overviews (AIO) citations changed by −4.6%. That was the only significant result, and it was negative; the other two platforms showed no significant change. For how to seal the template once, see → General Edition 2.7 Gates 3 and 4: indexing paths, and JSON-LD sealed once.
- **sameAs**: live fetching does not read JSON-LD (1.1), so however many registration numbers, association pages or Wikidata QIDs you chain together, the benefit on the path "AI reads this page live" is zero. Write the schema anyway, but note it as "indirect, Google leg only".
- **FAQPage and self-rated scores**: the former is built for rich results and does not feed into AI citations; the latter, on the strictly regulated side, also runs into the ban on testimonials and ratings

(Statute text in the example industry).

- **View counts**: views, likes and subscriber counts are almost uncorrelated with citation frequency. For the kind of long video to make instead, see → General Edition 6.7 Long videos narrated by the named expert.
- **Paid listings**: they are the shortest bar in the bar chart in 1.1. Put the same spend into earned media and institutions, government and associations instead, and the difference is an order of magnitude.
- **The six-category ledger**: the six buckets do not produce a single action. Effort allocation reads the page-type column, and its eight categories are: comparison / listicle articles / category directories / profiles / government registers / review sites / communities / own website. **The categories are not a ledger; they are the basis for assigning tasks.**
- **Notarised screenshots, permutation tests**: notarisation is valuable for comparing "after" with "before", and during the baseline period there is only "before". At n ≤ 10, a permutation test does not change any decision.
- **Pure symptom questions**: for pure symptom questions, AI cites public-health websites and encyclopaedias and almost never names a practitioner.
- **Phrasing variants**: a new page only dilutes the signal for the same intent cluster; see → General Edition 5.5 The page unit: one intent cluster, one page.
- The samples and sources behind each point are in → General Edition A.1 Evidence for mechanics, the door and identity and → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests.

Things that make it worse

Do not do (it backfires)		Do instead
Append an allow group to the end of robots.txt	→	Copy the wildcard group's Disallow lines, one by one, into every named group
Rewrite an old page	→	Only expand it: leave the URL, the H1 and the paragraphs carrying ranking keywords untouched
Edit Wikipedia's body text to add yourself	→	Only add checkable sources to an existing entry; do not add your organisation's name
Create a Wikidata entry with no third-party material	→	Create one only once you have a news report, a bylined journal article or an association announcement
On a regulated side, actively ask for reviews, buy a list spot or write a "from" price	→	Check 5.2–5.3 for your side first
Send out material that names peers	→	Keep it for your own eyes only
Change the question, the engine or the number of passes partway through	→	Leave what's frozen alone

Figure: Seven things that make it worse, and what to do in their place

- **robots**: the matching rule is "the most specific group wins, and every other group is ignored". Append a group directly and the named crawlers break away from the wildcard group completely: all the protection you had on the admin area, shopping cart and site search stops applying to them. For the merge steps, see → General Edition 2.6 Fixing the door: nosnippet, the four-step robots merge, the WAF allowlist.
- **Rewriting an old page**: an old page has history, links and existing rankings. The expansion cap is relaxed to double the original length. Before editing, export the page's top 10 queries from Search Console, circle the matching paragraphs one by one, and do not delete them. See → General Edition 5.5 The page unit: one intent cluster, one page.
- **Wikipedia**: self-referencing edits get reverted and leave a trace in the edit history, so the risk outweighs the benefit. See → General Edition 3.4 Third-party credential tiers, Wikidata and the five business profiles.
- **Wikidata**: it is not "no threshold". Admission requires that the subject "can be described using serious and publicly available reference material". Create an entry without that and the community will nominate it for deletion. See → General Edition 3.4 Third-party credential tiers, Wikidata and the five business profiles.

- **The two regulated-side rows** are not an efficiency issue; they are a **compliance issue**, and "disclosing that it is paid" is not a release from liability. **Comparison material that names peers is never sent out.** For whether and how each cell may be written, check → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers and → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions for your side. For the stop-and-escalate and sign-off rules when you are unsure, see → General Edition 0.4 The three labels for compliance sentences, and the stop-and-escalate rule: **a written sign-off resolves only the three stop-and-escalate situations; it never turns something banned as Statute text into something publishable.**
- **Changing the ruler**: change the question, the engine or the number of passes, and the before and after numbers can no longer be compared; see → General Edition 7.5 Ruler discipline: change the ruler and nothing is comparable.

Things that hide the cause

Wrong criterion		What to check instead
curl with a spoofed crawler UA, check what comes back	→	Use a browser UA to check content; use a crawler UA only to check the status code
Treat IndexNow as a switch for the ChatGPT leg	→	Paste the URL and have ChatGPT quote the price line back word for word
Check whether GPTBot is blocked	→	Check OAI-SearchBot
To judge AI Overviews, check Google-Extended	→	Check Googlebot and nosnippet
Add the two legs together or combine them into one score	→	Keep the three rows side by side; do not add them, do not use one to confirm another
Describe an API-leg result as what users see	→	Qualify it as "on the model API"
Cite "69% of AI crawlers do not execute JS"	→	Zero execution evidence; the only exceptions are Gemini and Applebot
Let bare category words make up most of the question pool	→	Location ≥60%, price words ≥30%, bare words ≤10%
Run it once and declare zero visibility	→	A single run on a single engine, n ≤ 3, is directional evidence only
Skip triage because pages on the list have risen	→	If seats have not cleared the noise band, always start triage at layer 1

Figure: Ten wrong criteria that leave you unable to find the cause, and what to check instead

- **Why each of the first four rows is wrong** is explained in one place only: the curl row in → General Edition 2.3 Gate 2: four identities, fetched live (read-only), where it is downgraded to a troubleshooting tool, not a criterion; IndexNow in → General Edition 2.7 Gates 3 and 4: indexing paths, and JSON-LD sealed once; GPTBot and Google-Extended in → General Edition 2.4 What each crawler is for, and judging the door leg by leg (read-only).
- **Adding them together**: the three rows measure three different systems, with different denominators and different failure modes. Add them up and any one of them getting worse can be hidden by another (1.2).
- **The API leg**: citation distributions on the API and on the web version have a systematic skew, so wherever you only have API-leg evidence, always qualify the statement as "on the model API" (1.2).

- **69%**: this figure is second-hand (reprocessed), and the actual finding is stronger; see 1.1.
- **Bare category words**: the wording of the question pool can produce a fake zero visibility. For how the ratios are fixed, see → General Edition 4.2 The question pool: where the 30 questions come from, how they are balanced, how they are signed; the comparison data behind this is in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests.
- **Pages on the list have risen**: pages on the list are the output of your own work, so +1 a month is close to the norm. Whether they have risen has nothing to do with whether you need triage; see → General Edition 7.3 Not-moved triage and next month's three points.

Chapter 2

Let AI read your website: check first, then fix the settings

Access logs, robots.txt, firewalls and nosnippet; save how things stand before you change anything

2.1 The five gates at a glance, and gate 0: access logs (read-only)

2.2 Gate 1 read-only audit: what to check in nosnippet, robots and the WAF

2.3 Gate 2: four identities, fetched live (read-only)

2.4 What each crawler is for, and judging the door leg by leg (read-only)

2.5 Gate 3, read-only: connect the two free AI reports

2.6 Fixing the door: nosnippet, the four-step robots merge, the WAF allowlist

2.7 Gates 3 and 4: indexing paths, and JSON-LD sealed once

2.8 The launch gate for every page

2.9 Door-layer troubleshooting and this chapter's acceptance checks

The door is a multiplier: if crawlers cannot get in, however many pages you write afterwards get multiplied by zero. With `nosnippet` switched on, the Google leg is counted as zero for the whole quarter; put the price in JS, and mainstream AI crawlers get not a single word of it. This whole chapter is free and takes half a day to a day, but you do it in two passes. 2.1–2.5 is the D1 read-only audit: get access, record the current state, and change nothing. Only after the baseline has been saved, following → General Edition 4.4 The web control leg and the frozen baseline (the book's only full spec), do you do 2.6–2.9 and fix the door. For how the three gates — door, identity, shelf — relate to each other, see → General Edition 1.3 Three gates and three paths.

Hard gate: until the baseline is saved, do not change the door, any profile or the facts page.

2.1 The five gates at a glance, and gate 0: access logs (read-only)

What you'll do in this section: first see clearly why the door is done in two passes; on D1, the first move is to send the access request, pull the *Crawler hit table*, and use the four ways of reading it to log problems as to-fix items. When you're done, you'll have a hit table and a to-fix list, and you won't have touched the door.



Figure: audit read-only first, freeze the baseline, then fix the door — only then are before and after comparable; reverse the order and what you freeze is a shelf you've already changed yourself

Why it has to be this way: what the baseline records — page-type mix, the four states per URL, question shape — is all measured on the shelf "while the door is still as it was". Fix the door first and freeze second, and your own changes are already mixed into the before sample; after that, whenever something "went up", you can't tell whether the door did it or it would have risen anyway. So the build order (door first) and the measurement order (freeze first) are two different things — never mix them. On D1 there are only three kinds of thing you can do that do not count as a change: request access, turn on logging, record the current state. The book's week numbers all count from the day you fix the door; see → General Edition 0.1 The book in one sentence, and the 90-day reading order.

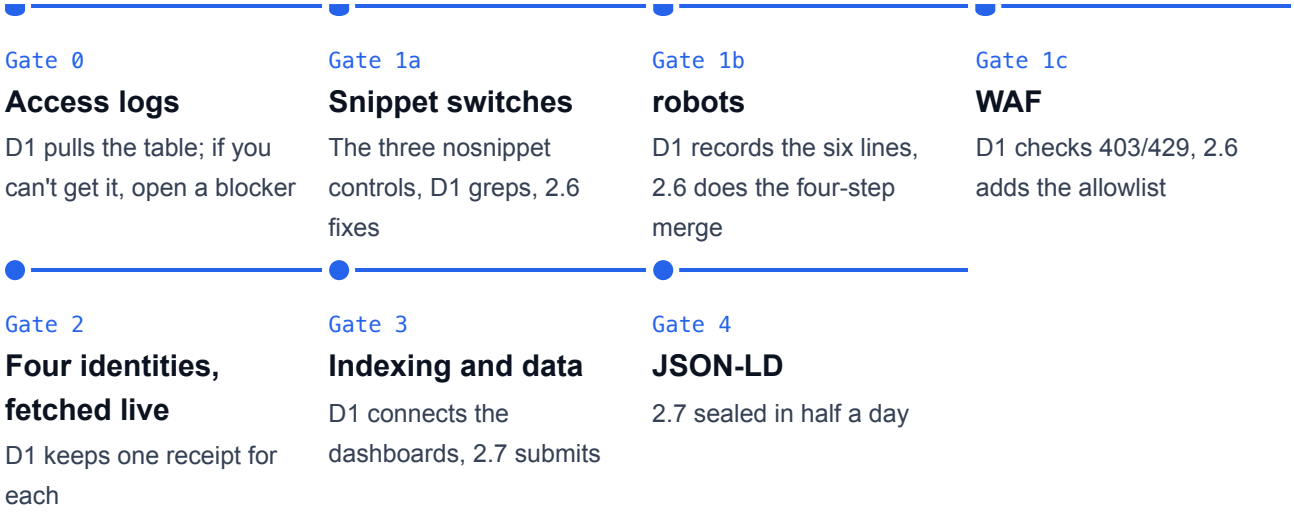


Figure: go through the five gates in the same order twice — D1 only looks, changes nothing; once the baseline is saved, fix them in this same order

Two ordering rules:

1. **Gate 0 comes before every other gate.** It is the only direct evidence, and the only piece of work where "the doing takes half an hour, but getting access can take days". In your first email, ask for all three kinds of access at once: **domain verification, permission to edit robots, and permission to export access logs.**
2. **Gate 3 connects the free first-party data before you spend money on probes** (see 2.5). Do it the other way round and you'll pay to "discover" something that was already free.

How to pull gate 0

Faking a UA, a `site:` query, eyeballing it in a browser — all of these are indirect evidence, and each can be wrong in either direction. Without access logs, the whole technical door layer rests on inference.

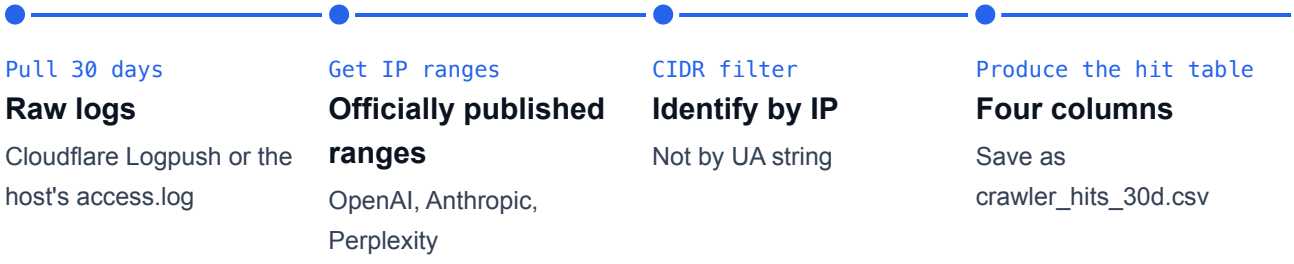


Figure: four steps to pull the logs; identify a crawler by its IP range only, because anyone can fake a UA string

OpenAI's three IP ranges are `openai.com/searchbot.json`, `openai.com/gptbot.json` and `openai.com/chatgpt-user.json`; Anthropic and Perplexity each publish their own official ranges. A hit count based on UA strings has scanners and spoofed traffic mixed into it, which invalidates the whole thing as evidence. For the template of the four-column hit table (crawler name, hits in 30 days, top 20 hit URLs, status-code distribution), see → General Edition B.1 Door templates.

Two fallbacks when you can't get it: no export access → open a task the same day and list a **red blocker**, and let gate 2 stand in with criteria b–d for now (see 2.3); the site never kept logs at all → turn on Logpush / access logs the same day, and **the first table only exists after 30 days**.

How to read the hit table

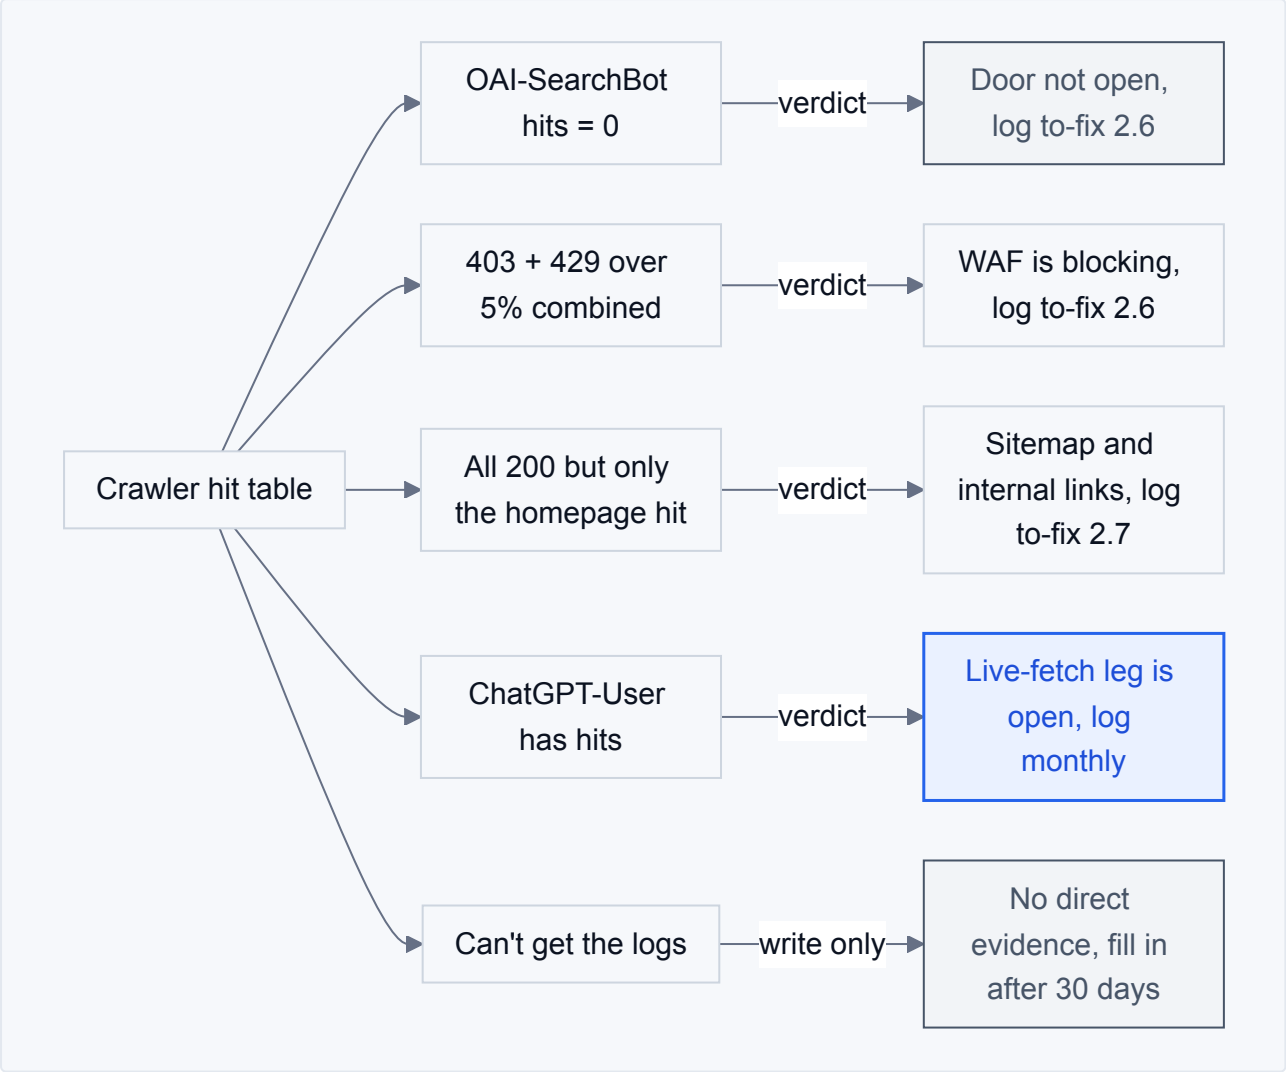


Figure: each reading on the hit table maps to exactly one to-fix item; before the baseline is saved, record only, change nothing

Three rules outside the figure:

- **At the layer where `OAI-SearchBot` hits are 0, do not write "wrong topic", and do not skip ahead to the "change topic" layer.** If the door isn't open, or the site simply hasn't been discovered, that's a door problem.
- **403 / 429 mean the WAF or a rate rule is blocking, not robots.** Editing robots will not fix it.
- **This table is three things at once**: gate 2's criterion a, the first check in the monthly troubleshooting pass, and the only direct evidence of whether a given leg is actually open. When you can't get the logs, don't say "the door is open"; note the catch-up date on the first line. Saying

you have no evidence when you have none is a book-wide rule of how you write things up (see → General Edition D.1 Number discipline: how to label numbers, and what stays internal).

For how to fix each reading, see the two diagnostic figures in 2.9.

2.2 Gate 1 read-only audit: what to check in nosnippet, robots and the WAF

What you'll do in this section: spend 10 minutes grepping the three nosnippet controls, then spend half an hour recording the current state of robots and the WAF in six lines. When you're done, you'll have a *Six-line door read* and 8 screenshots — record only, change nothing; every item you record gets fixed in 2.6.

Check nosnippet first, then robots

Google's own documentation states it plainly: to restrict how a page's content shows up in Search (**including AI Overviews and AI Mode**), you use `nosnippet`, `data-nosnippet`, `max-snippet` and `noindex` (developers.google.com/search/docs/crawling-indexing/robots-meta-tag). This is **the one switch on the Google leg that can zero you out with a single click**, and CMS templates and SEO plugins often turn it on by default. It takes 10 minutes to check — check it first.

Do this on 8 pages: the homepage, the price page, the `/facts` page, plus 5 template pages (a service page, a practitioner's or author's person page, an old blog post, a list page, a landing page). `curl` each page for **the HTML plus the response headers**.

Where to check	What a hit looks like
meta robots / googlebot in the HTML	nosnippet, or max-snippet:0 / :20
The HTTP response header X-Robots-Tag	Same as above; you need curl -I to see it
Tag attributes on key parts of the body	data-nosnippet; most often found on price tables, FAQ answers and fact blocks

Figure: the three nosnippet controls hide in three places; looking only at the HTML misses the response-header one

D1's receipt: the results of the three greps (all empty, or which one hit and on which page), one screenshot per page for all 8. **This switch gets turned back on after a redesign, a template change or a plugin update**, so it is also a standing item in the monthly door-layer check (see 2.9).

> **Note** Many people treat `Google-Extended` as the AI Overviews switch and spend effort arguing over whether to allow it, while missing this real switch. For why it isn't `Google-Extended`, see the four misconceptions in 2.4.

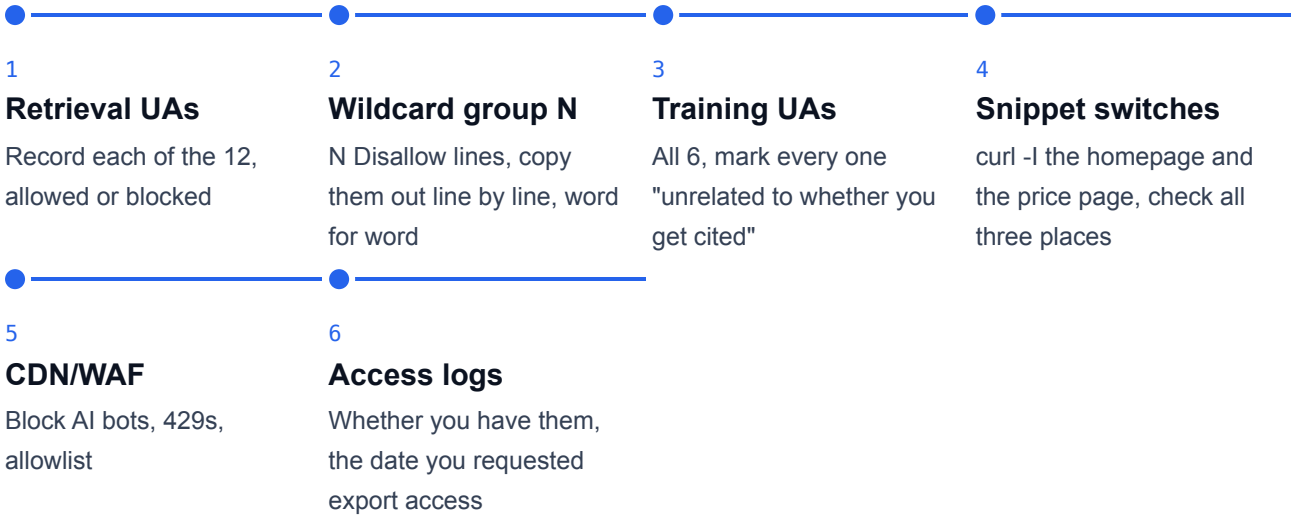


Figure: six lines for the read-only audit — miss even one and it doesn't count as checked; lines 2, 4 and 6 each sit on a failure path that can waste the whole quarter

The three failure paths are: skip line 2, and merging robots later strips away the protection that was already there (2.6); miss line 4, and `nosnippet` stays on with nobody noticing, so the Google leg counts as zero for the quarter; skip asking about line 6, and gate 0 can't start work the same day, leaving the whole door layer as nothing but inference — which is why this line should be asked as early as possible.

Line 5 needs three things nailed down: ① whether the "Block AI bots" / "Bot Fight Mode" switch is on, off, or doesn't exist; ② whether rate rules hit crawlers with 429s (check against the 429 share in the gate-0 table); ③ whether the IP range for `openai.com/chatgpt-user.json` is on the allowlist.

Where to check each line: for lines 1 and 3, open `https://<domain>/robots.txt` and read it group by group; for line 5, check the CDN / WAF dashboard; for line 6, ask ops and record "has Logpush" / "has `access.log`" / "kept no logs". The six lines take roughly 10, 5, 5, 10, 10 and 5 minutes in turn. For the list and purpose of the 12 retrieval UAs and 6 training UAs, see 2.4; for a blank record sheet, see → General Edition B.1 Door templates.

> **Note** **Allowed in robots does not mean actually let through** — the CDN / WAF layer can still return a straight 403 regardless of robots (**inferred, 85% confidence**). "Allowed" from the audit is only the paper-level first layer; the only evidence of actually being let through is gate 0's *Crawler hit table*.

2.3 Gate 2: four identities, fetched live (read-only)

What you'll do in this section: stop judging the door by looking at content through `curl -A`, and instead look at the same URL through four paths, then read the differences between the four to tell

CSR, WAF and a closed door apart. When you're done, you'll have four separate receipts — never merge them into one.

False negative: you conclude "blocked"	False positive: you conclude "open"
The WAF allows verified crawlers through by IP + rDNS	The site serves a degraded SSR version to non-browser UAs
You fake the UA from your office and get blocked as a spoofed crawler	The real retrieval path actually gets the CSR version
So you go fix something that wasn't broken	So you spend the whole quarter writing pages AI can't read

Figure: use curl -A to impersonate a crawler and read the content back, and you can get the verdict wrong in either direction

`curl -A "OAI-SearchBot"` only proves how your infrastructure reacts to that string. The kind of pass-through on the left isn't hypothetical: Cloudflare Verified Bots really does allow by IP + rDNS.

Four identities: one URL as four paths see it

Ranked by how trustworthy each one is; all of them are free.

Identity	Where you get it	Pass criterion	What it can prove
a · What the real crawler sees	Gate 0's access logs, CIDR-filtered against the official IP ranges such as `openai.com/searchbot.json`	This URL was hit by `OAI-SearchBot` within the last 30 days **and returned 200**	**The only direct evidence**, applies across every leg
b · The HTML Bing fetched	Bing Webmaster Tools → URL Inspection → the source returned by **Live Test**	**You can string-search the source and find the price figure and the body text**	The Bing / Copilot leg; also the free sample closest to "what a non-rendering crawler sees"
c · The HTML Google fetched	Google Search Console (GSC) → URL Inspection → the HTML under "**Crawled page**"	Same as above, string-searchable	**Only proves Google's side** (Google renders), never use it alone as your conclusion
d · What ChatGPT's live-fetch path sees	Paste this URL into ChatGPT, verbatim: "Open this link and copy out, word for word, the line on the page that gives the price."	It can **quote it back word for word** (not paraphrase it, not estimate it)	The `ChatGPT-User` live-fetch leg is open

`curl` is not a fifth identity — it's a troubleshooting tool, and it has to be split into two separate reads: ① fetch the HTML with an **ordinary browser UA** and check whether the body text and price are there —

this is the right way to judge CSR, and **CSR has nothing to do with which UA you use**; ② run it again with the `OAI-SearchBot` UA, **look only at the status code, not the content**, to judge the WAF. Screenshot each of the two separately, and write on the receipt "**for troubleshooting only, not a pass criterion**". The half-day quick-start version (→ General Edition 4.8 Two shortcuts and this chapter's checklist) also judges JS following step ①.

How to read the differences between the four

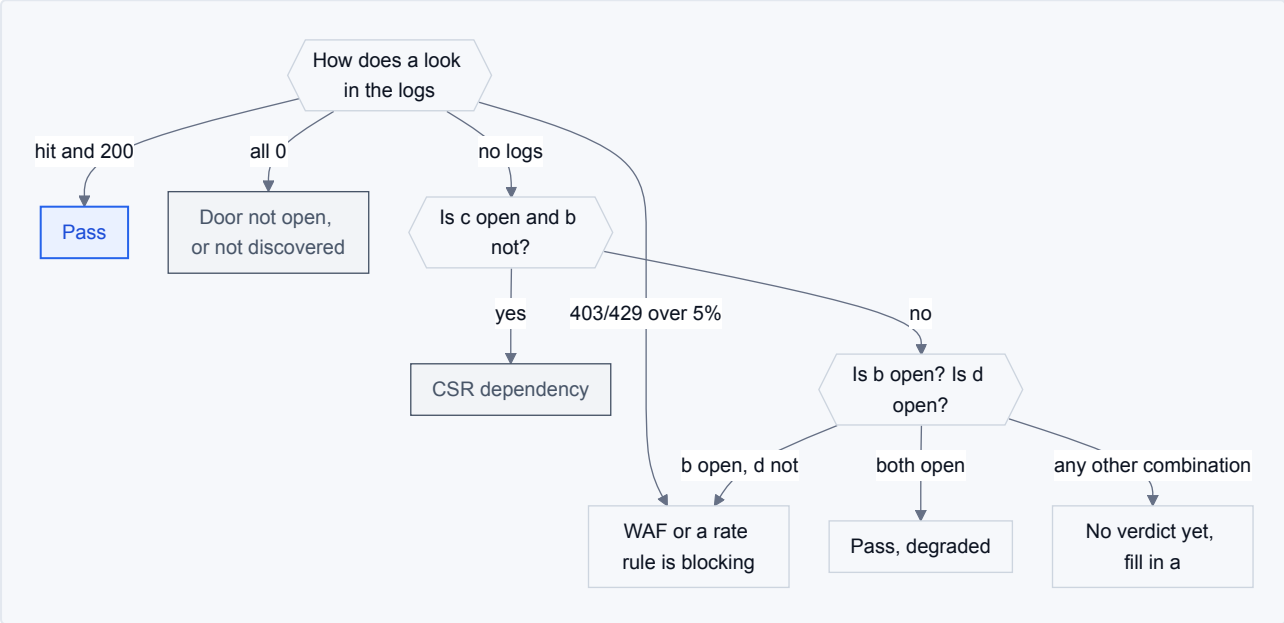


Figure: the four-identity decision tree — the conclusion comes from the differences between the four, not from ticking each one off

Outside the figure, also note:

- **"Pass, degraded" must have "no direct evidence, fill in a after 30 days" noted on the first line.**
- **CSR dependency** = Google's rendering can get it, but a non-rendering crawler cannot. Log it as a to-fix item and open an SSR / prerendering task the same day; push all page work back as a whole, and add no new content pages this month.
- **Never edit content for the WAF cell** — that belongs to the WAF allowlist in 2.6.
- **For the door-not-open cell, never write "demand doesn't exist"**, and never skip ahead to "change topic".

D1 is the audit. On day 7 after fixing the door (2.6), pull all four receipts again and lay them side by side against the D1 set.

2.4 What each crawler is for, and judging the door leg by leg (read-only)

What you'll do in this section: tell the 12 retrieval crawlers and the 6 training crawlers apart, and know which leg breaks when you block which one; when you hit a platform, multiple subdomains, a login wall

or a multilingual site, judge the door leg by leg. When you're done, you'll have a *Per-leg door table*: one row per host, one column per leg — never merge into one cell.

This is the single easiest thing in the whole chapter to get wrong: block the wrong one and nothing happens; allow the wrong one and the whole quarter is wasted.

Retrieval legs (12)	Training legs (6)
Decide whether you can appear in AI answers at all	Only about content licensing, unrelated to whether you get cited
Block OAI-SearchBot and the ChatGPT leg is zero for the quarter	Block GPTBot and you can still be cited
Block Googlebot and you lose AI Overviews and organic rankings together	Google-Extended has nothing to do with AI Overviews
Allow or not: allow	Allow or not: your own call

Figure: block a retrieval leg and that leg breaks; Disallow every training leg and not one retrieval leg is affected

12 retrieval legs: what blocking each one costs

Crawler	Belongs to	What blocking it costs
`OAI-SearchBot`	OpenAI	**The ChatGPT leg is zero for the quarter**
`ChatGPT-User`	OpenAI	Users who open your page from inside ChatGPT can't read it. Note **The real switch is in the WAF, not robots** (2.6)
`OAI-AdsBot`	OpenAI	Only affects ad scenarios; if you don't run ads, you don't have to allow it
`Bingbot`	Microsoft	Copilot breaks outright; one of ChatGPT's eight retrieval sources goes missing
`Googlebot`	Google	**You lose AI Overviews / AI Mode and organic rankings together**: all three share the same crawler
`Applebot`	Apple	**The whole Siri / Spotlight / Apple Maps / Apple Intelligence leg breaks**. Note Until this group is allowed, don't claim your Apple Business Connect profile
`Amazonbot`	Amazon	Alexa / Amazon-side retrieval breaks

Crawler	Belongs to	What blocking it costs
`PerplexityBot`	Perplexity	The Perplexity leg breaks
`ClaudeBot`	Anthropic	Claude-side access breaks
`Claude-SearchBot`	Anthropic	Claude retrieval breaks
`Claude-User`	Anthropic	Users who open your page from inside Claude can't read it
`DuckAssistBot`	DuckDuckGo	DuckAssist breaks

The 6 training legs: `GPTBot` · `Google-Extended` · `Applebot-Extended` · `Meta-ExternalAgent` · `CCBot` · `Bytespider`. Whether you allow them is a content-licensing question, not a visibility question.

Misconception	Fact
Blocking GPTBot protects your content, or gets you shut out of AI	GPTBot is a training leg; 88.2% of sites that block it still get cited
Allowing Google-Extended turns on AI Overviews	AI Overviews and AI Mode go by Googlebot + nosnippet
Block Applebot-Extended and the Apple leg breaks	It's only a training opt-out switch; the Apple leg goes by Applebot
No AI crawler respects robots, so editing it is pointless	All three Anthropic bots respect it; it's ChatGPT-User that doesn't

Figure: the four most common misconceptions — all of them mix up training legs with retrieval legs, or mix one company's rules with another's

The real danger is the flip side of the first one: blocking `OAI-SearchBot` and `Bingbot` themselves. The second is the most expensive to get wrong, and it rests on two of Google's own sentences. The crawler documentation: "**Google-Extended does not impact a site's inclusion in Google Search nor is it used as a ranking signal in Google Search.**" *AI features and your website* nails down the entry requirement: "**a page must be indexed and eligible to be shown in Google Search with a snippet**", and "**robots.txt directives for Googlebot is the control**". So the gate for AI Overviews is `Googlebot` plus the three nosnippet controls. The point of the fourth is that it saves you work: for the three Anthropic bots (`ClaudeBot` / `Claude-SearchBot` / `Claude-User`), editing robots genuinely takes

effect — you don't need to touch the WAF. 88.2% comes from BuzzStream; for the rest of the sources, see → General Edition A.5 List of sources.



Figure: judging a leg means checking the right UA; check the wrong one and you get one of two opposite wrong verdicts — "allowed" or "can't get in"

The zero-JS-execution corollary: mainstream AI crawlers do not execute JS; the only exceptions are Gemini (which runs on Google's infrastructure and renders in full) and Applebot (evidence in → General Edition 1.1 How buyers ask, and what AI reads). So the price, the core facts and the body copy must be in HTML that's already been rendered server-side — fetch the source with a browser UA and you should

be able to string-search it; write them in JS instead and the next 90 days are wasted. The flip side of the same corollary: live fetching doesn't read JSON-LD, so schema is a baseline, not a lever (2.7). This exception also gives the whole book its single most useful troubleshooting criterion: **the Google leg has seats while the ChatGPT leg is zero — that difference by itself is the fingerprint of CSR.**

Judging the door leg by leg

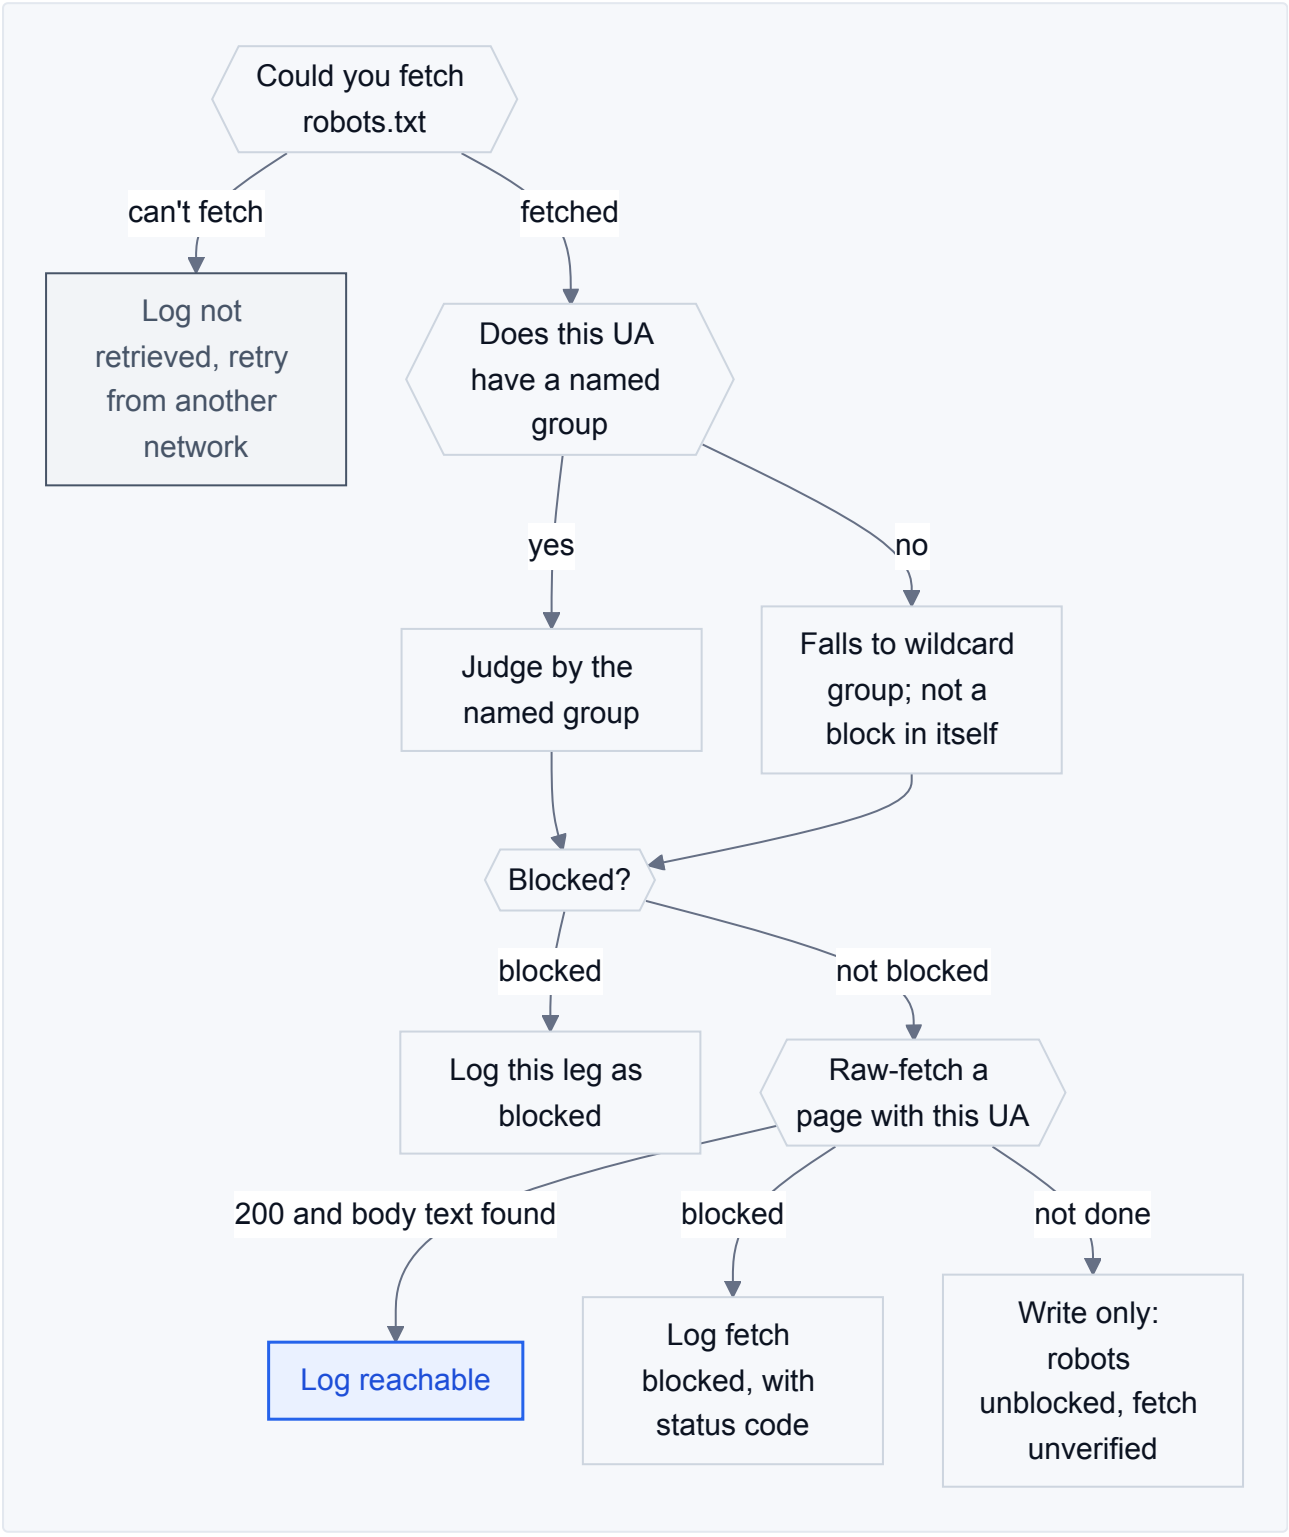


Figure: judging one leg's door — not blocked in robots does not mean it can actually fetch; only a raw fetch returning 200 with the body text found counts as "reachable"

Recording rules outside the figure — skip none of them:

1. `curl` **the whole of robots.txt yourself, download it, parse it group by group, and log it by date.** Grepping for one keyword alone will miss the single most important thing: that this UA was never given its own named group at all.
2. **Check robots on the host that the page actually lives on, not the main domain.** Detail pages, doc sites and help centres are often on a different host or subdomain, and robots is **not shared** with the main domain.
3. **One row per host, one column per leg — never merge into one cell.** The very same platform can easily have some legs blocked and others open; merge it into one cell and the whole chain gets thrown away.
4. `Google-Extended` , `Applebot-Extended` **and** `Meta-ExternalAgent` **are training / licensing legs — they don't go on this table.**
5. **The table must carry a "verification level" column**: `robots-verified` / `fetch-verified` . A raw-fetch receipt may only say "fetched 200 + body text found" or "fetch blocked (`<status>`)", never "allowed". robots only lifts a protocol-level ban; CSR rendering and WAF / geo-blocking (403s, captchas, redirects) are two other, separate gates, and they are not written in robots.
6. **A platform's robots can change**, so log it in the *Off-site asset decay table* and **rerun it once a quarter**. A refused connection, a timeout or a failed TLS handshake all get logged as "not retrieved"; if a retry still fails, leave it blank — **never guess and fill it in.**
7. **A platform closing its door ≠ that platform has no sales value.** All that's ruled dead is "citations through that one leg", not the business; other legs, and the shopping shelf that runs through data feeds, may still show up as normal. Never mix these two things up in anything you tell clients.

On a host where a given leg is blocked, schedule no content work on that leg, and don't let acceptance count it towards seats. What's being judged here is only whether the door is open, which is a different thing from → General Edition 4.6 Four states, two denominators and the abstain line judging whether "a given URL can be reached at all". For reproducible, per-platform robots commands, see → General Edition B.1 Door templates.

> **Example** (training) Measured on 2026-09-21, the robots.txt of TPGateway, Singapore's official training portal: `User-agent: *` → `Disallow: /` ; the `Allow: /` groups that follow list only `Googlebot` (including Image / Video / mobile), `Slurp` , `bingbot` , `Applebot` , `Twitterbot` and `Pingdom` . Not one of `GPTBot` / `OAI-SearchBot` / `ChatGPT-User` / `ClaudeBot` / `PerplexityBot` appears — all of them fall under `Disallow: /` . → The Google, Apple and Bing legs are open; the OpenAI, Anthropic and Perplexity legs are blocked: build for the first three legs only, accept work against the first three legs only, and don't count ChatGPT seats.

Five kinds of site, one extra check each

Kind of site	Extra check	Criterion and how to handle it
Single-page app / client-side rendering	Fetch the HTML with an ordinary browser UA, grep for one actual sentence of body text	Not found = the body text isn't in the HTML. Move key body text to server-side rendering or prerendering — no "wait and bundle it with the next SEO redesign"
Multiple subdomains	`curl` a robots.txt for each subdomain separately, one cell per subdomain × retrieval leg	For any subdomain with no receipt, its pages don't count as "door open" this quarter and can't be counted towards deliverables. Doc sites, help centres and trust pages need their door opened and to be server-side rendered too
Login wall / quote wall	Whether the key figures sit outside the wall	**Anything behind the wall is invisible to AI**: build a public, checkable summary page or a mirror on the main domain
Multilingual sites	Run the four-identity check once for each language	Never check only one language
Price or stock that changes very fast	Whether these fields are injected by a client-side script	Yes = as good as not there

Judge the two causes separately, and don't reverse the order: check the robots receipts first (one `curl` per subdomain × leg — the cheapest check). When one subdomain was missed, the blocked leg reads zero and the unblocked legs read normal — a one-line fix; only check CSR once robots is all clear. Reverse the order and you'll schedule a whole quarter of engineering rework for a problem that a single line in robots would have fixed. Only once both are fixed does that leg go from 0 to 1.

> **Example** (B2B software) A site naturally has multiple subdomains: `www` / `docs` / `help` / `blog` / `status` / `app` / `trust` . Check robots for each subdomain separately — a clean main domain doesn't mean `docs` is clean too.

2.5 Gate 3, read-only: connect the two free AI reports

What you'll do in this section: on D1, connect GSC's generative AI report and Bing Webmaster Tools' AI Performance, actually read each one once, and check that the site field they return really is your own site. When you're done, you'll have two sets of data that are the only free, first-party, unfakeable data in the whole engagement; only open the probe budget after you've read them.

1
Request access
Domain verification, in the same email as the log access request

2
Open the GSC report
Generative AI performance, export the CSV and screenshot it

3
Open the Bing report
Webmaster Tools AI Performance, same as above

4
Actually read it once
Check the site field — get it wrong and you'll silently read another site

5
Save it
Two connection screenshots; how to read them is in 4.2

Figure: five steps to connect the two dashboards; reading the wrong site is worse than not connecting at all, because the tool will silently read another site and still return ok

Three rules:

- **Get the site field wrong = this box isn't done — don't mark it complete.** When the site parameter isn't configured correctly, the tool will read a different site and still return ok, which is worse than not connecting at all: you'll be treating someone else's numbers as your own baseline. The same goes for other read-only dashboards such as GA4.
- **Can't get domain-verification access → list a red blocker.** It blocks more than just this box: Bing's indexing submission (2.7), and Bing Live Test and GSC's crawled-page view (identities b and c in 2.3) all need domain verification first.
- **This section is connect-and-read only.** How to read each of the two reports, what each one gives you, and what to watch out for are written up once, in → General Edition 4.2 The question pool: where the 30 questions come from, how they are balanced, how they are signed; checking Bing for indexing gaps and submitting via IndexNow are changes, and belong in 2.7.

Another free thing to start the same day, in passing, is copying out the actual words buyers used, from your front-desk records or CRM, for the last 15 deals before they closed — it's the question pool's only class A source, and how to do it is also in 4.2.

2.6 Fixing the door: nosnippet, the four-step robots merge, the WAF allowlist

What you'll do in this section: first confirm the baseline is saved, then on the same day make all the fixes in the order `nosnippet` → robots → WAF, and log the day you fix the door as the split day. The hard gate still applies: **until the baseline is saved, do not change the door, any profile or the facts page.** When you're done, you'll have a merged robots.txt, a receipt with the N × 12 line written on it, and one WAF allowlist rule.

Precondition
Baseline saved
Baseline Top 20, noise band, the 36 brand-six runs

Gate 1a
nosnippet
Revert all three places, grep comes up empty

Gate 1b
robots
Four-step merge, receipt has the N × 12 sentence, word for word

Gate 1c
WAF
Add the chatgpt-user.json IP range

Log split day
The day you fix the door
Written into the page register

D+7
Re-pull and recheck
Check the logs for ChatGPT-User, re-pull the four identities

Figure: on the day you fix the door, go in this order — the split day counts from this day; if the precondition isn't met, take none of the steps after it

The day you fix the door is the split day for the whole engagement; every "before vs after fixing the door" comparison is drawn against it. Each page also has its own split day: the day it goes live (2.8). The book's W1 also counts from this day.

Gate 1a · nosnippet

Where it hits	How to fix it
nosnippet or a low number in meta robots / googlebot	Change it to max-snippet:-1
The response header X-Robots-Tag	Remove it at the CDN / server layer, or change it to max-snippet:-1
data-nosnippet on a body tag	Delete this attribute

Figure: each of the three hit locations has its own fix; only when a rerun of the grep comes up empty on all three does it count as passed

Any hit left unfixed for four weeks → the Google leg counts as zero for the quarter — write it on the first line, and mark the Google cell "switch not lifted" in the monthly record.

Gate 1b · The four-step robots merge

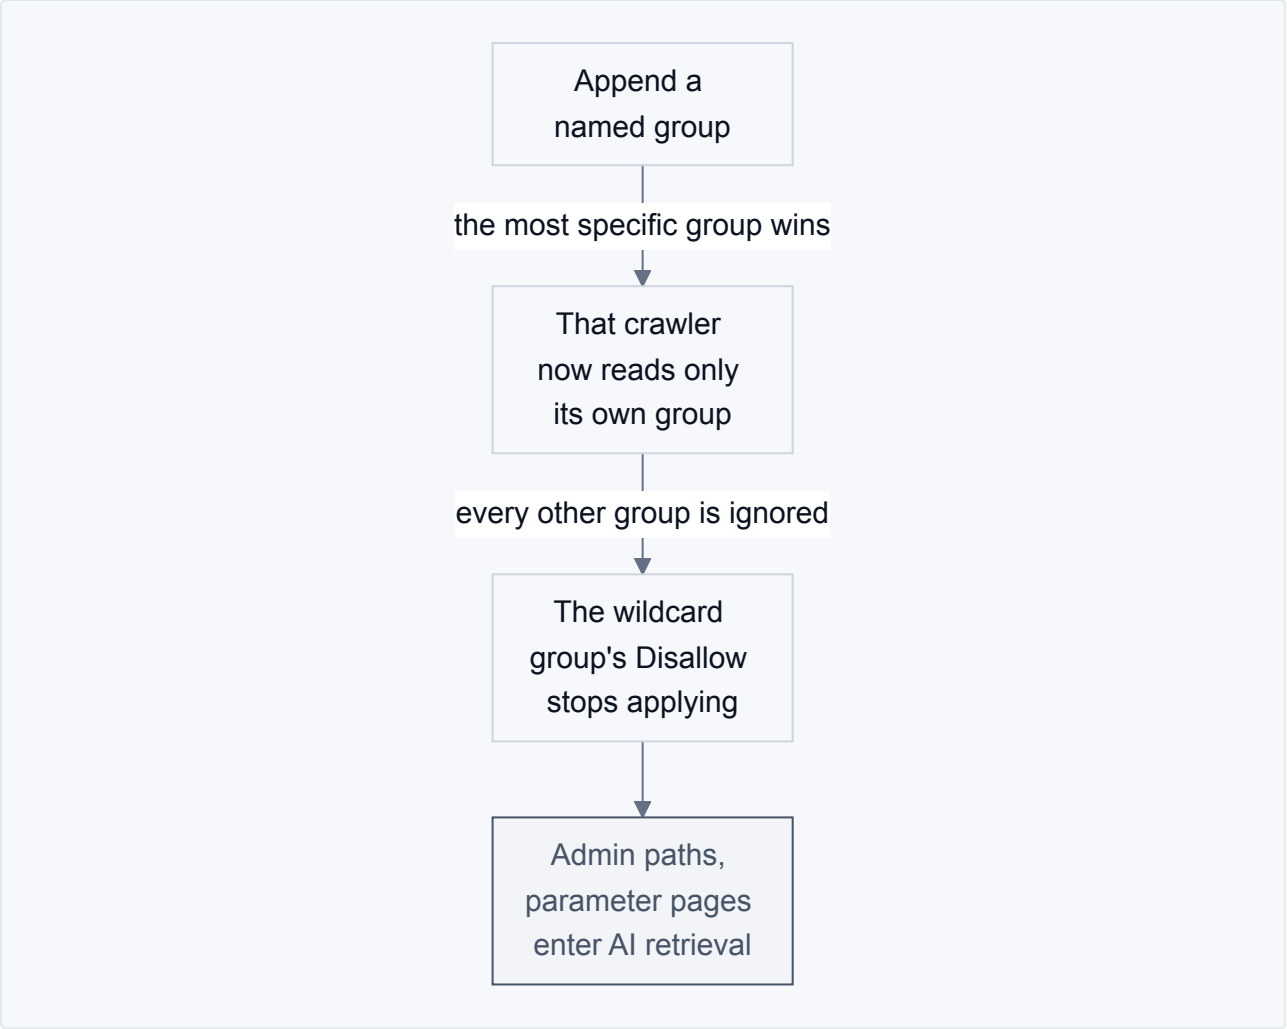


Figure: simply appending a named group to the end of robots pulls these crawlers out of the wildcard group, and every protection it gave them stops working

What gets let out is usually paths like `/wp-admin/`, `/cart/` and `/?s=`. The consequence isn't a stalled ranking — it's admin paths and parameter pages ending up in AI retrieval.

1
Copy the wildcard group
Copy the Disallow lines word for word, record the line count N

2
Write the retrieval groups
12 named groups, each with Allow: /

3
Copy the N lines
Into every named group, until the count reaches N × 12

4
Finish up
List training groups separately, put the Sitemap line at the end

Figure: four steps to merge robots, never reverse the order; skip step 3 and you've lifted the protection that was already there

The `Sitemap:` line goes at the end because it doesn't belong to any group. Whether to allow the training crawlers is your own call, but any training group that says `Allow: /` needs those same N lines too. **The receipt must record this sentence word for word:**

> Original wildcard `Disallow` lines = **N**, copied line by line into **N × 12** places (12 named retrieval groups); training groups with `Allow: /` also copied.

If this sentence isn't on the receipt, gate 1 doesn't count as passed. The only full text of the robots master template — 12 retrieval groups plus 6 training groups — lives in → General Edition B.1 Door templates; hand it to whoever edits the site and have them copy straight from there. The `<N lines of Disallow>` in the master template is a placeholder and must be replaced with the original text copied down in step 1 — not one line short, and in every group.

Won't edit robots → all site-layer work is downgraded to read-only monitoring.

> `Perplexity-User` does not go in the master template. Perplexity's own crawler documentation says: "Since a user requested the fetch, this fetcher generally ignores robots.txt rules." (https://docs.perplexity.ai/guides/bots). Like `ChatGPT-User`, robots can't control it; the only way to allow it is to add its official IP range in the WAF (published on the same page as https://www.perplexity.com/perplexity-user.json).

Gate 1c · WAF: the real switch for ChatGPT-User

`Write Allow: / in robots`	`Add the official IP range to the WAF allowlist`
Harmless, but has no effect on ChatGPT-User	The real switch, verifiable in the logs after 7 days
OpenAI's own words: robots.txt rules may not apply	Name the rule allow-chatgpt-user
Fix only this side and you'll think the door is open	The live-fetch path is actually still closed, ChatGPT-User hits stay at 0

Figure: for ChatGPT-User, Allow in robots is only a declaration — the real switch is in the WAF

OpenAI's own complete sentence is: this kind of fetching is triggered by a user action, "Because these actions are initiated by a user, robots.txt rules may not apply". On the WAF side, do three things the same day, and screenshot each one:

1. A screenshot of the "Block AI bots" / "Bot Fight Mode" switch's state.
2. Rate rules: move verified crawlers out of the rate rules, screenshot it, and attach the 429 share from the gate-0 table.
3. Allowlist: add the IP range for `openai.com/chatgpt-user.json`, name the rule `allow-chatgpt-user`, screenshot it.

D+7 recheck: re-pull the logs and check whether `ChatGPT-User` has any hits; on the same day, re-pull the four receipts from 2.3 as well and keep them side by side with the D1 set.

2.7 Gates 3 and 4: indexing paths, and JSON-LD sealed once

What you'll do in this section: handle indexing as two separate legs — use IndexNow for the Bing leg, and profile claiming plus pasting the URL for a quote-back on the ChatGPT leg; then spend half a day sealing the site's JSON-LD once, after which it never goes into a per-page check again. When you're done, you'll have two indexing receipts and one sealed schema template.

Gate 3 · Indexing in two separate cells

Bing leg	ChatGPT leg
Check indexing page by page with site:, submit any gaps via IndexNow	Foursquare claiming: pick categories word for word from the official category tree
Check IndexNow Insights in Bing Webmaster Tools	Get the Foursquare business status right; claim Yelp too
A submission-to-indexing report exists and can serve as a receipt	Paste the URL and have ChatGPT quote the price line back word for word
Copilot uses this leg directly	Local questions mainly go through Foursquare, Yelp and OpenAI's local index
IndexNow only feeds Bing and Yandex	The criterion is identity d from the four identities; Bing indexing is not a criterion here

Figure: indexing is two separate cells, never merge them — being indexed on Bing doesn't prove the ChatGPT leg has read you

"ChatGPT runs through Bing's index" no longer holds as of 2026-09: ChatGPT has multiple retrieval sources, and Bing is only one of them (evidence in → General Edition A.1 Evidence for mechanics, the door and identity). So still do IndexNow, but it isn't the switch for the ChatGPT leg. Claiming a profile is itself a change, and is only done once the baseline is saved; for the full method for the five business profiles, see → General Edition 3.4 Third-party credential tiers, Wikidata and the five business profiles. Having `Applebot` allowed, and verified, is a precondition for claiming your Apple Business Connect profile.

Gate 4 · JSON-LD sealed in half a day

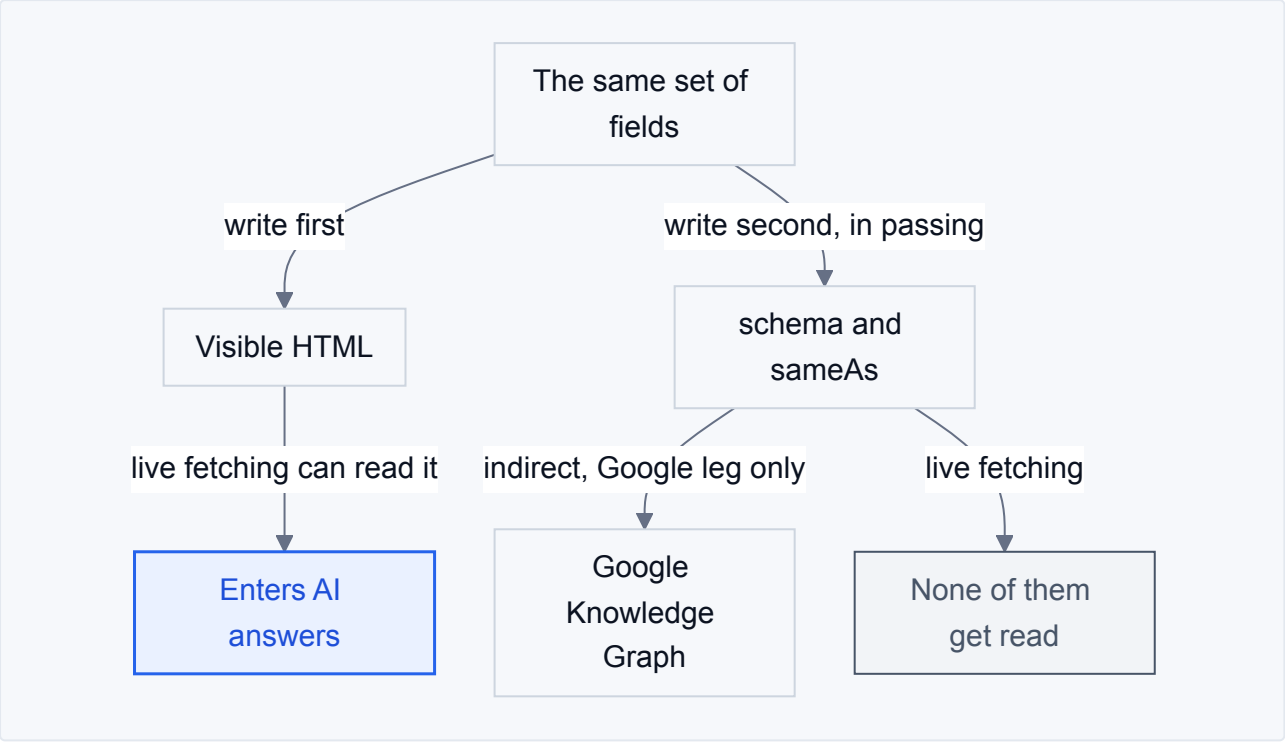


Figure: text only gets read by live fetching once it's in visible HTML; schema only takes the indirect route through the Google Knowledge Graph

Reverse the order and you end up with a site whose fields are all complete but that AI still can't read. Even schema's small indirect effect only exists on the condition that the same fields are already written into visible HTML; do it on its own and the return is close to zero. After adding JSON-LD, AI Overviews (AIO) citations moved −4.6%, the only significant result, and it's negative; the other two platforms weren't significant (evidence in → General Edition A.1 Evidence for mechanics, the door and identity), so this is a one-off task:

Rule	What it means
One-off template	Organization (or your industry's schema subtype) + a Person for every practitioner + sameAs. Set it up once when the site is built, **sealed in half a day**
Not written per page, not checked per page	Doesn't go on the per-page launch checklist, and doesn't go on the review checklist
Only four things need to be right	① the legal name matches `/facts` word for word ② the address and phone number match the five business profiles word for word ③ each practitioner's registration number ④ the sameAs string
A confidence note	Write one line in the template file — "**indirect, Google leg only, low confidence**" — so the next person doesn't mistake it for a main lever

The sameAs string links these: regulator register or registration-number lookup pages, association profile pages, Google and Bing business pages, ORCID / Scholar, conference speaker pages, LinkedIn, the Wikidata QID.

Two rules in passing: × don't stack `FAQPage` schema (it's built for rich results and doesn't feed AI citations); × don't add a self-reported `aggregateRating` (on the strictly regulated side this also runs into the ban on testimonials and ratings: Statute text in the example industry). **Do build question-style H2s and a short Q&A at the foot of the page; don't stack schema.**

The one exception: when structured data is used as a **data channel** (for example, a product feed feeding a shopping graph), it is useful — but that's a different route, one that doesn't go through live fetching. "Structured data is useful" is only true in this sense, and it is no reason to change anything on your landing pages. For the approved wording to give a client on this point, see → General Edition B.2 Approved wording and scripts.

2.8 The launch gate for every page

What you'll do in this section: on the day each page goes live, run it through three checks, then do a quote-back test on day 7; only once every check passes do you mark it "live", and log that page's split day. When you're done, you'll have a timeline for every page from launch to D90, and a page register logging the launch date.

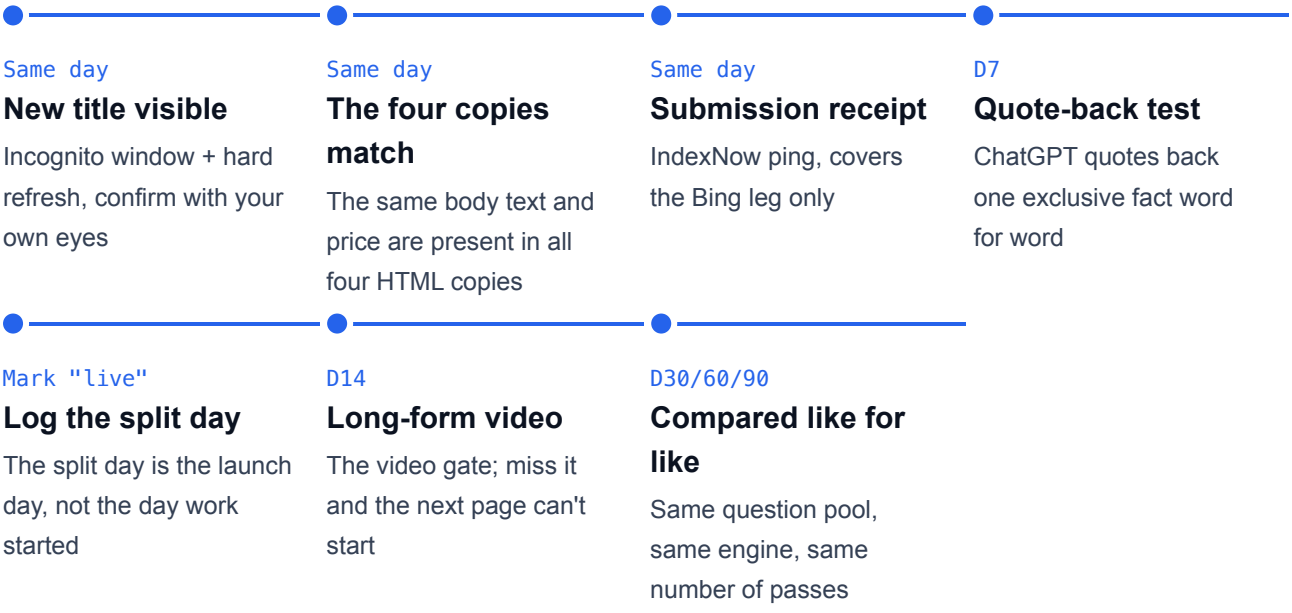


Figure: a page from launch to D90; miss any one of the first four items and you may not mark it "live" — this page's split day is logged as the launch day, not the day work started

What happens if each step is missed:

- Miss **new title visible**: whoever opens the page sees the old version.

- **The four HTML copies match**: fetch the source once with each of four UAs — `OAI-SearchBot`, `Bingbot`, `Googlebot`, an ordinary browser — and the same body text and the same price string must be present in all four. Missing this has two possible outcomes: showing machines a different version gets judged as manipulation, or the price never makes it into the answer at all. What this step checks is "are the four copies the same version", which is a different thing from the four identities used to judge the door in 2.3.
- The **D7 quote-back test** is the real acceptance check for the ChatGPT leg: ask ChatGPT once, using the exact wording, about one exclusive checkable fact on that page, or paste the URL and have it quote that line of figures back word for word; screenshot it for the record. n = 1 is valid here too — it tests "was it read or not", not "what position it ranks".
- **D14 long-form video** doesn't hold up "live" — it holds up starting the next page. It belongs to writing pages' three mechanical gates; specs in → General Edition 5.4 The 30-day writing order and the three mechanical gates and → General Edition 6.7 Long videos narrated by the named expert.
- **D30 / D60 / D90** must be measured to the same spec at all three points: same question pool, same engine, **the same number of passes**. Seat counts scale linearly with passes — baseline at 5 rounds, D90 run at 10 rounds, and the count doubles even if you did nothing at all; for why, see → General Edition 7.5 Ruler discipline: change the ruler and nothing is comparable.

When each page goes live, also log 2–3 checkable fact strings that are unique to it, and look for them again during each monthly retest — this is the only signal in the whole process that can give you page-level cause and effect; for how, see → General Edition 7.2 The noise band, the page-level signal and the ten monthly steps.

2.9 Door-layer troubleshooting and this chapter's acceptance checks

What you'll do in this section: whenever a monthly retest shows a seat increase that doesn't exceed the noise band, start checking from the door layer, and use the two diagnostic figures to map every symptom to one fix; check it even when seats went up. Only once every one of the ten items on the checklist is ticked does the door count as done.

When to check: two points in time. ① Whenever the seat increase doesn't exceed the noise band, start checking here — **never skip a layer** (for the rest of the triage layers, see → General Edition 7.3 Not-moved triage and next month's three points). ② **Check the door layer once a month, unconditionally, even when seats went up**: its failures are silent — while a page is not being read at all, seats can perfectly well be rising for some other reason, and you won't discover a whole wasted month until the next one.

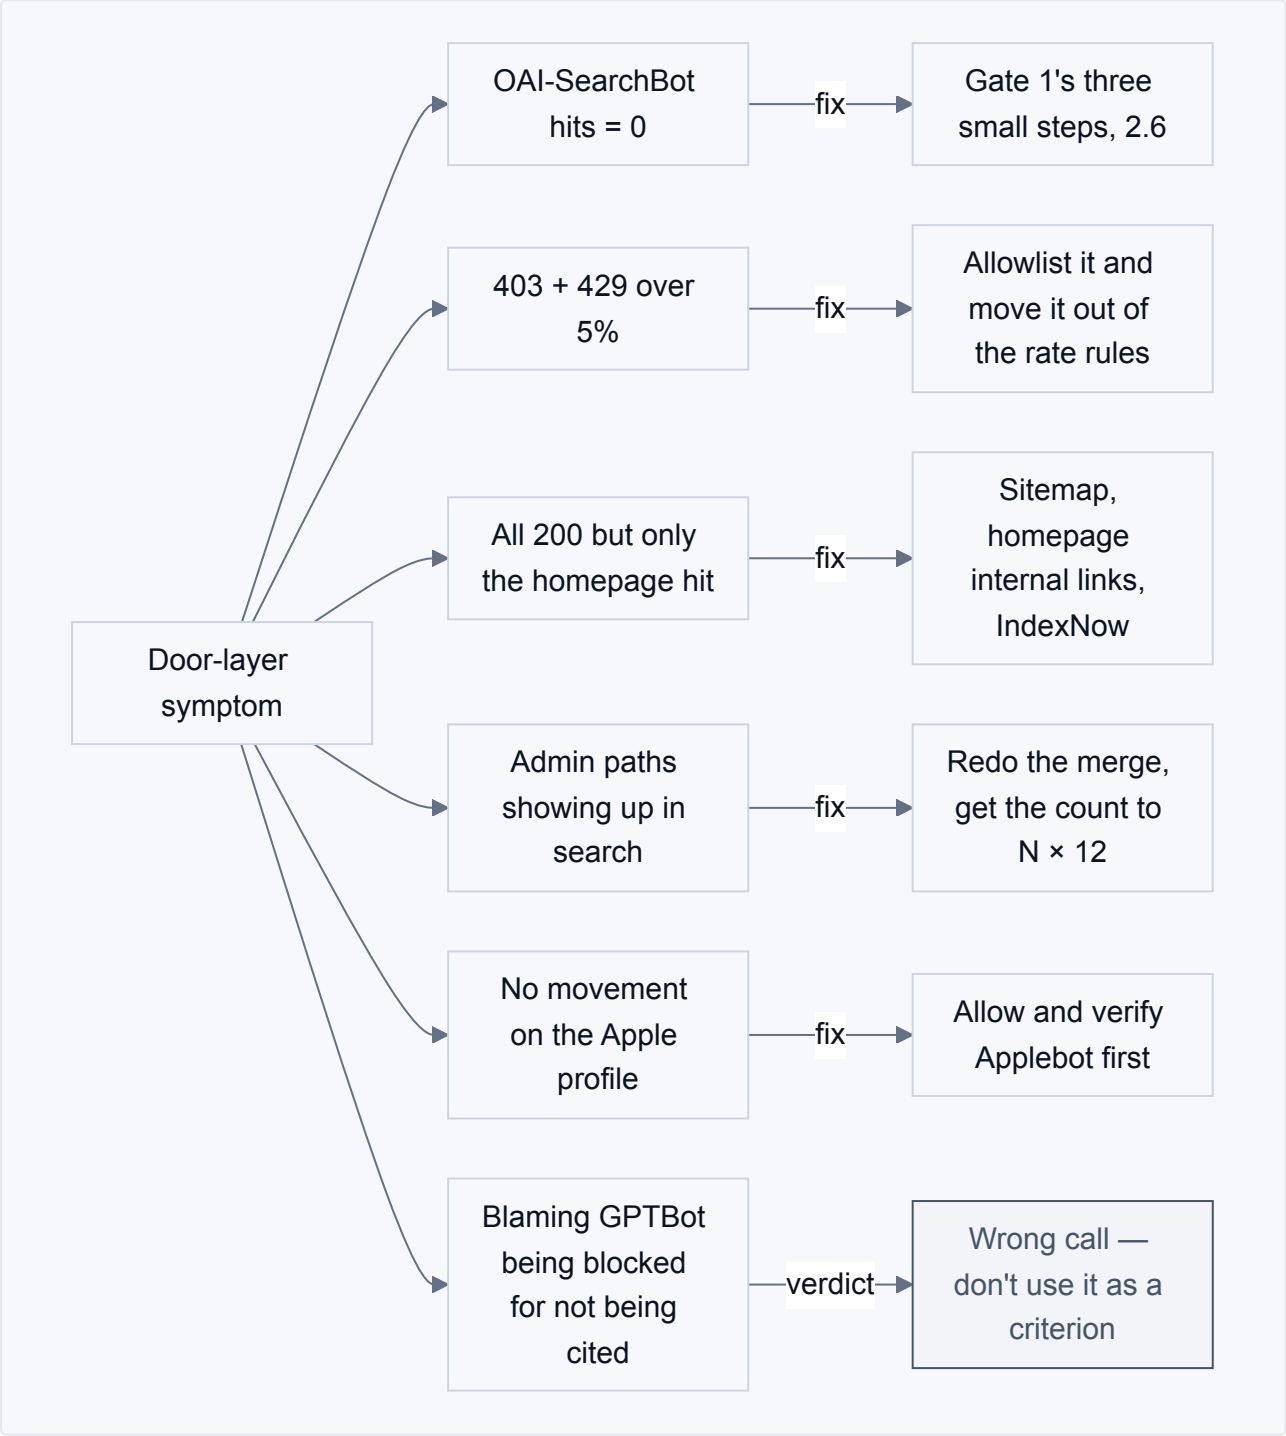


Figure: at the logs / robots / WAF layer, every symptom maps to exactly one fix

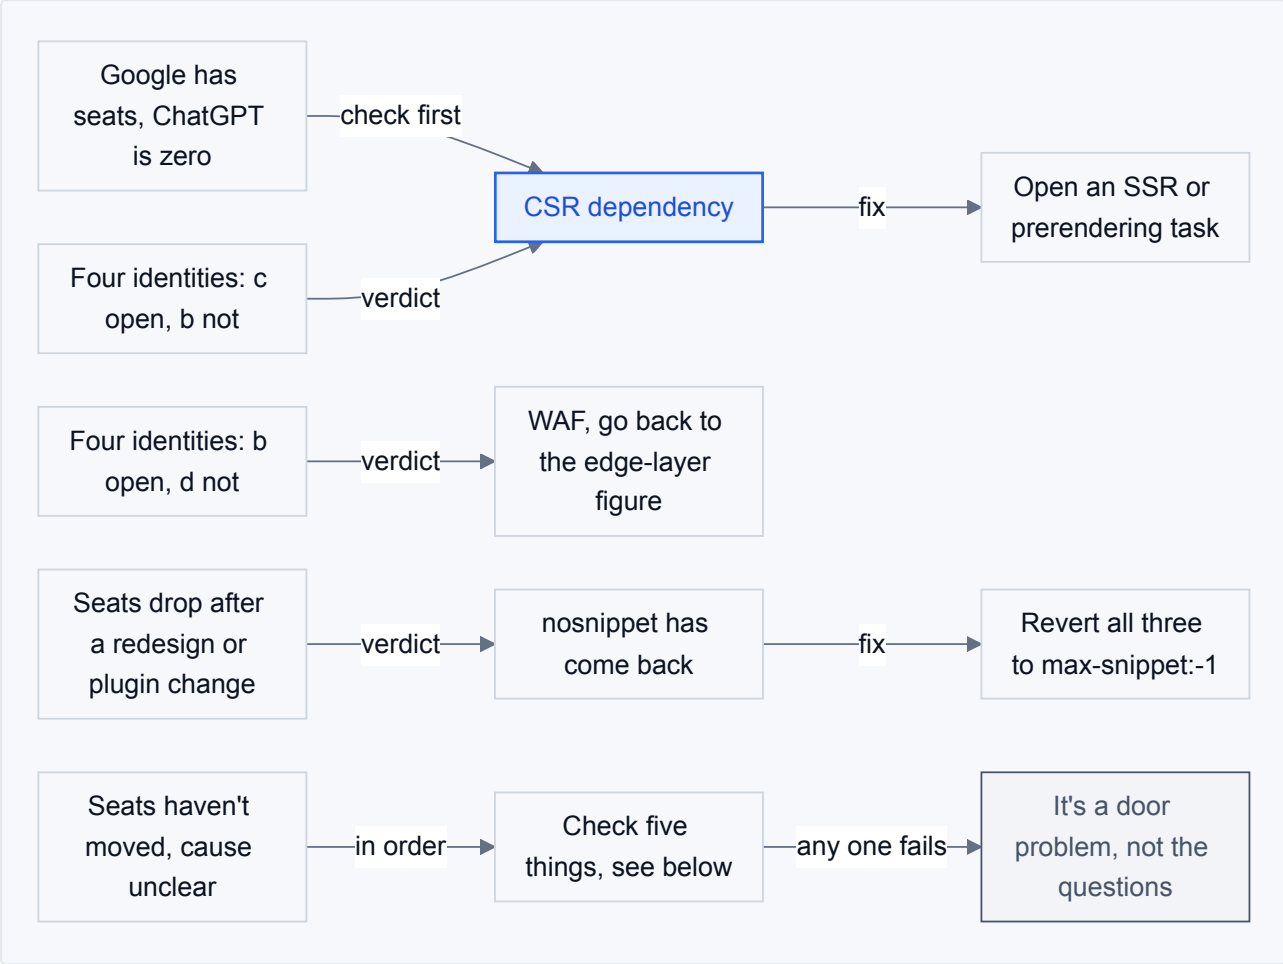


Figure: at the rendering / indexing / nosnippet layer — Google has seats while ChatGPT is zero: check CSR first, not the choice of questions

When the "cause is unclear", check these five things in order: ① whether the access logs show a real `OAI-SearchBot` hit, and what status code it returned ② whether the three nosnippet controls have been switched back on by a CMS or plugin ③ whether the four HTML copies match word for word ④ Bing indexing checked with `site:` ⑤ paste the URL into ChatGPT and have it quote the price line back word for word.

Rules outside the two figures:

- **Anything that fails: fix it the same day, recheck after 7 days, and don't change topics or add new content pages this month.**
- **Never edit content for the 403 / 429 cell.**
- **Never skip past the door layer and go straight to writing "demand doesn't exist". Stuck on the same layer, unfixed, two months running is an execution problem, not a problem with the questions.**
- **All five gates are mandatory in every industry.** Two more get added depending on the kind of site: a **CSR check** (for content rendered by front-end components such as price lists, course schedules, doc sites, filters and pricing widgets — fetch the HTML with a browser UA and check whether the body text is there); a **visible-text check** (raw-fetch with a retrieval-leg UA, strip the

tags, and count the visible body characters and the number of stand-alone statement sentences — a page with numbers but no sentences can't get onto the answer shelf at all).

> **Example** (e-commerce) The visible-text check is mandatory: a page with fewer than 12 statement sentences may not be marked "live". For a measured case where a category page's HTML was large but, once the tags were stripped, only a small scrap of visible body text remained, see A.1.

Acceptance checklist for this chapter

The door is a multiplier, not a variable: it only has two values, 0 and 1, and with no evidence it counts as 0. Tick these off, and only once every box is ticked does the door count as open:

- [ ] All four columns of the *Crawler hit table* filled in, saved as `crawler_hits_30d.csv`; where logs couldn't be obtained, a red blocker is listed with the catch-up date noted
- [ ] All three `nosnippet` greps come up empty, one screenshot for each of the 8 pages
- [ ] `robots.txt` has all 12 retrieval-class named groups, and the receipt has this written on it: "**Original wildcard `Disallow` lines = N, copied line by line into N × 12 places**"
- [ ] The `Sitemap:` line is at the end of the file
- [ ] The WAF allowlist rule `allow-chatgpt-user` has been added, with a screenshot on file; the 7-day recheck shows `ChatGPT-User` hits
- [ ] Each of gate 2's four criteria has **its own separate receipt**, none merged into one; the two `curl` reads are each screenshotted and marked "for troubleshooting only, not a pass criterion"
- [ ] Both dashboards — Bing Webmaster Tools AI Performance and GSC's generative AI report — are connected, each with a screenshot
- [ ] The JSON-LD template is sealed in half a day, and is never written per page after that
- [ ] Every page that has gone live has passed the launch gate (three checks the same day + the quote-back test on day 7), with the launch date logged
- [ ] The door-layer troubleshooting checklist has been added to the standing monthly actions

Every part of the checklist that's about "recording the current state" is done on D1; every item that actually changes the door is only done once the baseline is saved, and the day you fix the door is already logged as the split day.

Chapter 3

Get AI to recognise your business: the same facts everywhere

Name, address, phone, licences and other facts, word for word the same on your site and every profile

3.1 Why this comes before writing pages · The two-hour checklist

3.2 The fact table and the organisation facts page

3.3 Person pages and six-trace alignment

3.4 Third-party credential tiers, Wikidata and the five business profiles

3.5 The brand six questions, and what to do when you find an error

3.6 The positioning-statement myth, and the exits from this chapter

3.1 Why this comes before writing pages · The two-hour checklist

What you'll do in this section: confirm the baseline is saved, lay out the route for the identity gate's five fixed moves, and get through the six-part checklist in one sitting, in two hours — then decide whether you can take the 30-minute pass. When you're done, you'll have a first-draft entity-error list, a 14-row fact table, an inventory of the five business profiles, and an exclusions list.

Hard gate: until the baseline is saved, do not change the door, any profile or the facts page. How the baseline is run and saved, including that one set of 36 brand-six runs, is in → General Edition 4.3 Coarse screen and piles, full rounds, and the brand-six baseline.

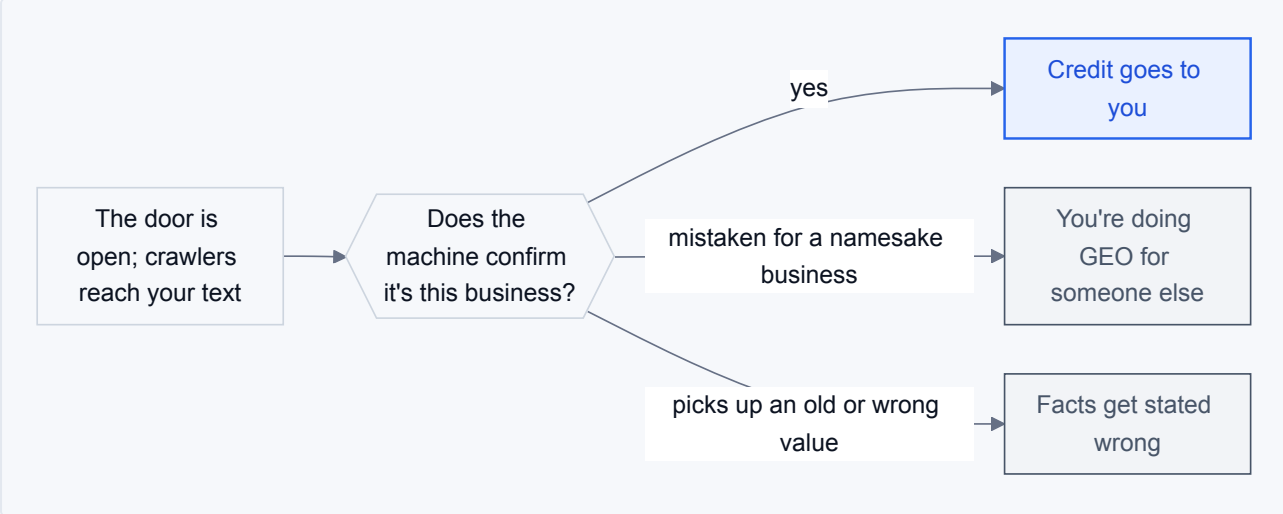


Figure: When the machine identifies the wrong business, all later page work pays back zero, or even less than zero

What gets misstated is usually an old or wrong value AI picked up from a third-party page: price, address, opening hours, registration number, a practitioner's name and credentials, or whether the business has closed. This gate has two features found nowhere else:

- **Fast payoff.** Of every move in the set, it's the only lever where "once you fix it, AI's wording is often the first thing to change" (observed, 70% confidence; not a promise); writing pages, off-site work and video all pay back on a monthly timescale.
- **It bears on the conversion rate.** After a buyer gets a shortlist from AI, they circle back to check the brand — search the name, look at the website, check the profiles — before making contact.

The only firm number that tells you whether this gate is done: the number of factual errors, X → Y. Seat counts move with sampling, and "on the list and cited" waits on a third party saying yes; this is the only cell that gives a firm number every month, and you can reproduce it yourself by asking the questions again (→ General Edition 7.1 What a retest produces, and the rulers).

How far the evidence goes for each claim in this chapter is set out in full in → General Edition A.1 Evidence for mechanics, the door and identity. The conclusion in one line: **AI's live fetch only reads visible HTML,**

so every fact in this chapter has to be written as visible HTML, one field per line. When you cite that evidence, copy the qualifier along with it — never copy just the number.

The route: five fixed moves

1
Unique spelling
Fix one spelling for the organisation's legal name and each named expert's name

2
The /facts page
Visible HTML, one field per line (3.2)

3
Six-trace alignment
Open each original page, copy it cell by cell and compare (3.3)

4
Same-day correction
Send a letter the same day for every cell that doesn't match (3.3, 3.5)

5
sameAs last
For the knowledge graph; comes after the visible HTML

Figure: The identity gate's five fixed moves; sameAs always comes last

The order can't be reversed: sameAs is for the knowledge graph, not for live fetch. Write the schema before the visible HTML and you end up with a site that's "fields all present but AI can't read it" (the schema template is sealed once, see → General Edition 2.7 Gates 3 and 4: indexing paths, and JSON-LD sealed once).

The two-hour checklist

Before you start
Baseline saved
If it isn't saved, don't change the facts page or the profiles

0–25 min
Read the brand six
The 36 baseline runs; copy out the errors and negatives

25–55 min
14-row fact table
Source, verified date, publishable?, owner for each row

55–75 min
Three states of the five profiles
present / absent / squatted; whether the URL points to the canonical domain

75–95 min
Recheck the door read
Check the receipt from 2.2

95–110 min
Backend site fields
Check the receipt from 2.5

110–120 min
Set the exclusions
Topics not to raise, prices not to publish, banned phrases

Figure: Two hours, six parts, with a hard gate up front

The 0–25 minute block is read by the person running probes; the rest is done by the person who edits the site. This part **does not run a separate probe**: it reads that one set of 36 raw answers from freeze week. A separate quick run is not the same ruler as the 36 runs; if X comes from a quick run, X and Y can't be compared. The 75–95 and 95–110 blocks were already done in the door chapter (2.2, 2.5); here you only check the receipts. Canonical domain means the www and https form matches the official site.

Time block	Output	Hit any one of these and you can't shorten it
0–25 min	`entity_errors.md` draft	Any factual error, being mistaken for a namesake business, or tendency judged negative / excluded → jump the queue that same day (3.5)
25–55 min	`facts-table.csv`	On a regulated side, with a published price → check the price row against the internal ledger (the exported price list), tier by tier; don't just copy the old value off your own page — anything stated publicly must be verifiable
55–75 min	`profiles-5.csv`	Two domains coexisting, a name change, a move, or any profile URL pointing to the homepage instead of the landing page → do it in full, onto week 1's completion list
75–95 min	`gate_read.md` receipt	Any one blocked, or "Block AI bots" switched on → back to → General Edition 2.9 Door-layer troubleshooting and this chapter's acceptance checks, push all page work back
95–110 min	Three JSON files + a screenshot of the site fields	When the site parameters aren't matched, the tool silently reads a different site and still returns ok. **If the site fields don't check out, this cell isn't done — don't mark it complete**
110–120 min	`exclusions.md`	Switch A is on (→ General Edition 0.2 Decide which side you are on first): with no written advertising authorisation, not a single outbound letter may go out afterwards

What each line of the six-line door read checks is in → General Edition 2.2 Gate 1 read-only audit: what to check in nosnippet, robots and the WAF. One of them matters for this chapter: **if Applebot isn't allowed through, don't claim the Apple Business Connect profile** (3.4).

On top of this (not part of the 120 minutes; counted per person): one person page for every named practitioner whose name carries trust, 40 minutes to align + half a day to build the page (3.3). An organisation with ≥3 registered practitioners, or any one holding specialist registration (on sides other than the strictly regulated side: or a third-party credential tier) → this must be done in full.

When a 30-minute pass is enough



Figure: Only when all four conditions hit does the 30-minute pass apply; if any one misses, do it in full

How to take the pass: fill in only the 14 hard fields of the fact table, check only that the URL field matches across the five profiles, and move every other step into the monthly review column. **Log the decision to take the pass in the first line of that week's work sheet — it's a call you made, not something you missed.**

3.2 The fact table and the organisation facts page

What you'll do in this section: first build a 14-row fact table as the single source of facts for the whole project; then, following the blueprint, build an organisation facts page (/facts) that machines can copy

whole rows from; on the day you change a price, sync the facts page, the thick pages and the five profiles the same day.

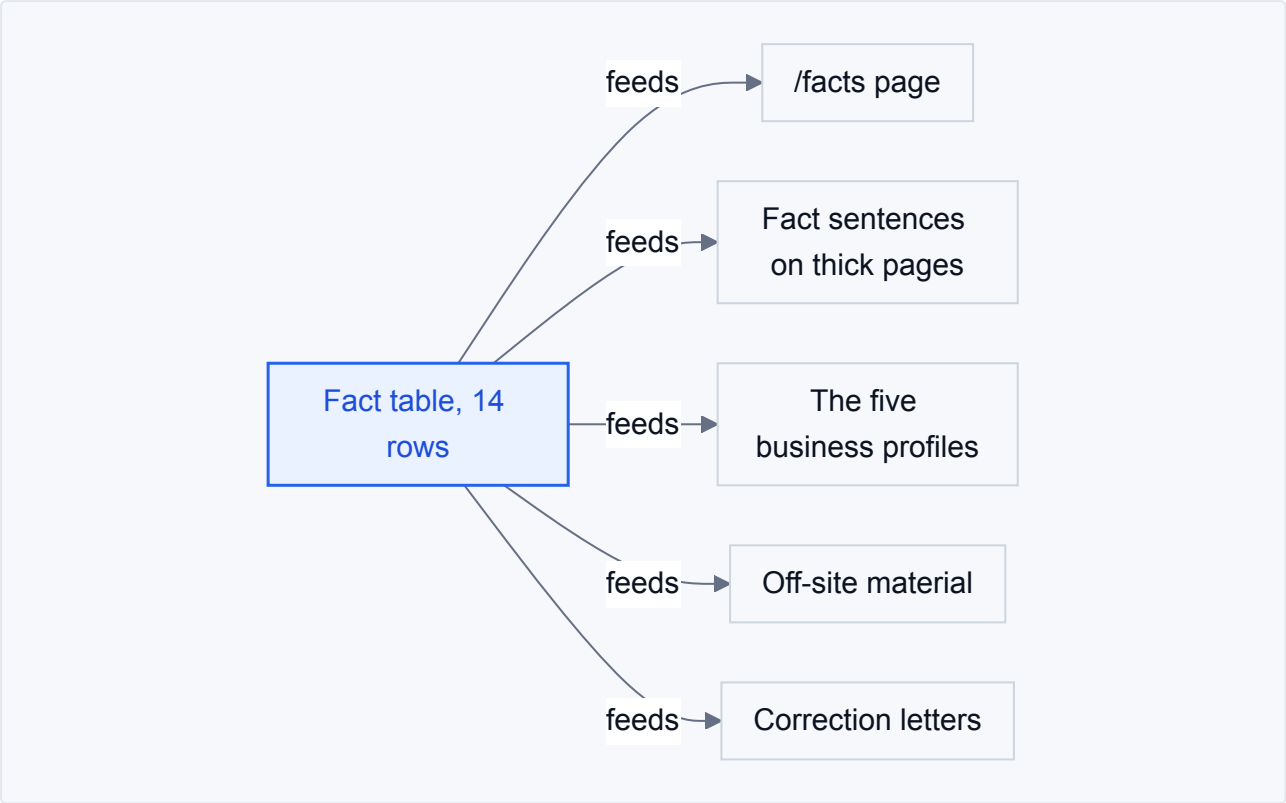


Figure: One fact table is the only source of facts for the whole project; the other five places may only take figures from it

The fact table has four fixed columns: **Source / Verified date / Publishable? / Owner**.

The /facts page: built for machines to copy whole rows, not for error-proofing

The official website's About page doesn't solve this — an About page is made of adjectives; AI can't copy it, and won't dare to. /facts has to deliver: copy it and it's right, copy it and it's complete, copy one row out and it still reads clearly.

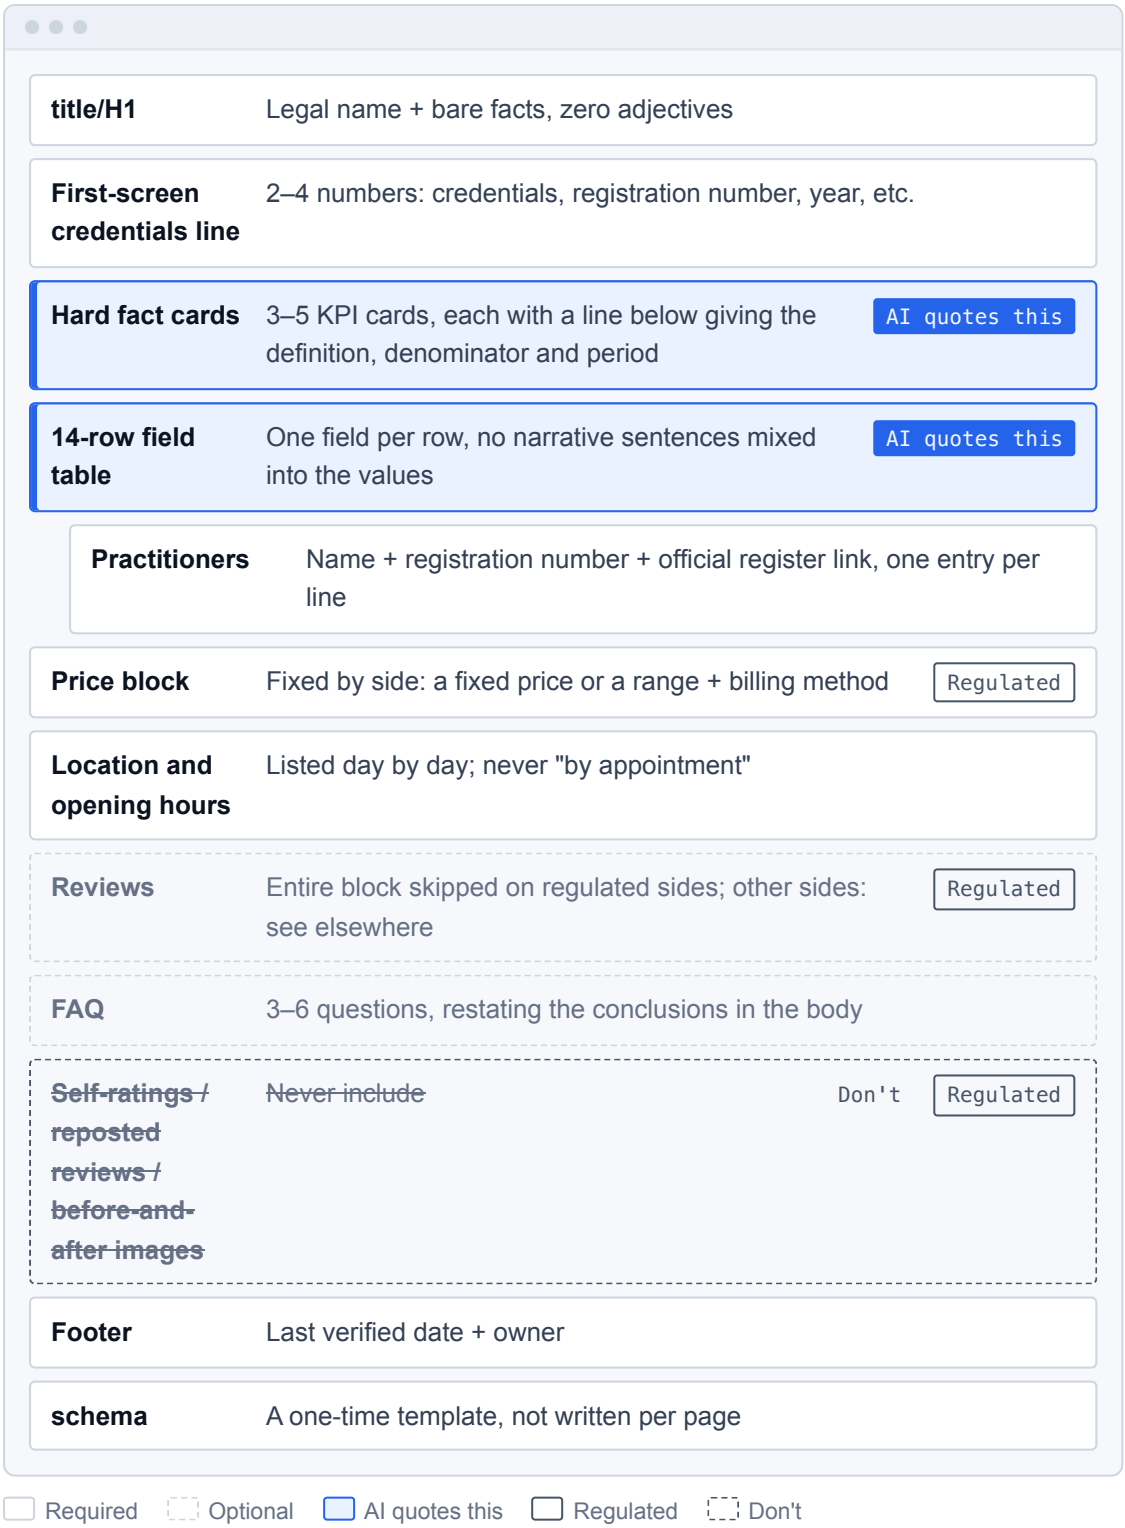


Figure: Block order for the organisation facts page; the field table and hard fact cards are the rows AI copies whole

It needs its own URL, has to be in the sitemap, and be linked from the footer on every page. The form has to be clean HTML — a `<table>` or a `<dl>` — never an image, never a PDF, never JS-rendered: the mainstream AI crawlers run zero JavaScript (only two exceptions; the mechanism is in → General Edition 2.3 Gate 2: four identities, fetched live (read-only)), and if it can't render, the page might as well not exist. The schema follows the one-time template (organisation type + a person type for each practitioner

+ sameAs); it's not written or checked page by page — that was already sealed in → General Edition 2.7 Gates 3 and 4: indexing paths, and JSON-LD sealed once and isn't repeated here.

How to handle reviews differs by side; it's written once, in full, in → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Here it's only marked by position: on regulated sides, skip the block entirely.

This page type relies on numbers, not a conclusion paragraph: the measured organisation facts pages had no conclusion paragraph, no date and no byline before the first H2, and still got cited (evidence in → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types). So every hard fact card needs a line below giving the definition, denominator and period; miss one and it fails the six checks for numbers (→ General Edition 5.10 Shared parts for all page types: six checks for numbers, no public price, billing units, dates and schema).

> **Example** (education and training): An organisation's homepage H1 was a slogan; AI didn't copy that line — it copied the results card right after the first screen, citing the numbers directly. The page has two flaws: one card says "100%" with no denominator, so it fails the number check; and fees aren't in the homepage HTML. The format is worth copying (a number card right after the first screen, readable as plain HTML); the execution isn't (no denominator, price not in the HTML).

Write H1 as `<Legal name> · Organisation facts file`, bare facts, zero adjectives; title can use `<Service/subject> <organisation type> Singapore | <one checkable fact>`, or `<Service> in Singapore — <Organisation name>`. Neither one carries a slogan.

> **Example** (education and training): A title that has been cited: "JC H2 Math Tuition Singapore | 90% A/B" — the format is worth copying (service + location + one number), the wording isn't (90% has no definition or denominator).

14-row field table (hard fields)

Each row = field name + value + last verified date, with no narrative sentences mixed into the value — a row has to be copyable on its own.

#	Field	How to write it
1	Legal name	Word for word as on the business registration, including capitalisation and the company suffix
2	UEN	The number

#	Field	How to write it
3	Industry licence or registration number	The number + the official register link
4	Practitioners	Each person's name + registration number + official register link, see 3.3
5	Address	One line + postal code, word for word matching the five profiles
6	Phone	Word for word matching the five profiles
7	Opening hours	Listed day by day; never "by appointment"
8	Year founded	The number
9	Service languages	Listed item by item
10	List of main services	Service names listed item by item, no adjectives
11	Devices / materials	Real brand + model + link to the manufacturer's product page; on the strict side, write only the model — no link to the manufacturer (3.4)
12	Insurance / subsidy coverage	A firm sentence
13	Follow-up and replacement policy	Only the time limit, the number of times, and what's already included in the quoted price; never "free", "complimentary", "at no extra charge" or "$0"
14	Last verified date	YYYY-MM-DD

> **Example** (dental and aesthetics): Write row 12 as a firm sentence like "this item is not payable under a given medical savings scheme", not "partly subsidised" or "depends on the situation".

The price block: fixed on the strict side

Banned: write it and you get fined		Rewrite: put this in the same cell
"starting" / "from" / "as low as"	→	Tiered, itemised fixed prices
A price range (S$X–S$Y)	→	Split into tiers, each with its defining conditions spelled out
Original price, discount, percentage change	→	Keep only the current price, with one effective date
Instalments, bundles, time-limited promotions	→	One price per tier + defining conditions + what's included

Figure: The strict-side price block; every banned item has a legal replacement that fills the same cell

The price has to be a fixed final price. The detailed wording for each side, and whether a range may appear, is set out once in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers; it isn't repeated here.

> **Example** (dental and aesthetics): A three-tier price list split into mild / moderate / complex, each tier spelling out the defining condition, the tax-inclusive fixed price, what's included and the exclusions — that's what strict-side wording looks like; the lightly regulated side doesn't have to copy it.

Price sentences come in two kinds: an own sentence gives the full price, the billing unit, the number of sessions and what's included; a market sentence gives a range and the variables that affect price. Whether a market sentence may be written at all, and whether it may sit on the same page as your own price, is set by side in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers.

> **Example** (education and training): Own sentence: "<Centre> charges S$480 per month for 4 group <subject> lessons of 1.5 hours each, including GST and notes." Market sentence: "<subject> tuition in Singapore generally costs S$X–S$Y per month, depending on group size, level, and tutor experience."

Some industries' advertising rules say this directly: when the ad space can't fit it, you must give a prominent link to the specific page on the official website, and the source it points to must not be spoken-only, with no record. **/facts and the price guide page therefore have a statute behind them — they aren't a spec we invented** (Statute text; example industry is training, rule text in the industry

edition's appendix — the training edition isn't out yet, and the *Dental & Aesthetics Edition* appendix doesn't cover this one).

How /facts and thick pages divide the work

/facts is a source that the density gate draws numbers from, not something the gate checks directly: it isn't bound by the thick-page word count, doesn't count against the intent-cluster quota, and doesn't count toward the page total. Every "checkable fact that can't be found elsewhere" on a thick page must point back to a specific row on /facts and its fact-table row number; one that can't be pointed back doesn't count. It also feeds the source gate: the register links and government pages that rows 2, 3, 4 and 12 already carry (plus, on the side that may link out to manufacturers, row 11's manufacturer product page) are the ready-made source for those 3-plus outbound links on the thick page. /facts is still an advertisement: on regulated sides it still has to pass the pre-publication compliance checklist — banned words, price wording, "free" wording, efficacy and peer comparison, regulatory endorsement — not one item is exempt; the checklist is set out in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers and → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.

On the day you change a price, sync /facts, the price table on thick pages and the price field in the five profiles the same day, and send an update letter to any third-party page that already lists your price. If the three don't match, you've smudged a fact you signed off on yourself.

How to check: /facts is live; fetch the page source with an ordinary browser UA and search it: the price string and the registration-number string are both there; price and registration number match the fact table word for word, 100%.

3.3 Person pages and six-trace alignment

What you'll do in this section: build a separate person page for every named professional; make the name, registration number, organisation name and address match word for word across six public traces; and when someone leaves, handle sameAs in one of three ways.

This applies to any role that carries the buyer's trust under a person's name — not just the organisation's founders; front-line service providers need one too.

> **Example** (dental and aesthetics): Every named doctor at the clinic counts, not only the founder.

On the regulated sides, almost every form of persuasive content is banned at the organisation level — but not at the person level. The regulator explicitly allows a practitioner to state credentials, scope of practice, practice arrangements and contact details; this is one of the few sources of persuasive content explicitly allowed on the regulated sides, so **on the regulated sides a person page isn't a bonus, it's**

the main line (Statute text; rule text in the industry edition's appendix (example industry: *Dental & Aesthetics Edition*, Appendix A)).

> Red line **Allowed ≠ unregulated**: personal information still has to pass the nine compliance criteria (truthful, accurate, verifiable, not exaggerated, not misleading, not inflammatory, not persuasive in tone, not comparative, not disparaging); and there's a separate ban on laudatory terms from the example industry's regulator (e.g. the Annex A laudatory terms table) — never describe this person-level allowance, in any external document or spoken statement, as "a back door around the organisation-level advertising restrictions." Title formatting must follow the format the regulator sets out, and a title not legally held may not be used; anyone not on the specialist register may not imply specialist status (the notice that sets out title formatting could not be machine-read; confidence in this rule rests on a manual check, and it must be checked against the original, clause by clause, before launch; Original text not obtained).

Pure directory listings are the lowest-citation-density page type in the measured data; the largest single measured effect is the entity-strength multiplier (a single third-party study, a US sample; evidence and qualifiers in → General Edition A.1 Evidence for mechanics, the door and identity) — money spent on directory listings buys nothing, money spent on the entity does.

> Red line **The output is a page, not an internal table**: pulling this together into an internal credentials table doesn't count as done — AI can't read an internal table, only visible HTML; not one of the five major crawling systems reads JSON-LD or hidden structured markup (evidence in → General Edition A.1 Evidence for mechanics, the door and identity).



Figure: Block order for the person profile page (this whole page type is not measured, so nothing is marked cite)

Fill in the external-anchor field when it exists; don't invent one when it doesn't — especially the academic identifier and the Wikidata QID; whether to build the latter depends on the threshold in 3.4. Third-party credential tiers give only the tier name and the awarding body — see 3.4. Write scope of practice and service prices the same way as the /facts price block; on regulated sides they likewise go through the compliance checklist once (see 3.2).

Each person gets their own URL (`/team/<name>`), not a paragraph inside the About page. The form matches /facts: a visible `<dl>` or `<table>`, one field per row, every external source a clickable `<a href>` — never JS-rendered, an image or a PDF. After that, every thick page, every video and every piece of off-site material carries this person's byline, and the byline links back to this page — that's the only move that ties page authority to the person's entity.

Fix title/H1 as `<Name>, <registered title>`; on regulated sides the title uses only the registered wording, never a self-styled one. On regulated-side professional content pages, write a reviewer line as the first line under the H1: the reviewer's name + degree + month and year, followed by the registration number, linking back to the person page; on other page types treat it as a bonus (the measured data on the byline line is in → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types).

> **Example** (dental and aesthetics): Under the H1 on every medical page, write `Medically reviewed by Dr X, <degree> · <month year>`, followed by the registration number.

The real mechanism behind "8x entity strength" (a single third-party study, a US sample; qualifiers in Appendix A.1) is not what's written in the schema — it's the same set of facts appearing word for word across several independent third-party pages that can be crawled.

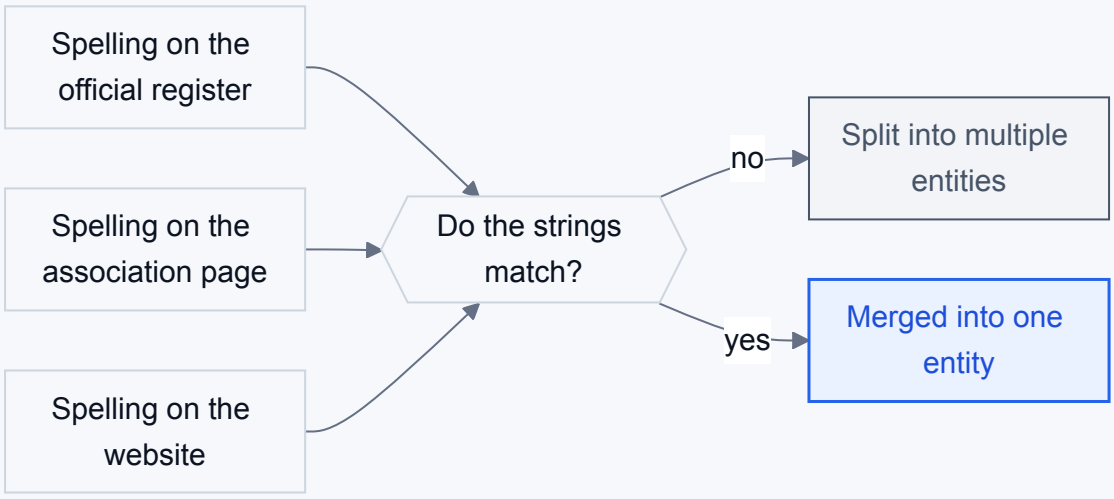


Figure: Three spellings of the same person; the machine sees three people

Three strings are three entities, and 8x becomes 1x. So the real lever is the six-trace alignment table below, not the page spec table itself.

The six traces change with the anchor layer

Anchor layer	Six public traces	Where the standalone page is built
Organisation	Official register, the five business profiles, industry association profile, the website, manufacturer locator, media coverage	/facts
Individual	Official register page, industry association profile, staff page on the website, conference speaker page, professional social profile, academic identifier	/team/<name>

The line that holds across every industry: an inconsistent spelling is the number-one reason an entity gets split into two.

> **Example** (e-commerce): The same brand goes by three different names on three platforms; the machine sees three entities, and none of the three sets of social proof adds to another.

Some industries have explicit rules for how the legal name itself must be written, and a non-compliant spelling also splits the entity in two: the spelling on business profiles, list entries and platform pages has to be changed at the same time, or you break the rule and split the entity at once (Statute text; example industry is training, rule text in the industry edition's appendix — the training edition isn't out yet, and the *Dental & Aesthetics Edition* appendix doesn't cover this one).

> **Example** (education and training): The registered name may not have "Singapore" added after it; to show the location, the only option is to write "located in Singapore" separately.

When an unrelated namesake entity is polluting your results, add one exclusive definition sentence to the first paragraph of /facts: product name + category noun + legal entity name + location, all in one sentence; then make that sentence the first sentence of the description in all six traces (the lift from this is not measured, but the mechanism is clear).

> **Example** (software): A search for a product's SOC 2 status returns an answer that mixes in an unrelated company with the same name.

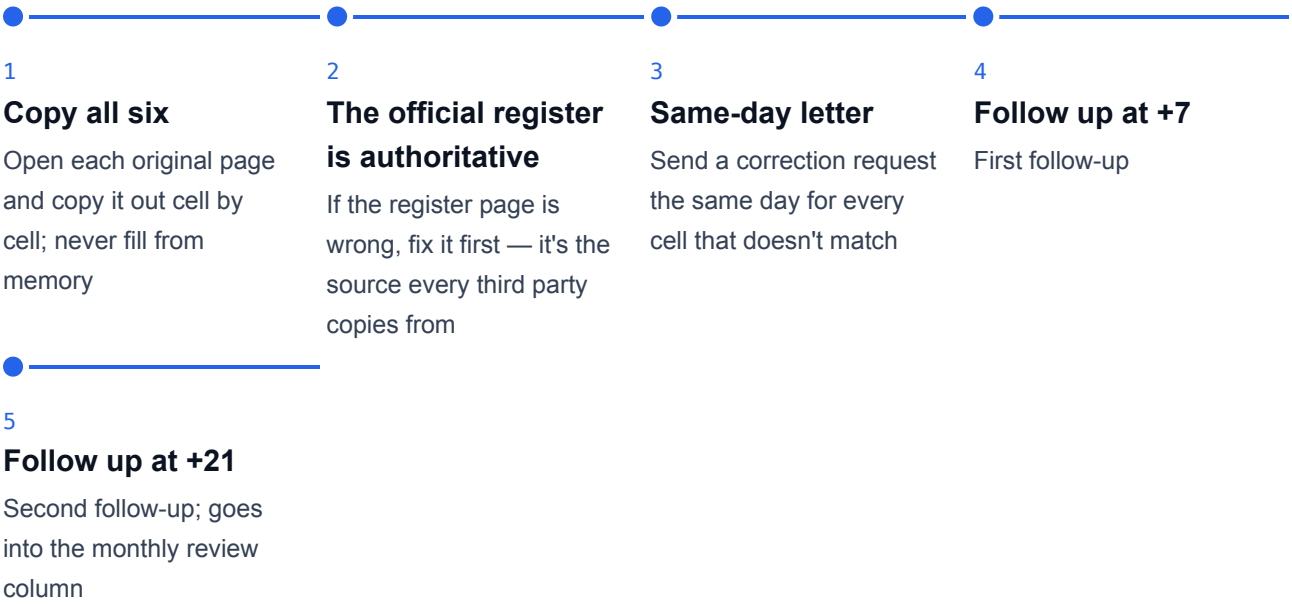


Figure: Align all six cell by cell, close out with two follow-ups

Source	Name, exact spelling	Registration / practising number	Organisation's legal name	Address
Official register page				
Industry association profile page				
Organisation's staff page				
Conference / speaker page				
Professional social profile				
Academic identifier (only if one exists)				

How long: 40 minutes to align each person, half a day to build the page. **Why it moves rankings**: machine entity disambiguation runs on strings matching consistently across multiple sources, not on how many times you repeat it on your own page. **How to check**: the person page opens on the public internet; a string search of the page source finds the name, the registration number and (on sides other than the strict side) the tier name; all six rows × four columns are filled in and match word for word, and every cell that didn't match has a correction letter already sent, with the date it went out; the byline on thick pages and videos clicks back to this page.

When someone leaves: keep the page, handle sameAs three ways

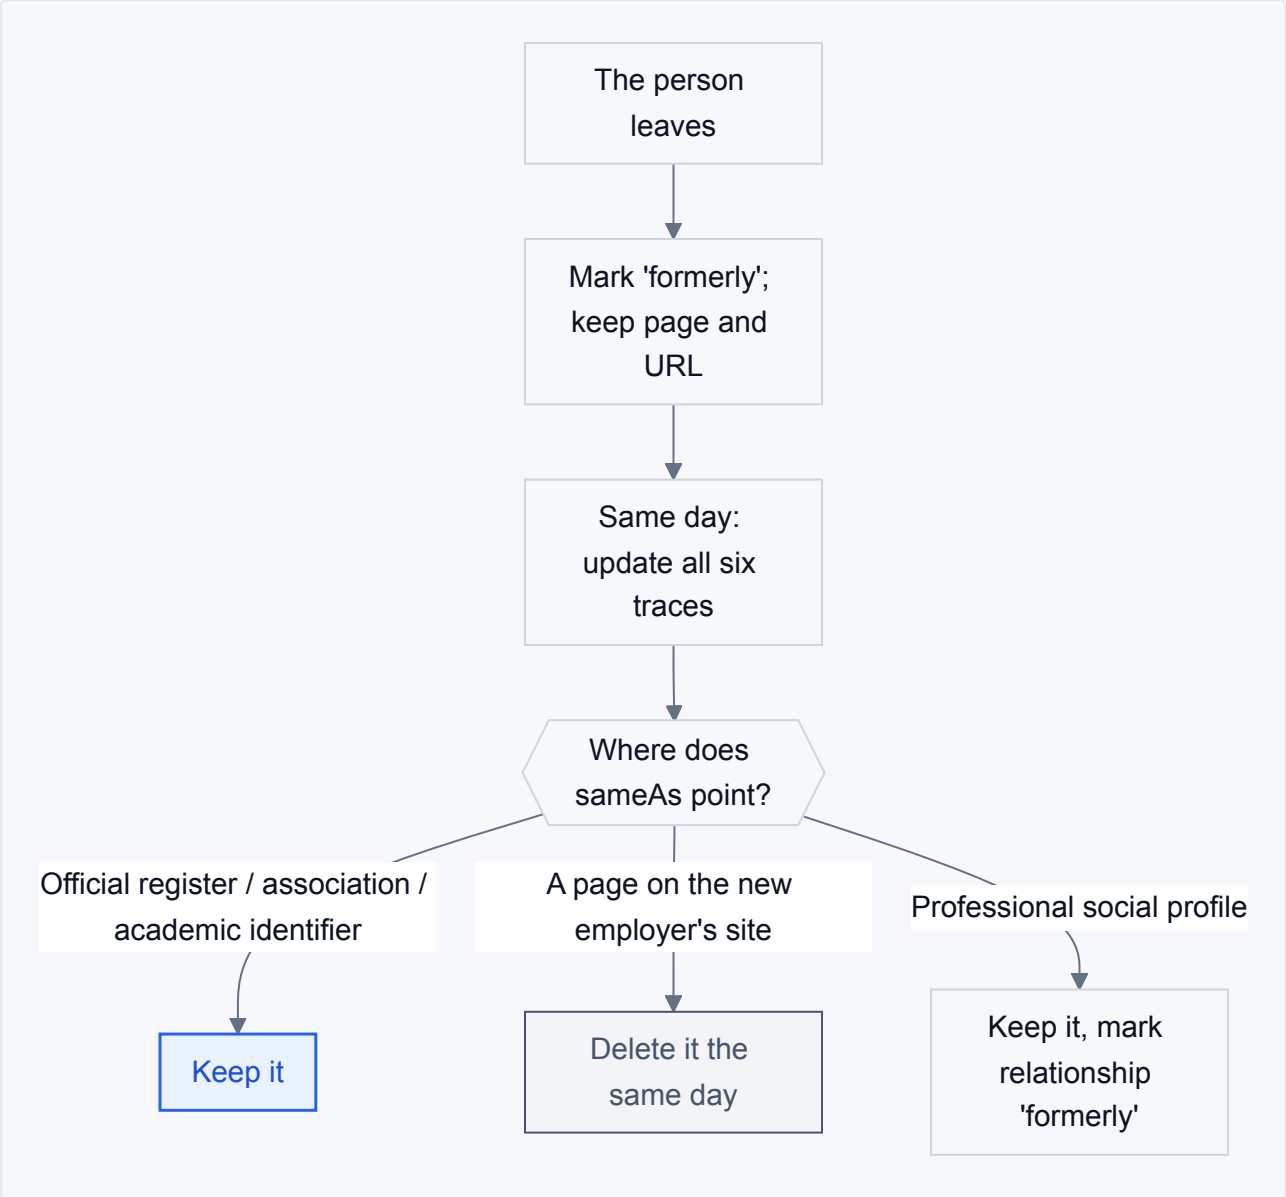


Figure: When someone leaves, keep the page and change its status; handle sameAs three ways

The page isn't deleted because deletion would break inbound links already built. A sameAs pointing to the official register, the association or an academic identifier stays because it proves who this person is; a sameAs pointing to the new employer's page is deleted because leaving it in place feeds authority to a competitor, and might even let entity resolution assign this page to the new organisation. This treatment isn't measured, but the mechanism holds.

There's no such thing as "301 a byline": a 301 acts on an entire URL; a byline field can't be redirected. Only a profile page with genuinely no remaining value gets 301'd whole to the team directory; every other page just has its current owner and author link changed (not measured).

The six places to finish in one pass, the same day: the status line + last-checked date on the person's own page, the organisation's current-staff list, any field across the five business profiles that carries

their name, the industry association or platform profile, list entries (send an update letter), and the opening link on their bylined content (point it to the "formerly" page, without deleting the content itself).

When someone new takes over, any sentence on any page that cites the outgoing person's credentials must be changed the same day, or the old and new versions sit side by side — which is itself a self-inflicted factual error.

Where the person is themselves the anchor, and their leaving would take the whole service line with them, the landing page must not be built under that person's name: the subject has to be the service itself, with the person only a replaceable "current owner" field — otherwise, once they leave, the whole page either gets deleted or turns into a wrong entity (not measured, but the mechanism holds).

> **Example** (education and training): For a one-on-one tutoring course page, the subject is "subject × level", and the tutor is only a "current teacher" field.

3.4 Third-party credential tiers, Wikidata and the five business profiles

What you'll do in this section: go get a third-party credential tier wherever you can — it's the substitute for when you can't ask for reviews; apply for manufacturer-locator listing wherever you're eligible; build a Wikidata entry only once the conditions are met; complete all five business profiles in one pass, with the URL field pointing to the landing page, then spend just 20 minutes a month on the three-state check.

Third-party credential tiers: the substitute when you cannot ask for reviews

On the light and unregulated sides, fill in a verifiable third-party tier wherever one exists, and if it can't be looked up, write the date; on the strict side, never write a manufacturer-granted tier (Conservative line (not statute text); basis in the industry edition's appendix; loosening this needs the client's compliance officer's written reasons) — write only technical-standard certification and registration status.

Three hard rules for how to write it:

1. Write only the tier name and the awarding body. "highest", "only", "leading" or "only X on the whole island" hits the banned-word list directly and is banned regardless of whether it's true.
2. On the strict side, case counts and procedure volumes are deleted by default — there's no confirmation process, and no "usable once confirmed" tier; on the other sides, write only your own absolute numbers, and pass all three gates at once: no adjectives, a ledger record signed by the owner, and never placed next to efficacy claims.
3. Tick each item off against the compliance checklist, with the owner's sign-off.

A written sign-off resolves only the three stop-and-escalate situations; it never turns something banned as Statute text into something publishable: the words in rule 1 can't go live even with a sign-off (the stop-and-escalate rules are in → General Edition 0.4 The three labels for compliance sentences, and the stop-and-escalate rule).

Law firms and similar can cite a third-party ranking as evidence of expertise, but which side it falls under is decided by the industry edition; it must pass that industry's publicity-rules review first, and cannot go live before legal signs off.

How to get a tier: apply to the manufacturer for listing in the official dealer / service-provider locator, using your purchase records as evidence; it's usually free, and one application pays off in two places — an inbound third-party link, and third-party endorsement of the real brand and model of the device you use.

Direction	Who links whom
Listed in a manufacturer locator	They link to you, inbound
The strict side's own page	Write only the model and registration facts, never link out to the manufacturer

Figure: A manufacturer locator is an inbound link — you don't link out to the manufacturer

Why it doesn't work the other way round: a strict-side page linking out to a commercial company runs straight into the ethical code's own wording (Statute text; rule text in the industry edition's appendix (example industry: *Dental & Aesthetics Edition*, Appendix A)); this same direction rule also governs where effort goes in the off-site chapter and the substitute tactics, see → General Edition 6.4 Three source groups, and the Group B actions and → General Edition 8.7 Cheap tactics on the unregulated side, and substitutes on the regulated sides — not repeated here.

Wikidata: required when the conditions are met, not open to all

Wikipedia is the single largest node in ChatGPT's citation set, and Wikidata is its structured anchor (numbers in → General Edition A.1 Evidence for mechanics, the door and identity). Its job isn't to get cited — it's to give the entity a unique ID: however long the sameAs chain gets, if there's no node in the graph to catch it, disambiguation still falls back on string matching.

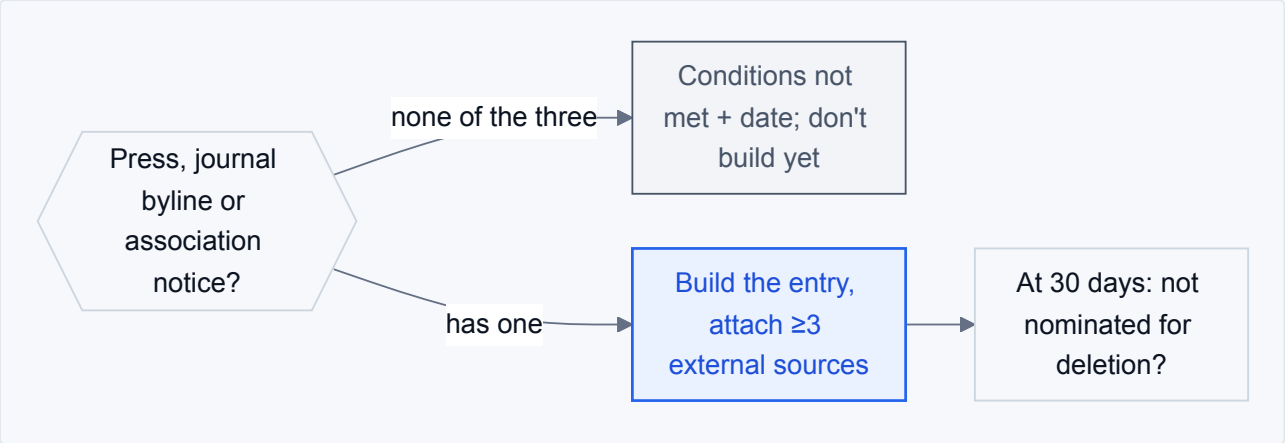


Figure: Whether to build Wikidata depends on whether there's a citable third-party source

The common misconception is "anyone can build one, and once it's built, it's there" — not true. Admission requires one of the three, and that one has to be describable using serious, publicly available reference material. A practitioner with only a website and a professional social profile who builds one will get it nominated for deletion by the community — you lose the time and also leave a deletion record in the edit history. With none of the three in hand, wait until off-site bylined contribution lands the first piece of third-party coverage, then build it — that's the interface between bylined contribution and this chapter.

Don't edit Wikipedia's body text to add the organisation's name or an outbound link — a self-citation gets reverted and leaves a record; the risk outweighs the benefit. The only thing worth doing is adding a verifiable source to an existing entry, without adding the organisation's name.

Building the entry takes 30–40 minutes; how to check: the QID opens on the public internet, the entry contains ≥3 external references, and a recheck after 30 days shows it hasn't been nominated for deletion.

The five business profiles: an entry ticket, not a lever

The verdict is fixed: this is an entry ticket, not a lever. Complete it once, then just run the three-state check once a month afterwards — never open an "optimisation" task.

1
Foursquare
Claim it + basic facts; pick the category word for word from the official category tree

2
Yelp SG
Claim it + basic facts; on the strict side, check the facts only, don't engage with reviews

3
Google Business Profile
Category + services + opening hours + attributes + photos; the URL is the only field with real leverage

4
Bing Places
Same as above; one of ChatGPT's eight retrieval sources, not the only one

5
Apple Business Connect
Precondition: Applebot must already be allowed through, see 2.2

Figure: Order for claiming the five profiles, ranked by which leg benefits first, not by platform fame

Why the order isn't ranked by platform fame: ChatGPT's local answers run mainly through the first two (70% confidence); on the Google leg, 95.5% of cited domains are exactly the URL registered on the profile (evidence in → General Edition A.1 Evidence for mechanics, the door and identity) — and that's the only field in this whole cell with real leverage. The URL on all five profiles should always point to the landing page, never the homepage, each one matched to the page for that buyer type.

Externally, say only the fixed line "a basic step, done in passing" (→ General Edition B.2 Approved wording and scripts) — never "do this and ChatGPT will recommend us": the mechanism confidence for Foursquare, Apple and Yelp is only 70%, and the researchers themselves say there's no data isolating which profile field contributed to the citation.

The compliance scope uses the same word list as the website pages: every free-text box, business Q&A, business post and photo caption on the five profiles is an advertisement. The two places most often missed are photo captions and business posts. Business post topics are limited to: a change in opening hours, a new service added, a change of address, an update to a practitioner's registration information — never a promotional post, never a case post, never reposting a review. Fill in the price field the way /facts does it, by side (3.2); on the strict side, only the fixed final price.

How to ask for reviews and how to reply differs between the two sides; it's written once, in full, in → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Here, just one line: whether you may actively ask for reviews follows the rule for your side, and the answer is in that section, not repeated here.

Monthly three-state check (about 20 minutes): ① present or not ② is the URL field correct ③ has a fact been changed by a third party or the platform. Tick all three off in the monthly report.

How long: completing it once takes about 2 hours; 20 minutes a month after that. **How to check**: all five are claimed; name / address / phone match /facts word for word; every URL field points to the landing page's canonical domain; each one's launch date is logged in the outreach ledger and added to the monthly review column.

3.5 The brand six questions, and what to do when you find an error

What you'll do in this section: starting from the 36 brand-six runs in freeze week, count factual errors every month with the same ruler; within 1 working day of finding an error, log it and notify; within 3 working days, fix every source you control and send a correction letter; when any of the five trigger conditions is hit, jump the queue and fix it the same day, without waiting for the monthly report.

What the six questions ask, and how the ruler is made up

The brand six is defined once for the whole book, and every mention of brand questions anywhere else in the book points back to this section — **change the composition and you've changed the ruler, and the factual-error counts you get out aren't comparable any more.**

#	Question	What it measures	Can it be used publicly?
1	`What does <Legal name> do`	Identity	Yes
2	`Is <Legal name> reliable / what's it like`	Tendency	Yes
3	`Where is <Legal name>, opening hours, phone`	Basic facts (measures the profiles leg)	Yes
4	`How much does <Legal name>'s <main service> cost`	Price misreporting — the most valuable of the six	Yes
5	`Does <Legal name> have complaints / bad reviews / disputes`	Negatives	Yes
6	`<Legal name> vs <competitor>, which is better`	Comparison	No — internal material, names a third party

Ruler composition (the one definition for the whole book): 6 questions × 3 rounds × 2 engines = 36 runs / month, frozen for the whole quarter. The baseline uses only the full-round set of 36 runs from freeze week (4.3), never a quick run — two sets with different compositions aren't the same ruler, and the numbers that come out aren't comparable.

Ruler	Composition and denominator
The brand six	6 questions × 3 rounds × 2 engines = 36 runs, denominator 36, the one definition for the book
The two Chinese questions (run only if a Chinese page exists)	Chinese versions of questions 1 and 4: 2 questions × 3 rounds × 2 engines = 12 runs, denominator 12

Figure: The brand six and the two Chinese questions are two separate rulers; their denominators are never combined in one table

The two Chinese questions run on a separate ruler, saved to their own file, and are never merged into the brand-six number: merge them and you get 48, but only someone with a Chinese page can ever produce 48 — the same "factual errors" count would instantly stop being comparable between businesses with and without a Chinese page. Both numbers go into acceptance, but each keeps its own numerator and denominator. **Print the two numbers separately**: "Brand six: factual errors X → Y (denominator 36)" and "Two Chinese questions: factual errors X → Y (denominator 12) · wrong-anchor N". Wrong-anchor logs whether the answer quoted a price from the wrong region, or treated an out-of-town business as a local one.

Log four dimensions, read by a person, no skipping

Dimension	Four states	How to log it
Identity	This business / namesake business / merged (two businesses' facts merged into one) / unknown	Namesake and merged both trigger jumping the queue directly
Facts	Checked one by one against the fact-table row number	Log the error count X — the only firm acceptance number
Tendency	Positive / neutral / negative / excluded	The last two states both trigger jumping the queue directly
Source	Every URL listed under the answer	Trace each false statement back to its source to decide who the correction letter goes to; anything that can't be traced is marked "source unclear" and retried next month

How long: running the full 36 takes about 1 hour; reading takes a person 3 hours. **How to check**: every raw answer is saved; every error in `entity_errors.md` has a screenshot of the original text + a source URL + the matching fact-table row number; not one is marked "unread". **Why it moves rankings**: it doesn't, by itself — it's this chapter's only ruler, and the only signal that can prove both legs move once the source is fixed. AI's factual errors are copied from the source pages it retrieves; fix the source, and both legs move.

After you find an error: the correction path

T+1 working day

Log and notify

A written record + notification, with a screenshot of the original text and the sampling spec

T+3 working days

Fix every source you control

/facts, the five profiles, the relevant website pages; keep before-and-after screenshots

T+3 working days

Send correction letters

A correction letter signed by name to every third-party page with the error, each one filed

T+7 / T+21

Two follow-ups

Log the outcome of each

Next month's retest date

Retest the same question

Log the change in the error count

Figure: Timeline after finding an error, closing out with two follow-ups

Whether and when a third party acts on it is their call, not a promise you make; what you can promise is two follow-ups, checking every month whether the error is still there, and re-checking AI's wording question by question. How to write the correction letter and the template wording is in → General Edition B.2 Approved wording and scripts; here, just one line: point out only the error, attach a source link to /facts, never mention rankings, never slip in promotion.

Precondition for sending a letter (a hard gate on regulated sides): if switch A (agency liability) is on, missing either the written advertising authorisation or the sign-off on external material means not a single correction letter may be sent. For list pages whose title is itself laudatory (Best / Top N types), regulated sides make no submission and buy no paid listing — disclosing the payment is not a release from liability; the regulator has specifically answered this question (rule text in the industry edition's appendix (example industry: *Dental & Aesthetics Edition*, Appendix A)). If you are already listed on one, the strict side asks the publisher to remove you and does not send a correction letter (reasons in → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions); the light side follows its industry edition's rules and may correct only the one line that is wrong.

If the same error hasn't cleared after two months in a row: escalate for a review, and suspect first that you haven't found the right source — the error may be coming from a third-party source you haven't identified yet.

Make before-and-after screenshots anyone can reproduce

Element	Spec
Comparison	Same question, same engine, same region setting, same time window; before and after side by side, the sentence circled in red
Qualifier	Burned into the image, not written beside the body text: "Single sample · single engine · directional, not acceptance evidence; the same question gives different answers at different times"
Wording	No causal language: only two parallel statements of fact are allowed — "On <date X>, Y was changed on source Z" and "On <date X>, this is what AI said" — never "because / therefore / it worked / it's taking effect"
Closing caption line	This image is not acceptance evidence; acceptance goes by the three numbers in the monthly report
Try-it-yourself card	Kept alongside every screenshot: the exact question, the engine name, and the steps to set the region — in three steps or fewer, so anyone can reproduce it. This card is a deliverable, not an extra
Cases naming a third party (question 6)	Marked internal material, the qualifier cannot be removed, and it may only go external with the name redacted

Three rules for anything public: ① Show only the current state — never "I can change it to X" — and no quantity words or timeframe words at all. ② Any price figure AI has made up is shown only as a risk illustration, never written into any external deliverable, and nothing may suggest that anyone can control AI's output. ③ Question 6 (versus a competitor) never goes external.

Entity errors jump the queue

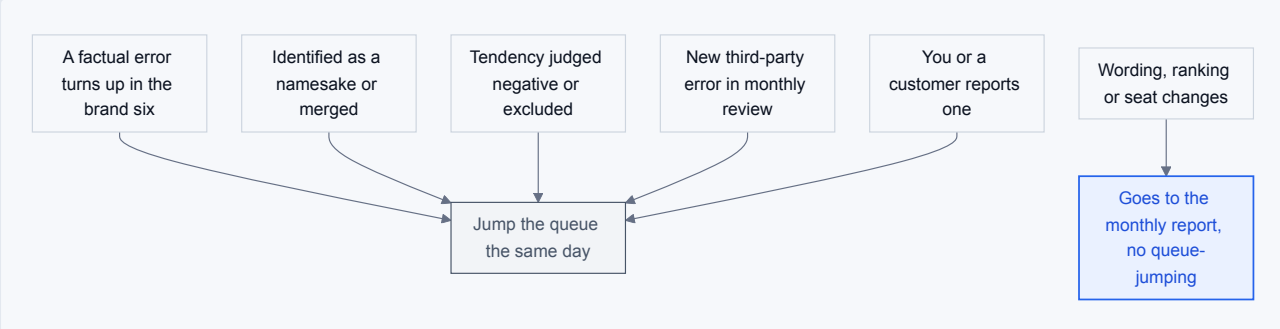


Figure: Five situations jump the queue and get fixed first; everything else goes to the monthly report

Also not a trigger: an increase or decrease in being listed as a source, or a competitor entering or dropping out — these go to the monthly report, not the queue.

Cost and red lines: jumping the queue draws on that month's content-page hours — it's an open cost, not extra time found from somewhere; it doesn't change the acceptance measure, and even once it's fixed, the only thing you may write is "baseline X errors → this month Y errors", never "I got AI to change its story"; no new content pages during a queue-jump. **The clock starts at "discovery": discovery = anyone reading it in any batch of answers.**

3.6 The positioning-statement myth, and the exits from this chapter

What you'll do in this section: don't spend 3–4 hours writing one "positioning statement" — put that time into 6 answer sentences for the first intent cluster instead; then check against the checklist at the end of the chapter to decide whether you can move on, or whether you need to jump the queue, push everything back, or do website-side work only.

Why a positioning statement does not work

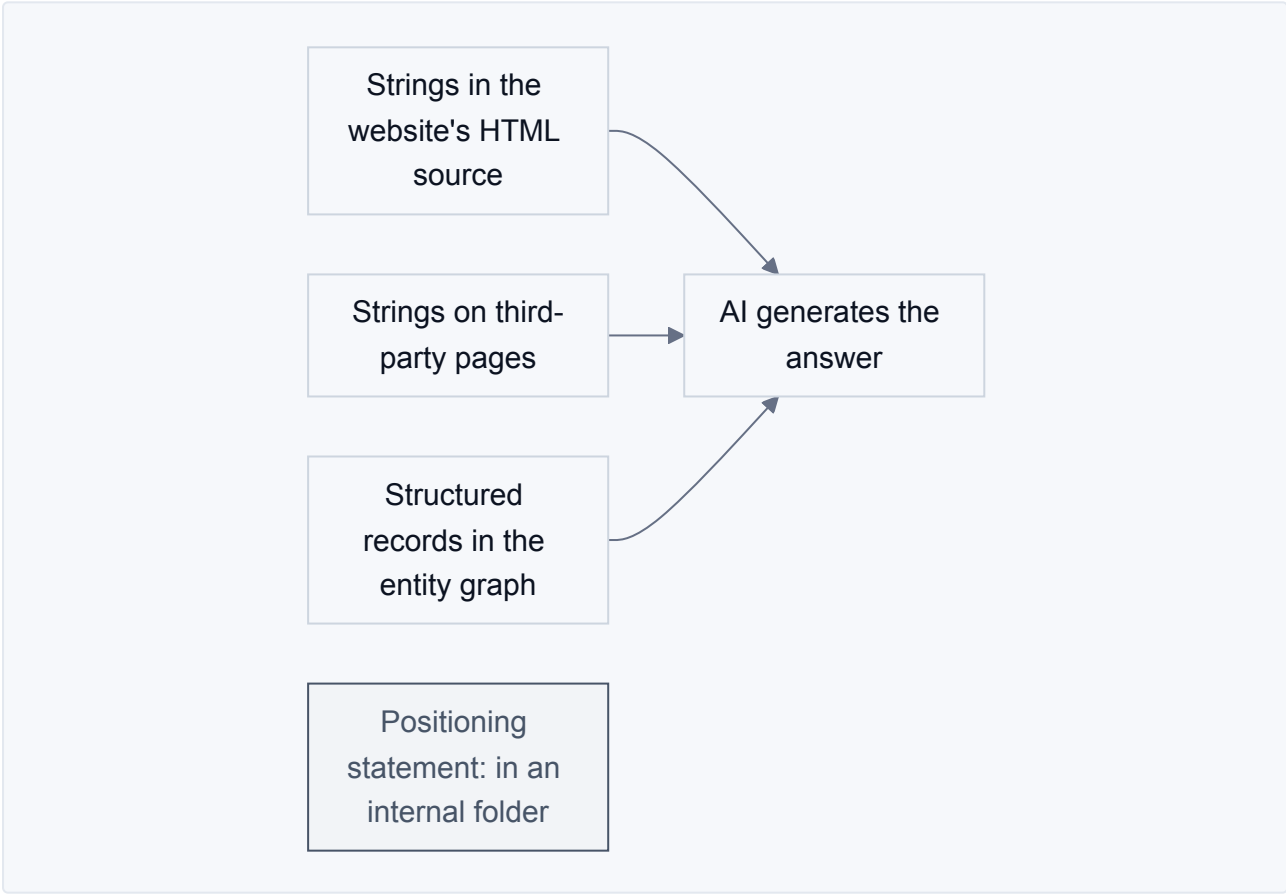


Figure: AI reads only three things, and the positioning statement isn't one of them

However solemnly the positioning statement is signed off, it still only exists in a folder. The real, legitimate outlet for that need is the **answer sentence**: the first sentence under every H2 on a thick page, containing the brand name plus one number or one clear condition, and readable without context — that's a string on the page, and AI can copy it. A positioning statement isn't.

Take the 3–4 hours that would have gone into the positioning statement, and spend it instead writing 6 answer sentences for the first intent cluster — finalised and added to the thick-page brief the same day; use the time you save for the person pages in 3.3 (40 minutes × 3–5 people).

A red line on wording: never say publicly "I'll get AI to describe us this way" — the alternative wording is in → General Edition B.2 Approved wording and scripts. Writing it correctly is your job; whether it gets copied is the machine's.

Three more rules for regulated sides: every public statement must be verifiable, with a source and verified date on each line; personal credentials are not a back door around organisation-level advertising restrictions; on the strict side, price may only be written as a fixed final price — never a range, never "from".

Exits from this chapter

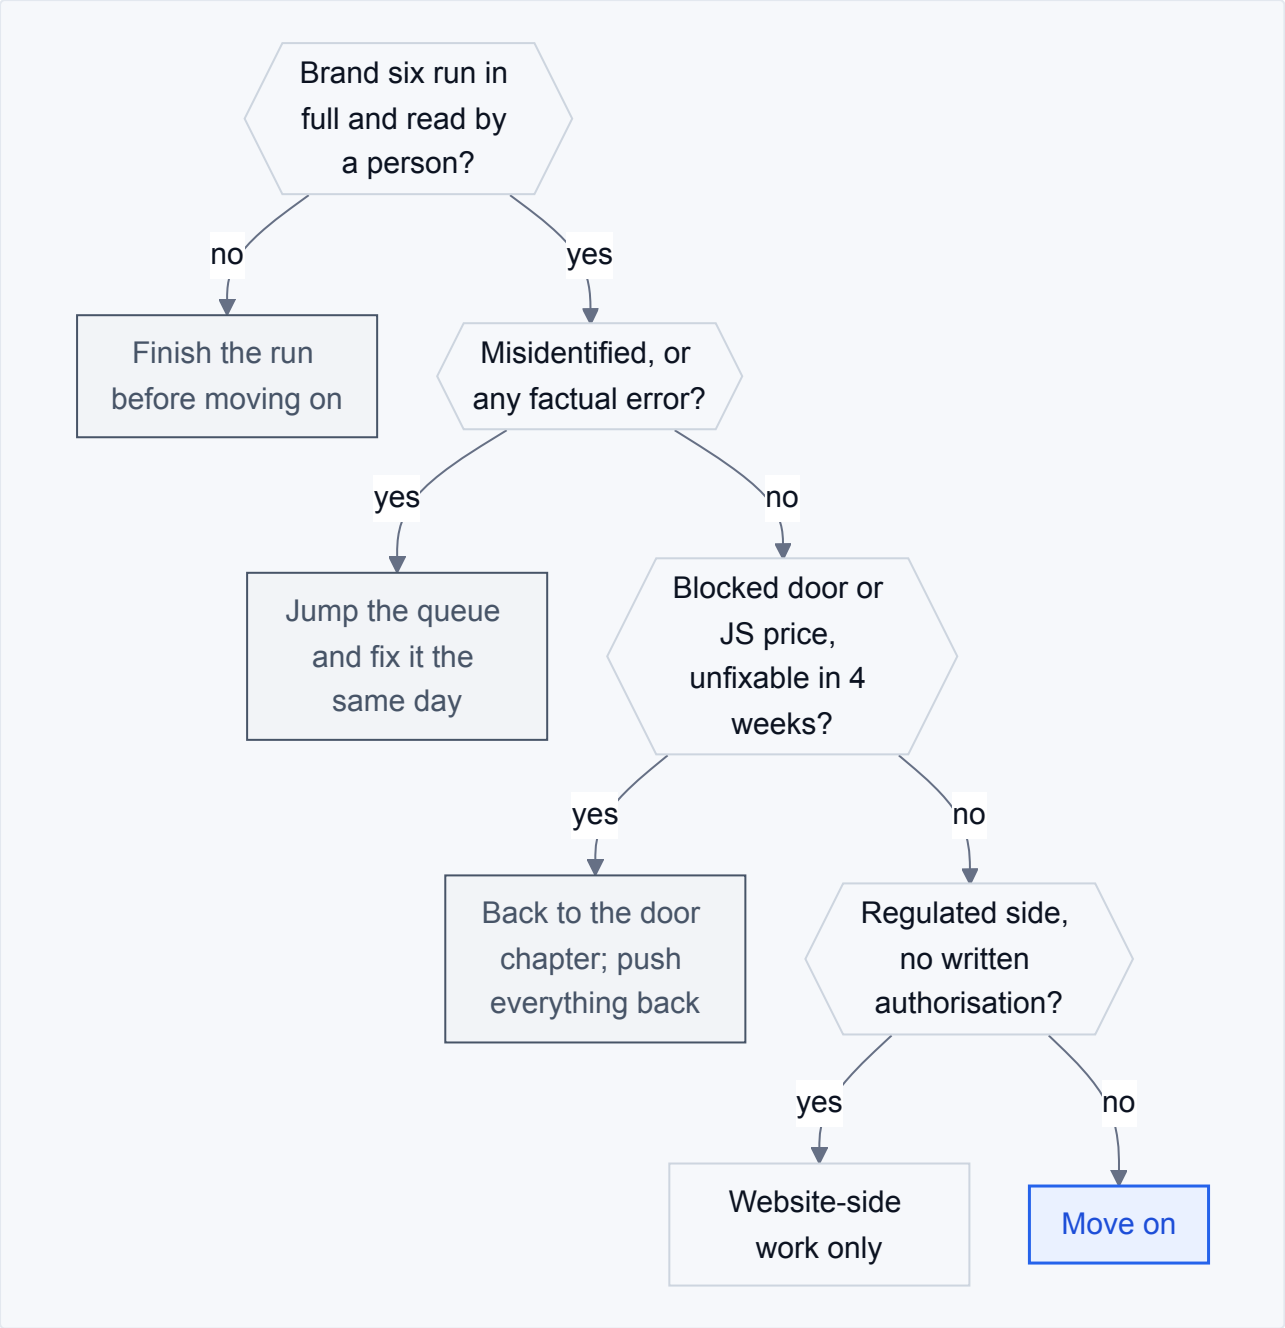


Figure: This chapter has four exits — which one are you at right now

The four exits map to:

- **Fix first, then move on**: a wrong identity, a merged entity, or any factual error → start jumping the queue the same day (3.5), fix every source you control and send the correction letters, then retest the same question two weeks later; **don't start page work on that intent until it's fixed.**
- **Push everything back**: the door layer is blocking retrieval crawlers, or the price is in JS, and it can't be fixed within four weeks → back to → General Edition 2.9 Door-layer troubleshooting and this chapter's acceptance checks; push page and off-site work back as a whole; while the door is shut, however well this chapter is done, it's multiplied by 0.

- **The whole lever fails**: on a regulated side, you can't get written advertising authorisation → not one piece of outbound material can go out, not even a correction letter; at this point only website-side work is possible (/facts + self-filled profiles), and the whole main line of six-trace alignment and third-party correction is gone.
- **Move on**: with everything else in place, go to → General Edition 4.1 The full target-picking process and the twelve pre-start questions.

Checklist for this chapter

Item	Bar to clear
Brand six	All 36 runs done and read by a person
Fact table	All 14 hard fields complete, each row with a source and verified date
/facts	Live, and the price string can be found in the HTML source
Five profiles	Every URL field points to the landing page's canonical domain
Person pages	Every practitioner has a separate page, six-trace alignment table all green
Door read	All six lines complete
Backend fields	Checked and confirmed as your own site

Common mistakes

Mistake	Consequence	Fix
Using an AI-invented price as evidence that "I can change it"	Crosses the red line on lying	Show it only as a risk illustration, the three rules in 3.5
Swapping the brand six for a different question count mid-stream	Changing the ruler = the before and after numbers can't be compared, and the whole quarter's acceptance is void	Freeze it for the whole quarter; if you really must change it, stamp a ruler version and note that before and after can't be compared directly
Folding the two Chinese questions into the six and reporting one number	Businesses with and without a Chinese page instantly stop being comparable	Print the two lines separately, each with its own denominator noted after it
Jumping the queue for seat fluctuation too	Queue-jumping becomes routine and eats the content-page hours	The "not a trigger" list in 3.5

Mistake	Consequence	Fix
Spending 3–4 hours writing a positioning statement	It never enters AI's input — written for nothing	Rewrite it as 6 answer sentences for the first intent cluster
Treating "complete the five profiles" as a visibility lever	Falls apart the moment someone asks	The fixed line: "a basic step, done in passing"

Chapter 4

Which questions to go after: pick 30 buyer questions and record how AI answers them today

Where the questions come from, which are worth winning, and saving today's answers as your starting point

4.1 The full target-picking process and the twelve pre-start questions

4.2 The question pool: where the 30 questions come from, how they are balanced, how they are signed

4.3 Coarse screen and piles, full rounds, and the brand-six baseline

4.4 The web control leg and the frozen baseline (the book's only full spec)

4.5 Buyer types, annual opportunity value and the three checks

4.6 Four states, two denominators and the abstain line

4.7 Actions by state, the ten question types and this quarter's slots

4.8 Two shortcuts and this chapter's checklist

4.1 The full target-picking process and the twelve pre-start questions

What you'll do in this section: first get the eight steps of picking targets clear: the first five come before the baseline is saved, the last three after. Then the decision-maker answers the twelve questions in person in 40 minutes; questions 2, 3 and 4 must each get a definite answer on the spot, written into the first line of the work order.

This chapter is not about picking "the questions I can win", it's about picking "**buyers who pay when you win**". The door and identity gates are ×0/1 switches; only the shelf gate adds points month by month (→ General Edition 1.3 Three gates and three paths). This chapter only manages the targets for the shelf gate: once the door and identity are fixed, this chapter decides which few cells the effort goes into. It does not touch the door, and it does not touch identity.

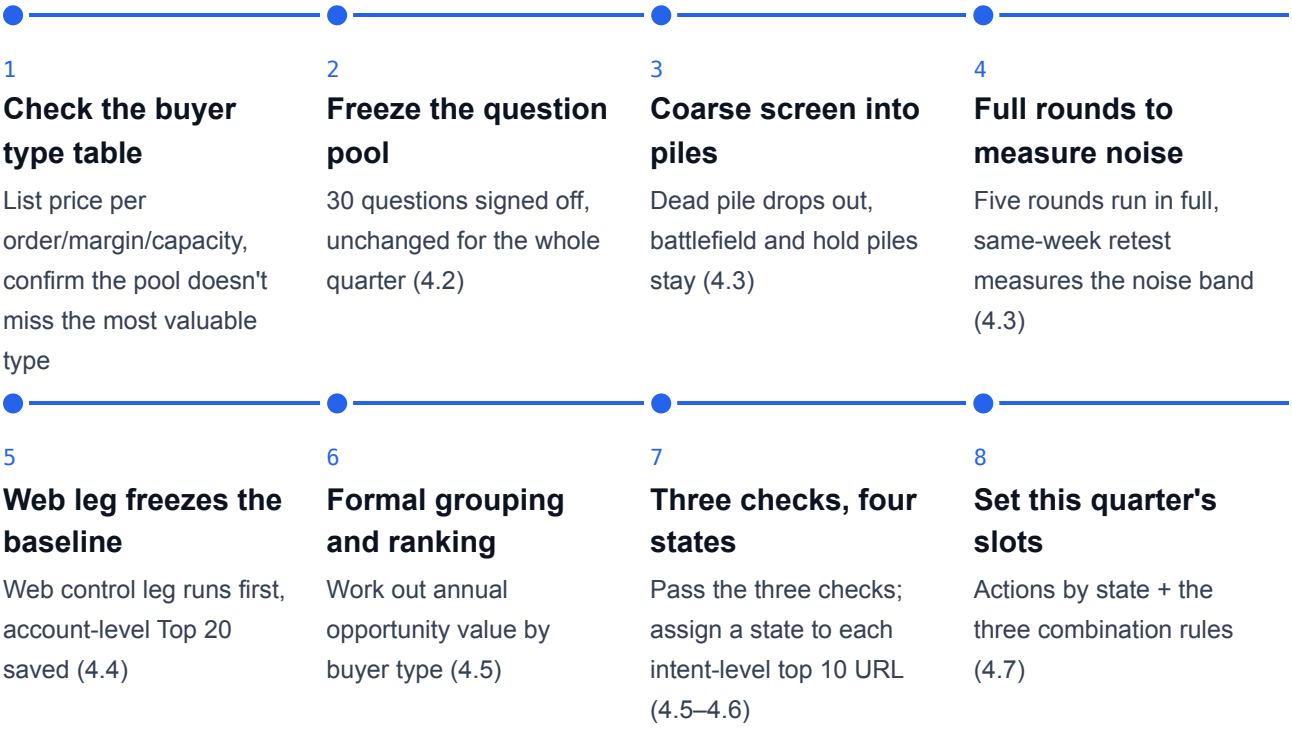


Figure: the eight steps of picking targets — the first five are done in D2–D6, before the door is fixed; the last three only after the baseline is saved

Step 1 is not the formal grouping; that happens in step 6 (4.5). It is only a check before you sign the pool: copy the price per order, gross margin and monthly capacity cap from the answer to question 5 and your own books, and use them only to confirm that the buyer type with the highest annual opportunity value has a question in the pool. Skip this check and you will freeze a batch of questions, only to find afterwards that the most valuable buyer type has no question in the pool at all — and the whole quarter cannot be judged.

Two premises the figure cannot show (the evidence is in → General Edition 1.1 How buyers ask, and what AI reads):

- The buyer's real path is: see the name → search separately → visit the website → check reviews → contact. Almost no one clicks the links in an AI summary. So only pick questions from the "who to buy from" stage; do not pick educational questions.
- Most of what AI cites is pages other people write about you; paid placements and advertorials are close to zero. So what the shelf snapshot needs to record is not "am I present" but which containers currently hold the seats for this question.

The twelve pre-start questions

Who: the decision-maker asks in person · **How long**: 40 minutes · **Output**: a one-page answer sheet · **How to check**: questions 2, 3 and 4 each have a definite, fixed answer written down on paper.

Tag every fact the other party (or you) states as "to be verified" first; the baseline (4.3–4.4) is what settles it.

#	What to ask	Why ask it	What to do once answered
1	Legal name, main market, licence/registration number	Seats and identity checks all go by the legal name	If you can't get the registration number, leave a gap in the fact table and fill it in before moving on
2	Which side you're on, which switches are on	Decides which set of side rules you follow; one entity can only be on one side	Write "side + switches" down on the spot as the first line of the work order; a business with mixed activities is built site-wide to the stricter side (switch D (spans two sides))
3	How far you can go with the price (asked by side; see below the table)	The question is "can you give it", not "are you willing to publish it"	Can give it → the price page is top priority; can't → downgrade the whole job to audit + monitoring only
4	Where switch A (agency liability) is on: will you sign the written publishing authorisation + the materials sign-off sheet, and who is the signatory	The publishing party shares liability with you	Won't sign → do not do it. This is the only hard refusal with no downgrade path
5	How many times in a lifetime a buyer buys, and how much per order	Buy once or twice in a lifetime, high price per order, buyer can't tell good from bad — most worth doing	High frequency, low price per order → downgrade to hold or do not do
6	Who can edit the website, how many days to change one page	If changing one page takes more than four weeks, the whole first month's schedule is void	More than four weeks and unwilling to change it → downgrade to hold

#	What to ask	Why ask it	What to do once answered
7	Who manages the five business profiles and the two webmaster tools, and can they give read-only access	Needed for the door gate; connect the free first-party data before paying for probes	Can't give read-only access → the first two items can't be verified at acceptance; list it separately as a risk
8	Is there a namesake organisation or namesake person	If AI identifies the wrong business, all the page work upstream returns zero or even less	Yes → the first batch becomes facts pages + profiles, not content pages
9	Are you buying optimisation services from someone else; are you using a platform priced per enquiry, per lead or by sale commission	If switch F (referral commissions banned) is on, this is a red line	Switch F is on and in use → write it into the compliance memo; stop using it first, then start work
10	Who you think your competitors are	Two reviews both found: the competitors you name yourself never show up in AI's answers	Tag all of it "to be verified"; the baseline decides
11	Are there Chinese-speaking buyers, are there buyers in other minority languages	A Chinese landing page is the only lever that goes "build one page → straight into the named set"	Chinese buyers → produce Chinese questions per 4.2; minority languages, see below the table
12	Whether orders mostly come in by form, chat or phone	Decides the acceptance tier (4.5)	Phone → observable tier + a mandatory front-desk field "how did you find us", never dropped from the pool for this alone

Three hard gates: questions 2, 3 and 4

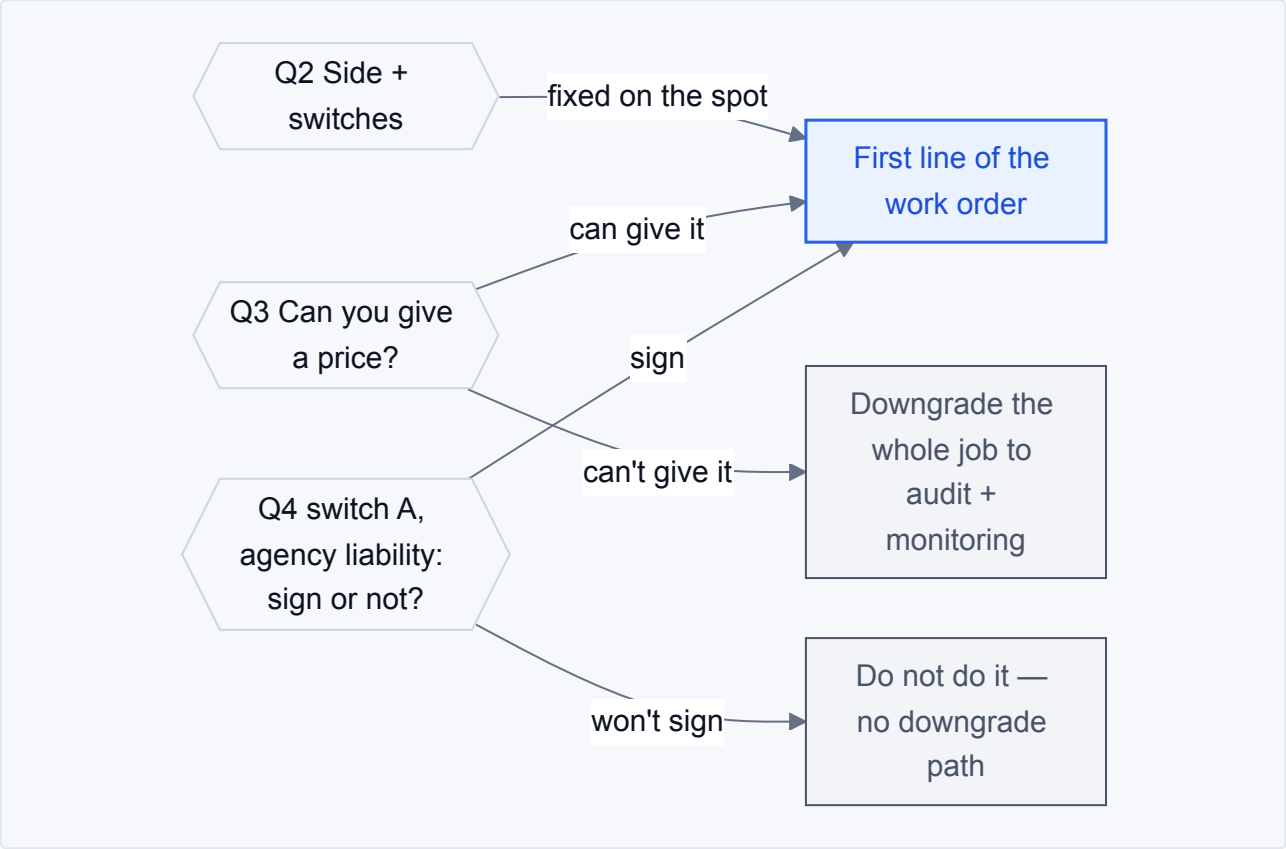


Figure: each of the three hard gates puts a definite answer into the first line of the work order; if question 3 can't be met, the job is only downgraded; if question 4 is not signed, the job is not done

The full definitions of the three sides and the six switches are in → General Edition 0.2 Decide which side you are on first. Question 3 asks something different by side: on the strictly regulated side it asks "can you give a fixed final price in tiers" — this side cannot even write a range, the rule is in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, and if you will only write "from", that counts as can't give it; on the lightly regulated side it asks "can you give a range + billing method"; the unregulated side has no such gate.

For question 11's minority languages, run a real search first: if the phrasing points to another country's home market, that is a wrong target, not inefficiency — add Chinese questions instead of forcing that language (example in 4.2). Questions 10 and 12 are not gates: competitors are judged once the baseline is out; the order channel only decides how the acceptance tier is split, not whether a question enters the pool.

4.2 The question pool: where the 30 questions come from, how they are balanced, how they are signed

What you'll do in this section: first spend half a day reading two sets of free first-party data and the buyer's own words from the last 15 orders, then use only the four classes of source with evidence to make up 30 questions, have someone who did not draft them count the ratios against the fixed targets, and finally have the decision-maker read every question aloud in person, make changes on the spot, and sign on the spot. Once these 30 questions are signed, they do not change for the whole quarter.

The 30 questions are the ruler for the whole quarter: picking targets draws only from this pool, and the retest and the settlement use the same set of questions. **The ceiling on the question pool's quality is the ceiling on the project's results** — so source discipline comes before probing.

Step zero: read the free data before paying for probes

Who: the person who edits the site gets access, the person reading the data reads it · **How long**: half a day · **Output**: two screenshots + a one-page summary · **How to check**: the decision-maker has read it; only then is the probe budget released.

Connecting the two back-end reports was already done in → General Edition 2.5 Gate 3, read-only: connect the two free AI reports; here we only cover how to read them and what to watch out for in each.

What you get	The one catch
GSC: generative-AI impressions by page/country/date	Impressions only; already included in total impressions, don't add them again; flag blank values exported as 0
Bing: cited pages, grounding terms, citation share	Grounding terms count only as a clue, not standalone demand evidence
Front desk/CRM: buyers' own words before purchase, last 15 sales	If there's no record, make it a mandatory field on the spot

Figure: step zero — what you get from the two free reports and the front-desk records, and the one catch each has

This step does not move rankings by itself. It decides whether all the effort that follows lands on real questions, not questions you imagined. Bing's grounding terms feed straight into the question pool as a class C source below; the third row is a class A source, and the mandatory fields are "how did you find us" and "what did you most want to ask".

Only sources with evidence enter the frozen pool

Who: the person setting the questions drafts it · **How long**: one day · **Output**: a question-pool table, each row with a source label and the original evidence (a link or screenshot) · **How to check**: a second

person checks every row, then the pool is signed in person.



Figure: each of the four source classes needs one piece of evidence to enter the frozen pool; everything else goes only into the reserve pool

Class C means real GSC queries, Google autocomplete, related questions, and Bing grounding terms. Class D only takes questions that are **already listed** on a competitor's comparison page, a directory site or a list page. The four classes are mutually exclusive, each entry is tied to one checkable piece of evidence, and any row without evidence is sent back. The reserve pool can be used as a reference for topics when writing pages, but **it does not enter the frozen pool and is not a ruler**.

If you can't make up the number: go back and find more class A words from real buyers — do not fill the gap with AI. An AI-generated sentence that gets into the frozen pool is the single most fatal, irreversible mistake of the whole quarter: it contaminates the ruler itself, not just one page's content. If you can't reach 10 class A entries, downgrade at this step to audit + monitoring: a business that can't produce even 10 real buyer questions in their own words can't support a ruler.

The ratios are fixed

A baseline conclusion that breaks the ratios is void; redo the question pool.

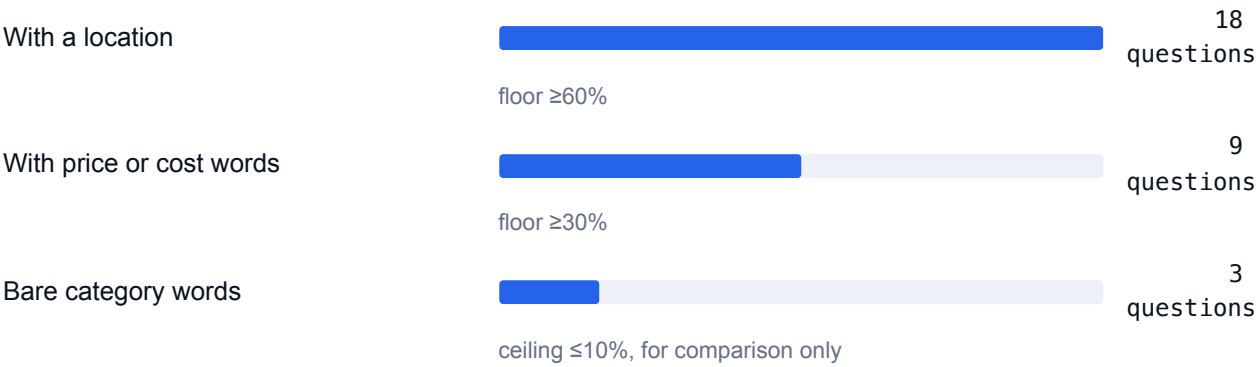


Figure: the ratio lines for the 30 questions — location and price are floors, bare words are a ceiling

Each of the three lines has one reason. The location floor: bare words can produce a fake zero — the same query with or without a location can go from "zero businesses" to "every business restored", and that "zero" in the report is one you made yourself. The price floor: price/cost-type phrasing has a high rate of triggering AI Overviews (4.7). The bare-word ceiling: a national head term with no location can't move within 90 days. Evidence in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests.

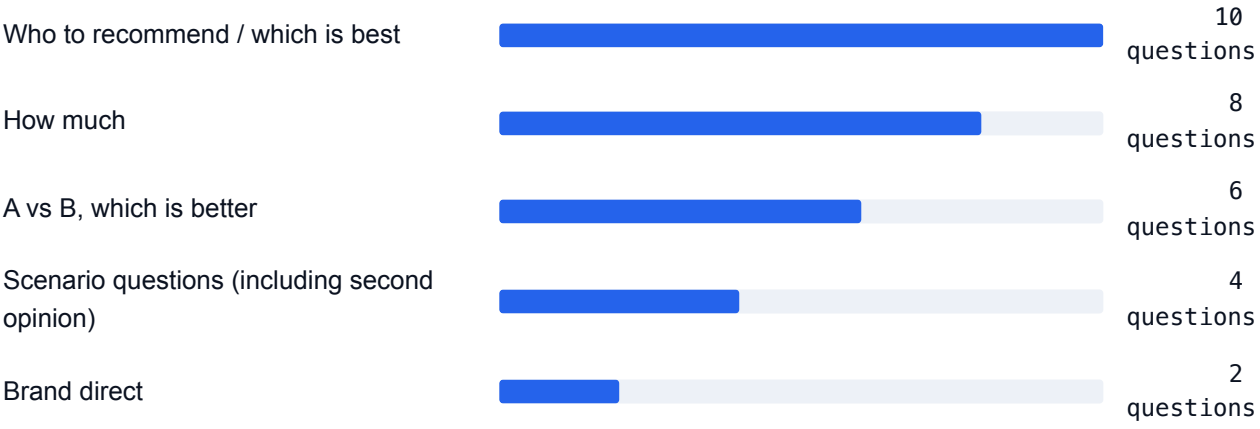


Figure: the same 30 questions split again by question shape — parallel to the ratio lines, not a separate batch

Where there are Chinese-speaking buyers, Chinese-language questions are the main line, not an add-on (already decided at question 11 in 4.1). Write the count into the work order; Chinese questions always carry a location, price-type ones count into the price-word row, **and count at the same time into the location and price numerators above — they don't get their own total or a separate table**. When the ratios can't be filled, cut from the 3 bare category words first; **never cut a Chinese question**. The bar for the Chinese cell is not getting onto a list; it is whether you have a Chinese landing page (evidence in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests).

> **Example** (legal) A divorce lawyer fee question phrased in Indonesian: all 9 citation sources were Indonesian local law firms, the answer quoted a price in Indonesian rupiah, and the engine itself added a line saying the result does not concern Singapore. The buyer's intent was misaligned, and even doing the question wouldn't get you into the local named set — that's a wrong target, not inefficiency. Add Chinese questions to the pool instead.

Second-person check and in-person sign-off

Who: someone who did not draft it checks it + the decision-maker signs · **How long**: 30 minutes to sign · **Output**: a signed PDF of the question pool · **How to check**: there is a signature on paper, and nothing changes for the rest of the quarter.

Any person who did not draft it counts through the ratios item by item and signs "ratios verified · date · name" at the end of the question pool document. Before signing the pool, check the buyer type table against 4.1 step 1 — if the buyer type with the highest annual opportunity value has no question in the pool, do not sign. The decision-maker reads every question aloud in person, makes changes on the spot, and signs on the spot; for the sign-off wording see → General Edition B.2 Approved wording and scripts. **Changing a question is changing the ruler, and once the ruler changes, the numbers before and after can't be compared** — swapping questions partway through is the easiest way to fail the whole quarter, because afterwards no one can tell whether the number rose or the ruler changed.

4.3 Coarse screen and piles, full rounds, and the brand-six baseline

What you'll do in this section: first run a coarse screen of 30 questions × 1 round × 2 engines = 60 runs to split the question pool into a dead pile, a battlefield pile and a hold pile; run only the battlefield and hold piles through five full rounds, then retest in the same week to measure the noise band; in the same week, also run the brand-six questions in full — 36 runs — as the baseline for factual errors. When recording answers, count "listed as a source" and "written into the answer" separately.

Coarse screen into piles

Who: the person running probes runs it, the person setting the questions reviews it · **How long**: half a day · **Scale**: 30 questions × 1 round × 2 engines = 60 runs · **Output**: a three-pile list.

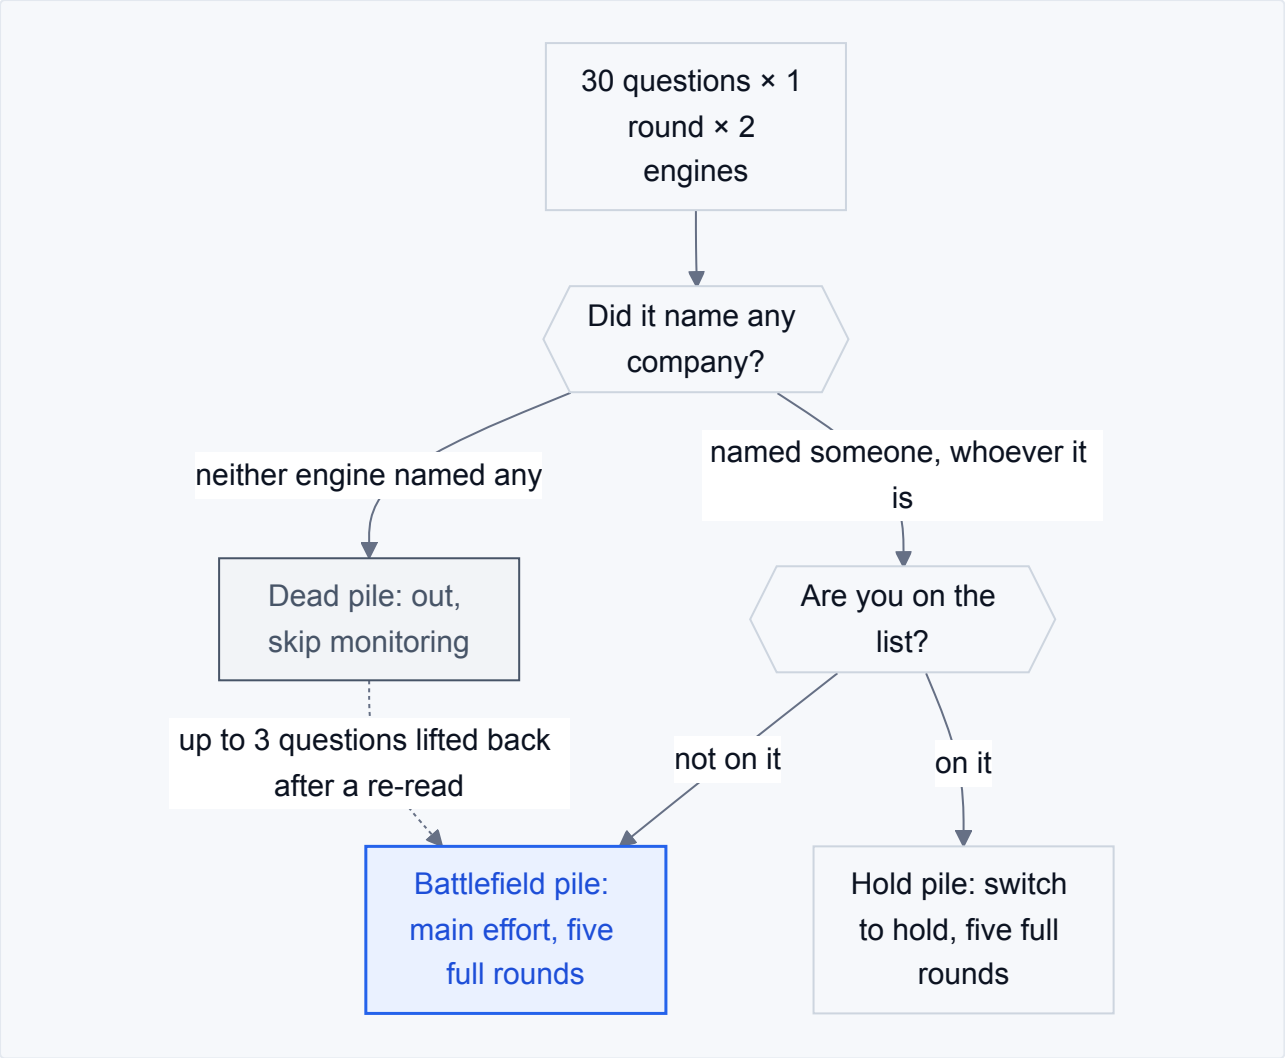


Figure: the coarse screen only checks whether AI named any company at all, splitting the pool into three piles — only two are worth paying to run in full

The dead pile means neither engine names any company at all — it only gives educational content, or just tells you to consult a professional. Once out, spend no more money on it, and it does not enter the monitoring layer.

The dead-pile criterion is an operational convention, not a measured finding: 70% confidence. Because a wrongful kill is irreversible, two safeguards are added: a second person must re-read the original answer text before the dead-pile list is finalised; the dead-pile list is reviewed together with the shelf snapshot, and up to 3 questions may be designated to lift back into the battlefield pile for a full round. **The approved external wording for the dead pile is fixed: n = 2, write only "no business named in the coarse screen", never "zero visibility"**. For how much the two-stage screen saves compared with running everything in full from the start, and what is lost, see → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests.

Full rounds and the noise band

Who: the person running probes runs it, the person reading the data compiles it · **How long**: 2–3 days · **Scale**: (battlefield pile + hold pile question count) × 5 rounds × 2 engines · **Output**: raw answers + a

summary table + `noise-band.md` · **How to check**: 2 hours of deep reading by a person + checking every item against the judgement checklist.

1
Validate the ruler first
Run a question you know gets cited; if the citation doesn't show up, stop

2
Lock region and language
Don't give AI your own URL; run Chinese and minority languages separately

3
Run the full rounds
Battlefield pile + hold pile, five rounds each × two engines

4
Same-week retest
Same questions, same engines, same number of passes, run 3 more times in the same week

5
Web leg + freeze
Full spec in 4.4, not repeated here

6
Keep two tables separate
Don't merge brand questions with category questions; the brand table stands alone

7
Sign the control agreement
Same-site pages not to touch, three competitors, the primary number to watch, when to stop

8
Save after the deep read
2 hours of manual deep reading can't be skipped; save it the same day, permanently

Figure: the eight steps of a full round — the baseline is all saved together in this one week; if step 1 misses the known citation, the whole round stops

When the ruler misses the known citation: do not issue a report, do not announce any "zero". Minority languages are run separately so they don't get mixed into the main language's basis.

How to measure and use the noise band: this month's noise band = the range (max − min) of the seat count across those 3 same-week runs, with a floor fixed at ±1 seat, written into `noise-band.md` and redone every month at the monthly retest (→ General Edition 7.2 The noise band, the page-level signal and the ten monthly steps). A monthly change in seat count may only be written as "increased" when it is larger than the noise band — report "up 1 seat" without measuring the noise band, and you are reporting noise, not results.

Once this round is done, the engines, the number of passes and the question pool are all frozen for good — every comparison for the next 90 days is against this set of numbers. Write the two legs' official names only as given in → General Edition 1.2 Two legs: ChatGPT looks for the source, AI Mode for second-hand summaries; never write "real user visibility".

What to look for in the deep read: whether a community mention is a namesake mix-up, check any negative claims once, and whether the industry wording is right. After saving, keep only the 3 most striking screenshots to make your case with; anything that names a third party is treated as internal material. Don't notarise the whole batch of screenshots.

Two things to count separately in every record

Listed as a source	Written into the answer
The page appears in the source area below the answer	AI writes your name or number into the body text
Ruler: on the list and cited X/20	Ruler: seat count
Not listed: get the page onto the list first	Listed but not written in: adding pages won't help, you need to change the sentence

Figure: one answer has two different kinds of "present" — the fix is completely different depending on which one is missing, so count them separately

These two numbers are not merged into one "hit" (basis in → General Edition 1.3 Three gates and three paths). Also record the naming position: getting onto a list AI already cites and ranking first on it raises visibility and moves your position in the answer earlier, but this set of numbers is only used to order the work — never to promise any result. Write numbers and their qualifiers in the unified form given in → General Edition D.1 Number discipline: how to label numbers, and what stays internal. How the four states and the intent-level denominator use this data is in 4.6. When recording each answer, also log these fields:

Field	How to record it
Who was named, whether you're in it, which position	Record the position
Total names in this answer	The total count of business names named in the answer body — record the actual count, don't estimate
What kind of site the citation comes from	own site / directory / list / review site / maps / forum / government / media; count listicles as their own category — share and sources in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests
Review field	Whether the answer body writes review counts or star ratings word for word; record recommendation-type and non-recommendation-type on two separate lines, never merged into one total frequency; once your own baseline is run, use your own numbers — reference numbers in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests
Forum source share	Only put it on the work list when Reddit's share of this month's cited URLs is >2%; at ≤2%, record only, don't invest; don't write a blanket "don't do Reddit" — reasons in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests
Sampling status	Keep valid/failed separate; remove failed samples from the denominator and list them separately

Field	How to record it
Citation count	Keep URL-level and answer-level separate; during the baseline period most will be 0 or 1, URL-level detail can wait until the first retest
Tendency	positive/neutral/negative/excluded, **a negative or excluded result jumps the queue and triggers an alert the same day**

How to write the numbers: write "asked K times, named X times" — never a single-question small-number form like "hit 2 out of 3".

The brand-six baseline

Run the brand-six questions in full, 36 times, in the same week: 6 questions × 3 rounds × 2 engines, defined in → General Edition 3.5 The brand six questions, and what to do when you find an error. The baseline takes only this one full round, never a quick run — two runs with different composition are not the same ruler. It is the only source of the count "factual errors X": the 25-minute read in → General Edition 3.1 Why this comes before writing pages · The two-hour checklist reads this same raw set of answers, no separate quick check.

4.4 The web control leg and the frozen baseline (the book's only full spec)

What you'll do in this section: run the first web control leg before you freeze the account-level Top 20, and force source domains found only on the web version into the candidate pool. After freezing, ask the same fixed 5 questions by hand every month, once each, and compute only the two overlap rates, to calibrate the two API legs. Only once all three parts of the baseline are saved do you unlock → General Edition 2.6 Fixing the door: nosnippet, the four-step robots merge, the WAF allowlist.

Why this leg is required

Both API legs measure the model API. The citation distribution on the API is not the same set of sources a person sees on the web version, and this is a systematic skew, not noise:

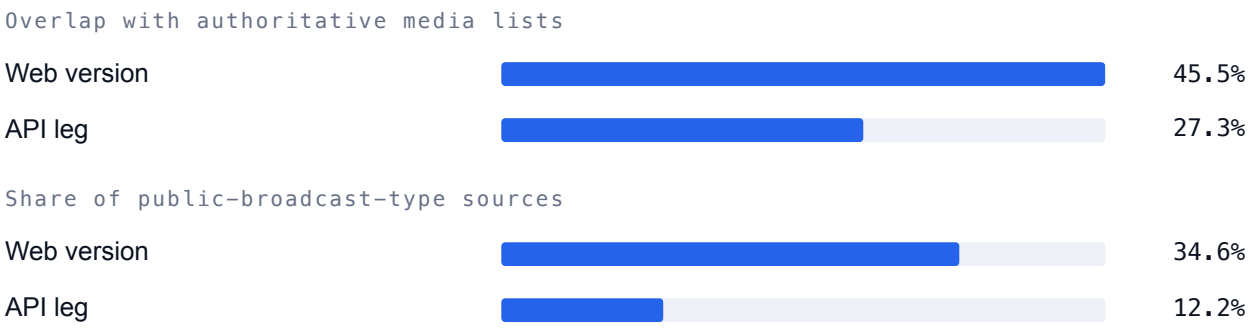


Figure: in both comparisons, the web version is far higher than the API — the two sides cite different sets of sources

A single domain can shift as far as position 1 on the web version and position 61 on the API. **The skew is just as large whether you run 5 rounds or 50** (sample and sources in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests). That last sentence is the whole reason this leg exists: adding more rounds does not fix a systematic skew — only switching the path does. Skip it, and the whole quarter's off-site effort aims at the API leg's account-level Top 20 while buyers are actually seeing a different set of pages. When the effort misses, nothing warns you: the on-the-list page count still climbs, and the report still closes out clean.

So two things are fixed: the ruler's official name is "**Model-API visibility baseline**", printed in the footer of every report; visibility for AI Overviews and AI Mode always uses the Search Console generative AI report's impression numbers instead (already connected in 4.2's step zero) — never pass off the Gemini model leg's numbers as a substitute. The Gemini API, AI Overviews and AI Mode are three different products.

How to run it (the spec is fixed; change nothing)

Item	Fixed
Questions	Take a fixed 5 from the frozen 30: ≥2 "who's best / near me" type, ≥1 price type, the rest from the battlefield pile. Once these 5 are picked, they don't change for the whole quarter
Path	ChatGPT web version, typed by hand in a browser. No API, no proxy interface of any kind
Identity	Logged out, or using Temporary Chat; personalisation and memory turned off
Region	An IP local to the target market (be there in person, or use a local exit node; e.g. an SG IP for the Singapore market)
Passes	1 pass per question. Don't rerun, don't pick the best answer, don't reword
Recording	A full-screen screenshot per question (including the source area below the answer — scroll and take more than one if needed), copied by hand into `web_leg_<month>.csv`

Item	Fixed
Fields	question / date / businesses named in the answer (each one, in order of appearance) / source URLs listed below the answer (each one) / screenshot filename
Who	The person running probes runs it; another person checks every row of the screenshots against the csv
How long	40 minutes within one day each month; not part of the monthly retest process table — schedule it as its own day at the start of the month

It produces two numbers, only these two, with no further processing:

Number	How it's calculated
Named-set overlap	(the set of businesses named by the web leg ∩ the API leg) ÷ their union, calculated only on these 5 questions
Source-domain overlap	(the set of source domains listed by the web leg ∩ the API leg) ÷ their union, calculated only on these 5 questions

Neither number is mixed in with the full pool of 30 questions; the denominator is the union, not the web leg alone. They are never used as acceptance evidence: they don't go into the seat count, don't go into the on-the-list page count, are never added to any other ruler, and never produce a standalone "web visibility" score.

The baseline run: web leg first, then freeze the account-level Top 20

The order is fixed; get it backwards and the whole quarter cannot be recovered.

1
Save the API leg
The full-round baseline is done, the raw domains_cited file archived

2
Run the web leg the same day
Fixed 5 questions × 1 pass, asked by hand on the web version

3
Fill in the candidate pool
Domains unique to the web leg are forced in even if the API leg cited them 0 times

4
Take the top 20 and freeze
Merge the candidate pool and take the top 20 by URL-level citation count

5
Mark and keep permanently
web_only=Y listed separately; monthly review continues as normal

Figure: the five steps before freezing — get them backwards and the whole quarter cannot be recovered; only after all five are done may you fix the door

Use the **raw** `domains_cited` from the full-round baseline to take the top 20 — not a version with your own site and generic domains already stripped out. Once taken, a second person spot-reads 3 pages. The frozen account-level Top 20 is the denominator for "on the list X/20", and also the whole quarter's off-site target — if the target is taken only from the API leg, it is only the API's shelf, not the buyer's shelf. Freeze first and run the web leg afterwards, and the domains it turns up can no longer enter the denominator — the whole quarter's target is then wrong.

How to read the two overlap rates

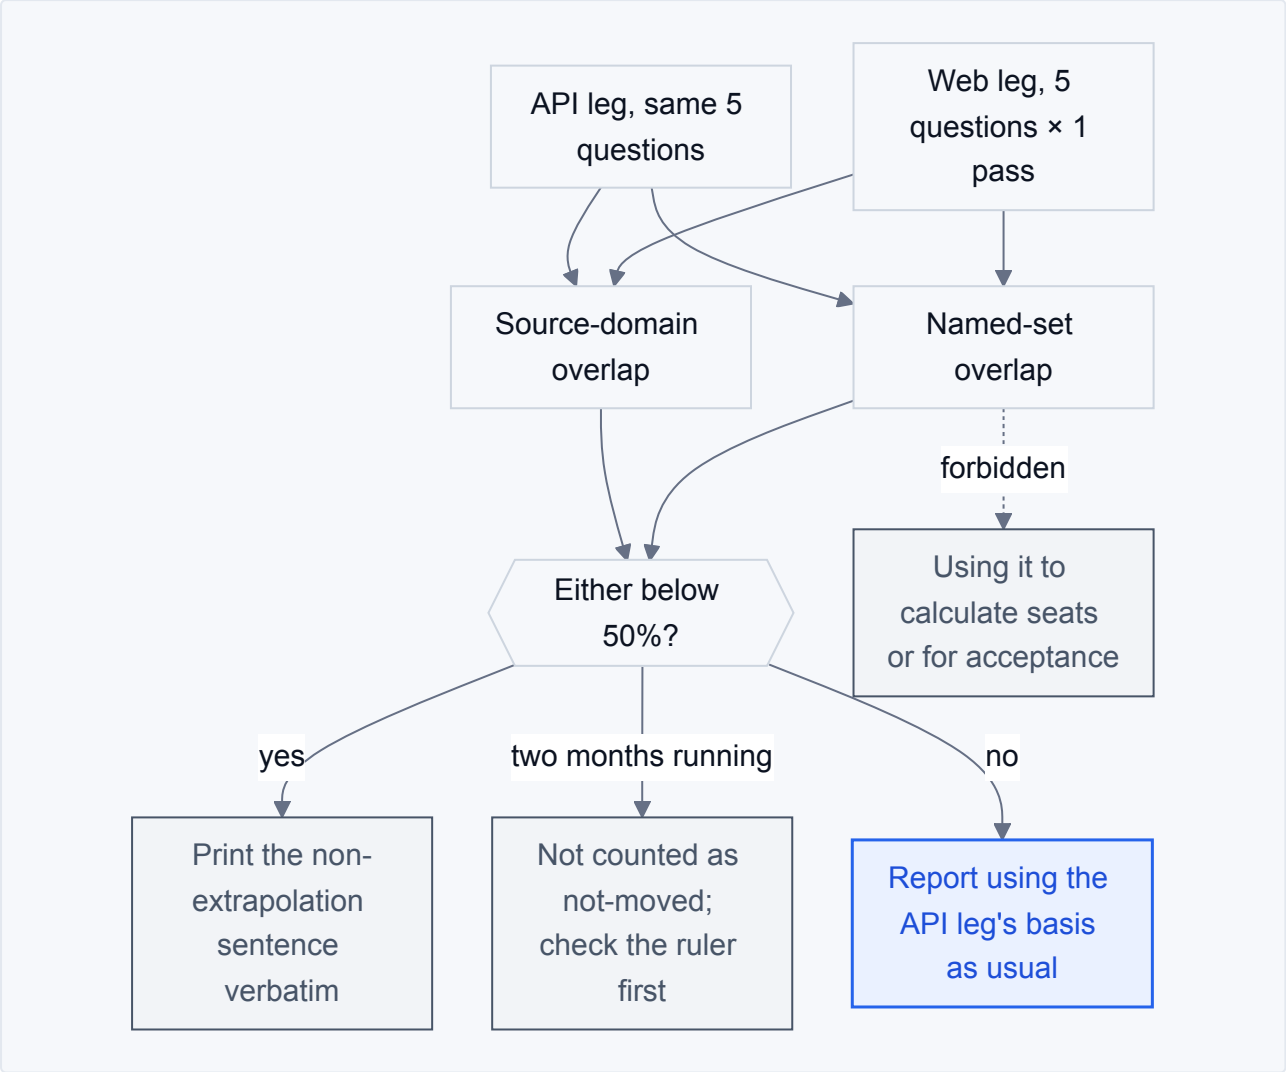


Figure: the two overlap rates only calibrate the API leg — below 50% only changes how the report is written, never any acceptance decision

The non-extrapolation sentence does not change a single word: "This month the two paths diverge (named overlap X%, source overlap Y%). The seat changes below were measured on the model API and are not extrapolated to what real users see on the web." **Never use it the other way round**: don't write "already visible" just because the web leg named you (n = 5, single pass), and don't use these 5 questions to calculate seats either.

How to check, and common mistakes

Miss any one of these and it doesn't count as done: `web_leg_<month>.csv` has 5 rows, every row has a screenshot filename, and the screenshot opens; the screenshot shows the "logged out / Temporary Chat" interface feature and a visible date; the numerator and denominator sets for the two overlap rates are listed item by item, not just given as one percentage; in a month below 50%, the non-extrapolation sentence is present, word for word; in the baseline run, `sources-top20.csv` has a `web_only` column with at least one entry forced in — if there are none, there are only two possibilities: it wasn't run, or the

two legs are completely identical (the latter must be stated explicitly, with the full domain list from both sides attached).

Common mistakes: running the web version through an API or a script instead (that's still the API leg — a wasted run); running it while logged in; running one question three times and picking the best-looking answer; swapping the 5 questions every month; reporting the overlap rate externally as a results number; freezing the account-level Top 20 and only running the web leg afterwards.

The account-level Top 20, the noise band, and the 36 brand-six runs — all three saved = the baseline saved. Next, go back to → General Edition 2.6 Fixing the door: nosnippet, the four-step robots merge, the WAF allowlist.

4.5 Buyer types, annual opportunity value and the three checks

What you'll do in this section: once the baseline is saved, formally group the questions by buyer type (not by question shape), rank them by annual opportunity value, then run every candidate through the three checks. When you're done, you'll have a candidate list grouped by buyer type, ranked by annual opportunity value, with every entry having passed the three checks.

Group by buyer type

Who: the decision-maker + the person setting the questions · **How long**: half a day · **Inputs**: the price list, the last 20 sale records, capacity · **Output**: a buyer type table · **How to check**: the price per order and capacity for each type has been confirmed in person by the decision-maker.

The same question is handled differently in different industries: in AI Mode, professional services win through educational content, consumer services win through review volume (75% confidence, sample doesn't include the local market, evidence in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests). For professional-services buyer piles, add long-form content and directory/list placements. Note But on the strict side the review-volume point must not turn into "go and ask for reviews": there, for consumer-type buyers, put the main effort into price pages, process facts and complete business profiles; reviews are only observed passively. The same playbook can't be used for every side.

Rank by annual opportunity value

```
annual opportunity value = price per order × gross margin × monthly capacity cap ×
```

All four numbers are copied straight from your own price list and sale records — anyone can recompute them. This is the only ranking axis; there is no scoring.



Figure: example — a clinic's own books, for ranking only, not a revenue forecast; when values are close, pick the one with less competition

Four rules: fill in the numbers yourself, confirm them in person, and put them somewhere visible in the plan; **the monthly capacity cap is a hard cap** — if you can't handle that many orders in a month, don't pick a question for that many orders; annual opportunity value is for ranking only, never described externally as a revenue forecast, and it must never be multiplied, divided, or shown side by side with seat counts or enquiry counts; when values are close, pick the one with less competition.

The reason for not ranking by a four-factor score: three of the factors measure the same underlying variable (which is equivalent to cubing it), and the fourth is really a proxy for search volume — it would let "ranking by search volume", the wrong method, back in through the side door, systematically pushing low-frequency, high-price work to the bottom. Dividing by cost on top of that adds only noise, because cost is roughly constant.

The three checks



Figure: the three checks, each with its own action when a question doesn't pass; check 2 must never be written off directly as "do not do"

Check 1's definition is fixed: naming means naming any company, not specifically you. A common misreading is "keep it only if you yourself are named at least 3 times out of 5" — read that way, every business that genuinely needs the work fails check 1, which turns away every target; and at n = 5, the difference between 3/5 and 2/5 is a coin flip.

Check 2 reads this intent's own intent-level top 10 and judges each URL's state one by one; Note it does not read the account-level Top 20. The four states and "reachable enough" are in 4.6; actions by state and "retry from another angle" are in 4.7.

Check 3 answers one question: the people this question brings in — will you take them, and will you make money? Questions like "the cheapest X" get dropped from the pool on purpose, with a record kept — write it into the plan so people can see what was turned away.

Whether you can give a price was already asked once at question 3 in 4.1, so it's not asked question by question here. Whether attribution is possible is also not judged question by question in check 3: it's a site-level capability, not a property of any one question — asking it question by question always gets the same answer, and judging by channel for real would systematically cut every high-price, phone-order question. So it is handled with two acceptance tiers instead, set on the day work starts; nothing is dropped from the pool here:

Attributable tier	Observable tier
Form/chat pre-filled, CRM tags complete	High price per order, phone orders, 3–9 month sales cycle
Accepted per enquiry	Accepted by seats + position + on-the-list page count, plus manual attribution via a mandatory front-desk field

Figure: the acceptance tier is set by the order channel; phone orders are never dropped from the pool for this alone

4.6 Four states, two denominators and the abstain line

What you'll do in this section: judge the state of every URL in the intent-level top 10, keep clear what the account-level Top 20 and the intent-level top 10 each govern, and count how many are "reachable enough" — fewer than 4, and this intent does not get worked on this quarter.

This section is the book's only authoritative definition of the four states and the abstain line; later chapters only refer back to it and do not repeat it.

How to assign a URL to one of the four states

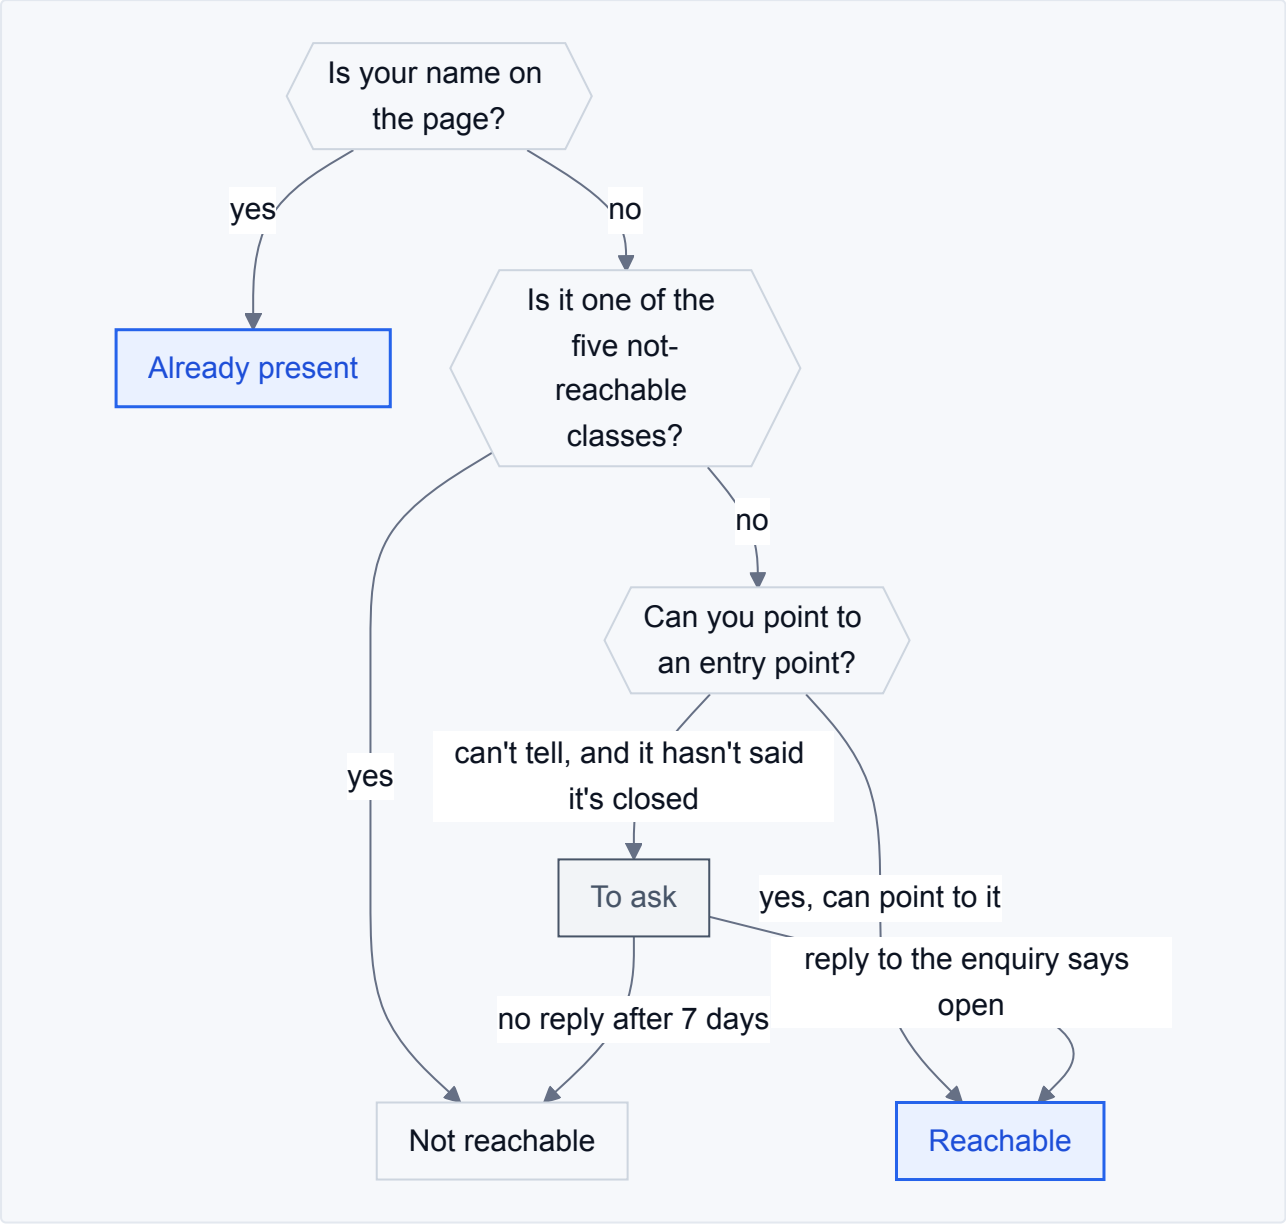


Figure: how a URL is assigned one of the four states; "to ask" is only a transitional state and must settle into "reachable" or "not reachable" before the baseline

"Reachable" means this page accepts listings, submissions, publication, claims, nominations or corrections, free or paid; you must be able to point to the entry point: one of an entry-URL screenshot, an email receipt, or the place on the page where it says so explicitly. "To ask" is only allowed to exist at the coarse-screen stage; it must be cleared to zero by the baseline stage — send one enquiry per item. If "to ask" is still in the baseline table, that baseline isn't finished and must not be used for picking targets.

The five not-reachable classes: government or statutory registers and notice pages; Wikipedia article text; pages under a competitor's own domain; platform-produced content pages (edited by the platform itself, explicitly not accepting outside submissions); closed-door ratings, or where a reply explicitly says it's not open, or 7 full days pass with no reply. Ruling Wikipedia's article text not reachable doesn't close

that lead: its structured anchor, Wikidata, is reachable — whether and how to build it is in → General Edition 3.4 Third-party credential tiers, Wikidata and the five business profiles.

Every judgement must keep four columns: state / judged on / basis / judged by. **A row with an empty basis column is always treated as "to ask"** — you can't tell whether it was "judged too generously at the time" or "that page really did redesign and close submissions", and the next action for these two is exactly opposite. **Judge by URL, not by domain**: under the same domain, some pages are open and some are not. How the four states map to the glossary is in → General Edition D.2 Glossary; how they map to the judgement axes is in → General Edition 8.3 The ten axes (1): page types, containers, anchors, regulation, question shapes — "who controls the container" there judges a different thing; don't mix the two up.

Two denominators, each with its own job

Account-level Top 20	Intent-level top 10
All intents merged, taken as the top 20 by URL-level citation count	Each intent takes its own top 10, from its own answers
Frozen, unchanged for the whole quarter; it judges no single intent	Grouped by intent_id; only exists after the baseline
Used only for the one number "on the list X/20"	Read for the abstain line, check 2, and the reachability side of not-moved triage

Figure: the two denominators each have their own job — never mix them

Why the two layers must be split: one shared account-level list always gives the same count, so using it to judge "do or don't do this intent" gives only two outcomes — every intent judged dead together, or none judged dead at all; the safety valve is effectively not installed. The data already exists; all that changes is how the summary script groups it — the raw citation records already carry "which question's answer this came from", it was just being merged away before. Intent-level takes the top 10, not the top 20, because a single intent's source pool is naturally shallow — URLs past rank 10 are mostly cited only once, and putting them in the denominator would only dilute the judgement.

The abstain line

The criterion counts the number "reachable enough": a URL only counts as reachable enough if it meets all four conditions **at the same time**.

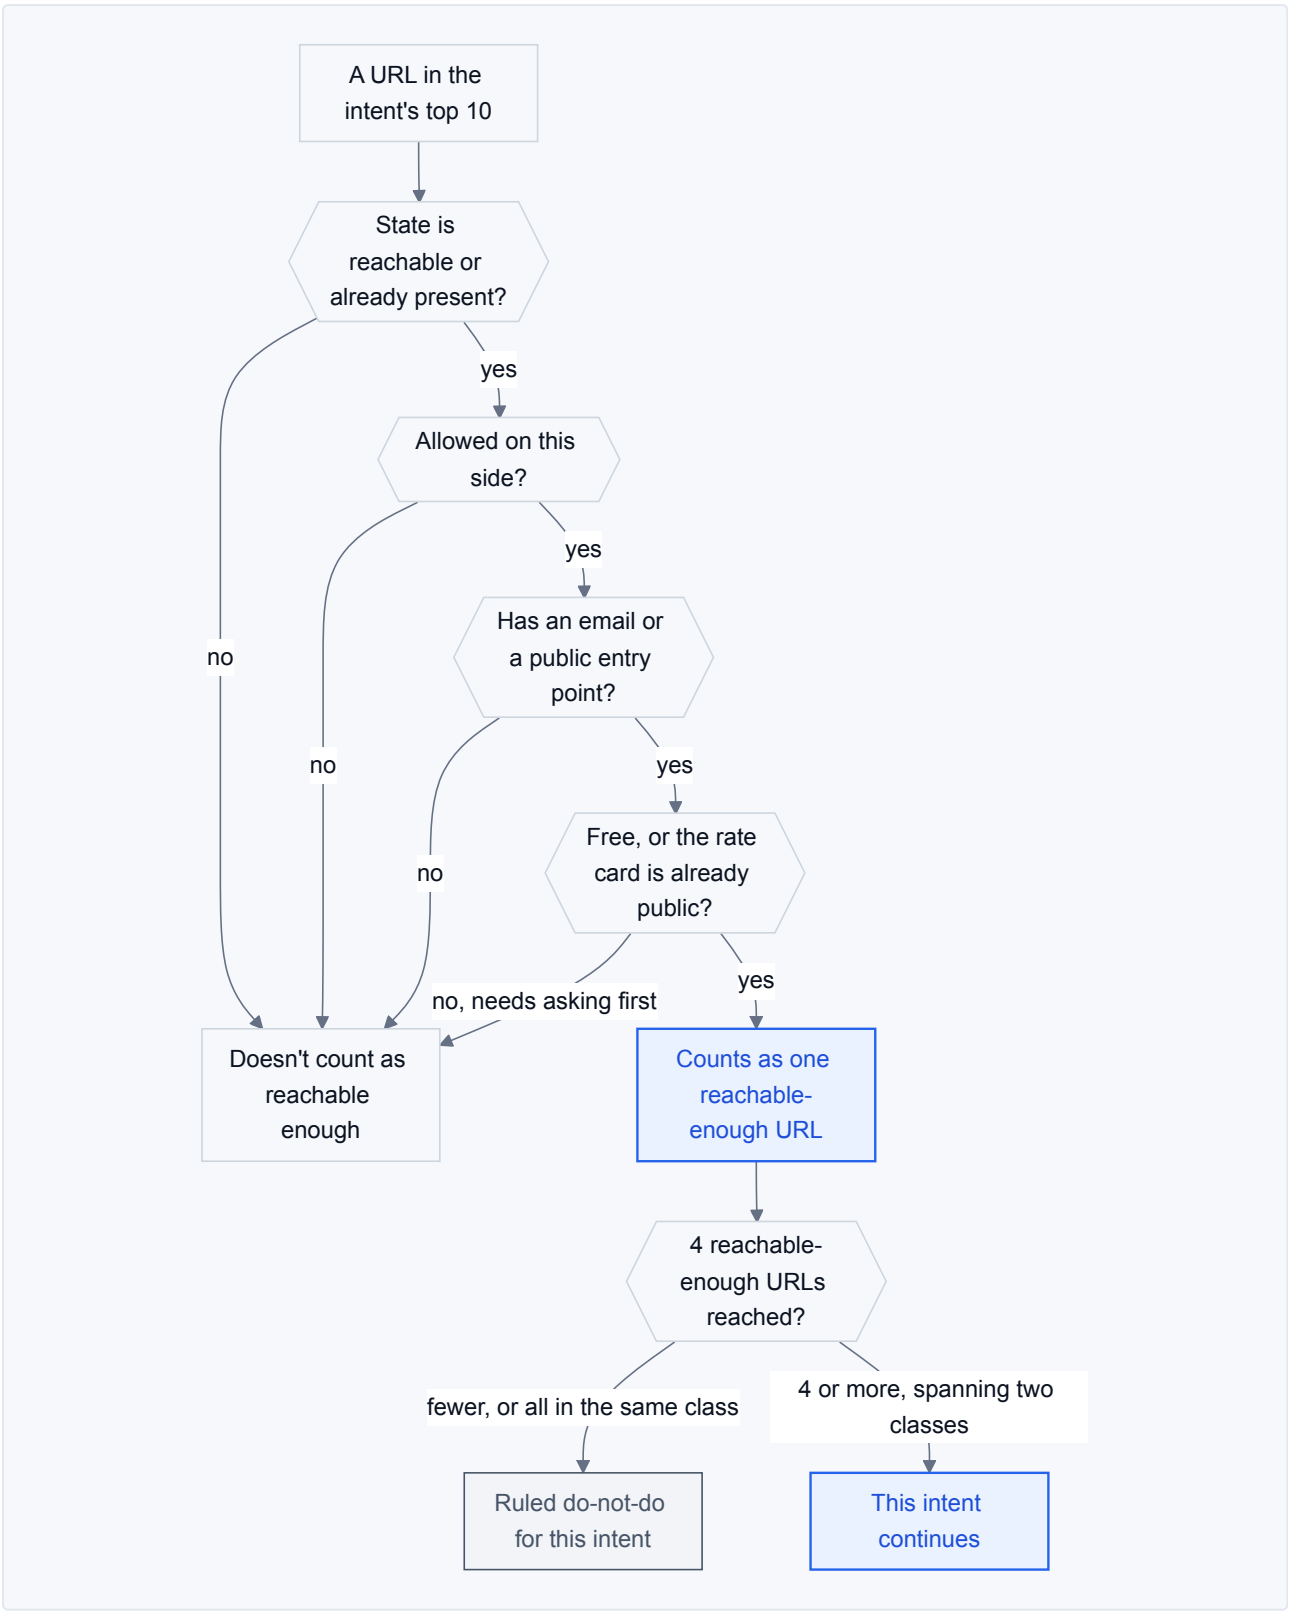


Figure: the abstain line — all four conditions must hold at once to count as reachable enough; 4 entries must also span two classes

Whether it's "allowed on this side" is checked by side: on the strictly regulated side, list pages whose title is itself laudatory (the Best / Top N kind) don't count as reachable enough; the verdict is "request removal": if you are already listed, send a removal request, and if you aren't, leave the page alone; the

full rule is in → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. If all 4 fall into the same action category, count them as 3.

1
Denominator
Intent-level top 10

2
Split into four classes
already present / reachable · lists and directories / reachable · maps and reviews / reachable · own site

3
At least two classes
Bet on only one class, and if that class won't move, the whole quarter has no fallback

4
At least two per class
Only 1 = betting on a single page — one redesign and it drops to zero

5
Multiply
2 classes × 2 entries = 4

Figure: how "4" was derived — anyone can recompute it; don't substitute a different number

Which list to read at which stage

Which of the three to read at each stage is fixed; never mix them up (the ban on writing a bare "Top 20" or "top 10" is in → General Edition D.2 Glossary):

Stage	Which list to read	How "to ask" is counted	Can it be reported externally
Deciding whether it's worth starting (4.8, shortcut 1)	A separate coarse-screen top 10 for each of the three questions, not merged	Counted as reachable (this half-day exists to talk you out of bad projects; better to re-judge at the baseline than to rule a URL out here)	No — n = 1, single engine, only serves the single decision "do it or not"
Picking targets and off-site	The baseline's intent-level top 10	Must already be cleared to zero by now	Yes
"On the list X/20"	Only reads the account-level Top 20, unchanged for the whole quarter	—	Yes

4.7 Actions by state, the ten question types and this quarter's slots

What you'll do in this section: assign an action to every URL by state — send an update letter first for anything already present; decide, for each of the ten question types, whether it means writing a page,

going off-site, or dropping from the pool; break each priority question down into the things AI will check, and match each one to a page or an off-site entry; then use the three combination rules to lay out this quarter's slots. When you're done, you'll have a list of actions by state and a finalised table of this quarter's buyer-type slots.

What action each state gets

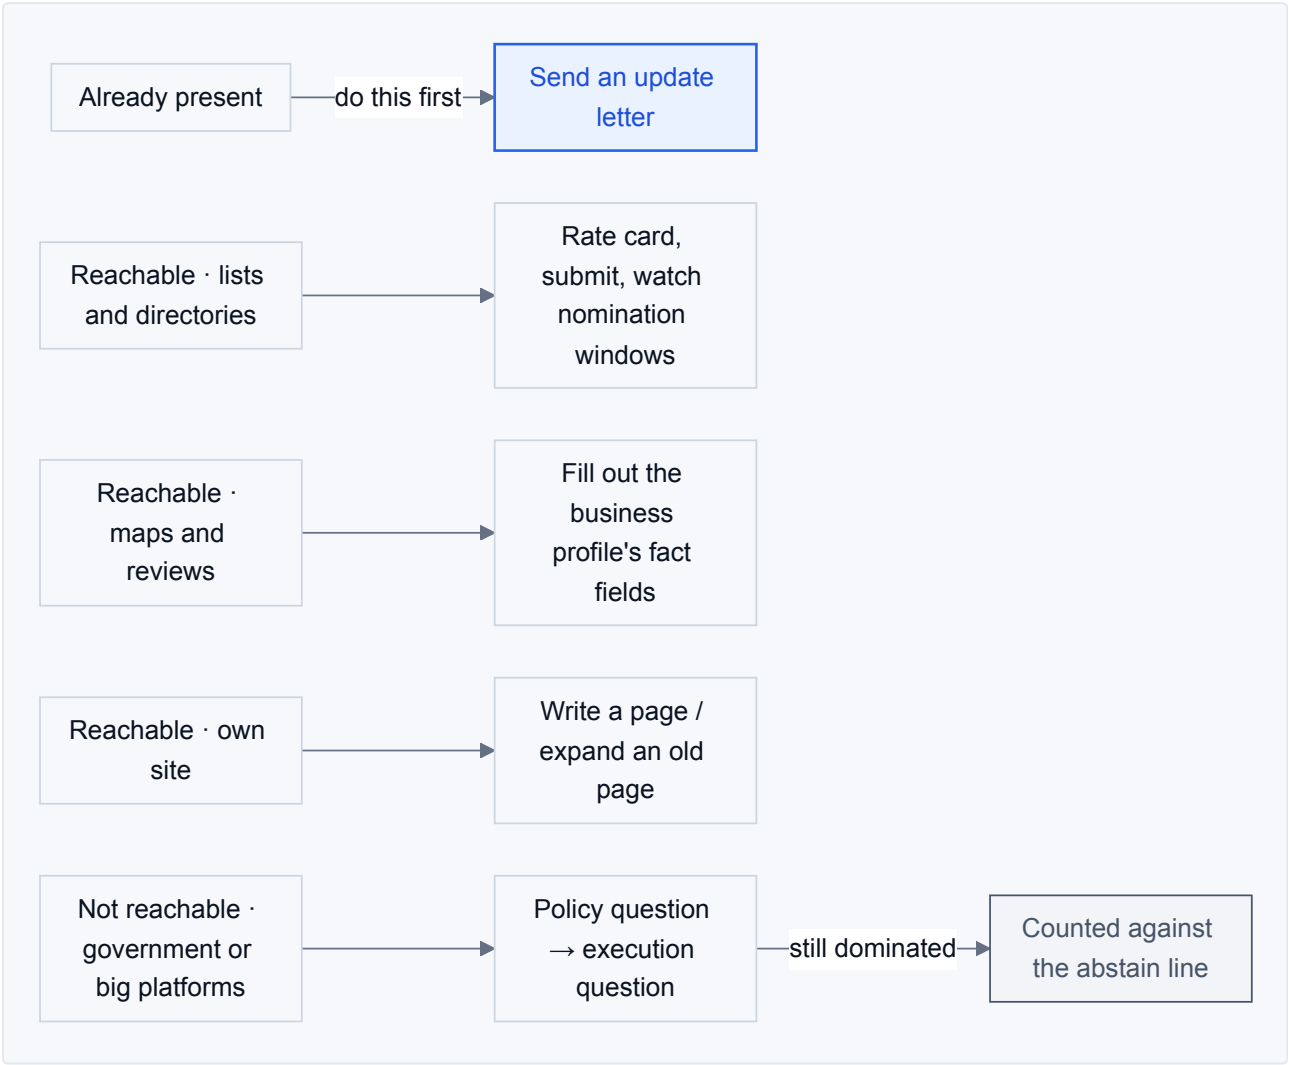


Figure: the action for each of the four states — already-present goes first; the own site is not the default main battlefield

For "already present", send the editor an update letter: new price (written per side), new credentials, new photos, an 80-word description. Adding content into a container that's already cited shows results in days; building a new page yourself takes 2–8 weeks (inferred, 65% confidence).

For "reachable · lists and directories", ask about the rate card, submit, apply for a listing, and watch for the entry window on nomination-type lists — schedule it separately; on the strict side, first apply the separate rules for payment and for wording, see → General Edition 6.2 The send gate, paid listings and the citation-slot table.

For "reachable · maps and reviews", only fill in fact fields, and also claim Apple Business, Foursquare and Yelp: all free, about 2 hours in total. ChatGPT's local answers don't read Google Business Profile directly, and retrieval doesn't go through Bing alone either (70% confidence, evidence in → General Edition A.1 Evidence for mechanics, the door and identity).

Only for "reachable · own site" do you write a page or expand an old one: most citations are on third parties, so the own site is not the default main battlefield — it becomes one only when this cell says so.

"Not reachable · dominated by government or big platforms" must not be ruled dead outright — doing so would wrongly kill a batch of regulated industries. Retry from another angle once first, turning the policy question into an execution question; only if it's still dominated after retrying does it count against the abstain line.

> **Example** (local services) "Can this subsidy be used for this service" is a policy question; turned into an execution question it becomes: "If you do it with us, how much do you pay after the subsidy is deducted, and do we help you apply?"

Where each of the ten question types goes

Question type	Where it goes
Second opinion / remedy	Priority; scheduled into the first batch (wording below)
How much / cost breakdown	Priority; write price by side
Symptom + option + price, three-part	Priority; this is a shelf question
A vs B, which is better	Priority; lands on a qualifying comparison page
Who to recommend / which is best	Priority; off-site gets you on the shortlist; your own pages get you through the check
Which one nearby / near me	Priority; lands on business-profile fact fields
Brand question / correction	A category of its own, takes slot 1 this quarter
Scenario question (does this suit my case)	Do it once it passes the three checks
Pure symptom type / what is it, should I see someone	Drop from the pool; this is an educational question, not the "who to buy from" stage

Figure: where each of the ten question types goes — shelf questions stay, educational questions drop from the pool

Price-type questions have a high rate of triggering AI Overviews, and mixed-intent questions an even higher one — numbers and sources in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests.

For "A vs B, which is better", the main action is "make a qualifying comparison page appear under this question": comparison pages' citation density is counted by page type, not by domain ownership — whose domain the page sits on is a secondary question; compare options, not businesses. Full rule in → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. "Who to recommend" works in two steps: AI first searches up a shortlist, then checks the businesses one by one (see "Question breakdown" below). The shortlist comes mainly from off-site — best-of blog posts already take up a large share of ChatGPT's citations (A.2), and directories and business profiles come in at this step too; for the check, AI searches with the named business's name or `site:` its website, and what it uses there are that business's own practitioner pages, service pages and price pages. So this cell needs work at both ends: off-site gets you on the shortlist; your own pages get you through the check. On the strict side, never seek a listing on Best/Top N pages or pay for one; if you are already listed, ask the publisher to remove you, see the same section 5.3.

Brand question / correction is the only one where "stopping means falling back", and also the cell you can see with your own eyes and reproduce yourself.

The wording for second opinion / remedy questions is fixed (applies together to the H1, H2, body, title and off-site material): the page's subject must be the buyer's situation, not another business's conduct. Write only the checks and how long they take, the available options and the conditions each one suits, itemised fixed prices, and what to bring. Do not write success rates, do not narrate what happened, do not show before-and-after comparisons; do not point at or disparage other businesses — named or unnamed is equally banned. In external documents, call this group only "second-opinion clients", never a label like "failed-case clients".

> **Example** (dental and aesthetics) Second opinions such as orthodontic relapse or a loose implant have a price per order of S$8k–18k, are the least price-sensitive, and the buyer is bound to search on their own — so schedule them into the first batch.

How to describe the current state of second opinion / remedy questions: in industries already measured, this type of phrasing shows no business named at answer level, and its citation slots are already occupied (internal ranking basis: region not locked, single engine, one-off, 85% confidence, evidence in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests). This is "no one has been named in the answer yet", not "no one is doing this" — it is the question with the best chance of an exclusive seat within 90 days, but the work must count those citation slots into the citation-slot table too, not just write an own-site page. For your own industry, go by what your own baseline shows. **Say only this one line externally**: "Who AI names for these questions today, and whose pages it cites as sources, we will count on the spot in the baseline." Never say "no one's doing this" or "competition is near zero", and never give a number of businesses — the moment the other side searches for themselves next week and sees several businesses, you lose their trust.

For symptom-type questions, keep only one distinction: a symptom question that includes an option and a price is a shelf question; a pure symptom question that only describes the symptom gets cited from authoritative educational sources by AI and almost never names a specific business — drop it from the pool directly, don't sit it in the monitoring layer first.

> **Example** (dental and aesthetics) "My gums are swollen: do I need the tooth pulled, and how much does an extraction cost" is a shelf question — priority; "why are my gums swollen" is a pure symptom question — drop it from the pool directly.

Question breakdown: what AI checks for each priority question

Before it answers, ChatGPT searches the web, and it doesn't just search the buyer's own words once. From 225 stored ChatGPT answers we pulled out the searches it actually ran (101 multi-industry, 60 law firm, 64 dental and aesthetics): for "who's best" questions, the median answer ran 11–16 searches (dental 11, multi-industry 16); for information questions, 6–8.5 (dental 6, multi-industry 8.5). Of the 743 URLs cited in answers to "who's best" questions, only 49 (7%) ranked in Google's top 10 for the buyer's own words, while 363 (49%) ranked in the top 10 for one of AI's own searches (Google is only a proxy; for how this is counted, see → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests).

The words AI adds are its own checklist, not the buyer's words. Of the 995 multi-industry searches for "who's best" questions (one search can fall into several categories), 373 were about credentials and official sources, 238 about price, 208 added a year, 105 were about reputation, and 247 used `site:` to limit the search to one website (about 150 of them to one business's own website); only 79 still carried best or top. The process roughly has two stages: first search up a shortlist, then check the businesses one by one. The named business's own pages are used mainly in the check stage.

So once each priority question on this quarter's list is settled, break it down once more:

What AI checks	Which page or off-site entry answers it
Credentials and official registers	Registration fields on the practitioner page + the official register lookup
Price	Price guide page, single-service price page
Trade terms (the buyer's words swapped for professional ones)	Service pages and definition pages written in trade terms
Reputation	Business profiles; handle reviews by your side
Checking each named business (by name or site: website)	The named business's /facts, practitioner page, price page
Finding candidates	Off-site: directories, profiles and third-party pages
Year	The page states its update date

Figure: every category AI checks needs a page or an off-site entry to answer it

How to break it down: run the question through the API at least 3 times (the ChatGPT web app doesn't show everything it searched), save the raw web-search queries from each answer, and sort each one into a category in the left column above; then check whether each category has a matching page or off-site entry in the right column, and fill whichever cell is missing. The baseline round has to be run in full anyway (see → General Edition 4.3 Coarse screen and piles, full rounds, and the brand-six baseline); save the searches together with the answers, and no extra probe budget is needed. For which pages each industry should have ready, see chapter 4 of your industry edition.

This is correlation, not causation: when the named business also has its own pages cited, it may be that AI first picked it on credentials and name recognition, then went to its website for material. Whether adding pages gets you onto the list has not been shown by any intervention experiment. So the question-breakdown table covers "when AI comes to check, do you have something to show it" — it doesn't promise a place on the list.

This quarter's combination rules and slot allocation

Who: the person setting the questions produces the list, the decision-maker signs · **Output**: this quarter's list (intent / buyer type / annual opportunity value / four-state action / acceptance tier) · **How to check**: the three rules are checked off one by one.

Once the list takes shape, check the three hard rules off one by one; it counts as final only if all three pass:

Rule	Why
Must have ≥1 hold-pile question	The hold pile's baseline isn't 0, so a change in position can be measured within 90 days
Must have ≥1 brand / correction question	Shows results fastest (what AI says often changes first once the fix is made, see 3.1), and you can reproduce it yourself
National head term with no location ≤1	Can't reach a determinable threshold within 90 days; each extra one you pick is one more entry guaranteed to end up written as "not demonstrated"

These three combination rules are the minimum for the quarter to reach any verdict, and they come before the annual-opportunity-value ranking — a quarter in which not one question can reach a verdict is an empty ledger, however valuable its ranking looks. But once the three are satisfied, every remaining slot follows annual opportunity value in descending order, and the order may not be adjusted for any reason.

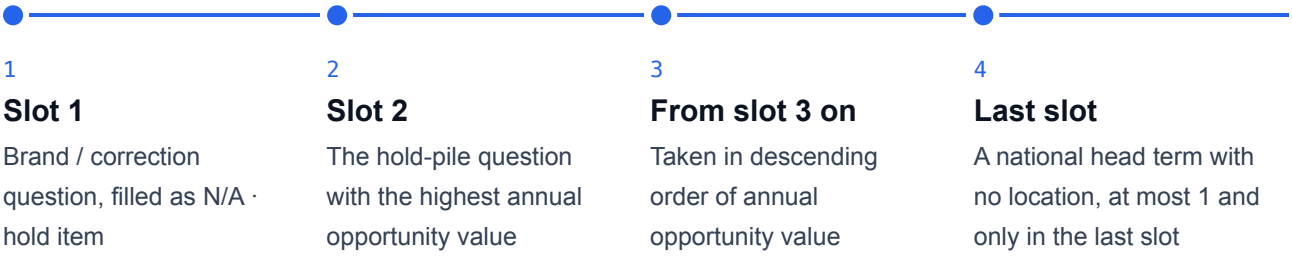


Figure: this quarter's slots are ordered this way — the combination rules come before annual opportunity value

Slot 1 is filled as N/A because a brand / correction question is not a buyer type, so it simply has no annual opportunity value to fill in. A high-value question squeezed out by the combination rules is logged into *Next quarter's first picks* by its annual opportunity value and taken up first next quarter — you don't go back and swap questions for it. Only when the three-pile list genuinely can't produce a hold-pile question or a brand/correction question do you go back to the three-pile list and swap in a question — that is not starting over. Every row in the list must state which buyer type it belongs to.

4.8 Two shortcuts and this chapter's checklist

What you'll do in this section: with zero probe budget, use shortcut 1 (1 hour 40 minutes) to first decide whether it's worth starting; when pressed for time or able to do only three things, use shortcut 2 (48 hours) to produce a baseline and a first work order. The output of both shortcuts never enters the contract or the acceptance record. Only once every item on the checklist at the end of the chapter is ticked off do you move into writing pages.

Shortcut 1: half a day, no probe budget, decide whether it is worth starting

Use this when you haven't yet decided whether to put this quarter's effort in at all. The criterion in one line: at the coarse-screen stage, take one top-10 source-URL list per question (three lists, not merged), and count the "reachable enough" entries per 4.6: fewer than 4 rules that question out, 4 or more always continues; at this stage, "to ask" is counted as reachable.

1
30 minutes
One cost question, one recommendation question, one scenario question; ask by hand, logged out, no probe tool

2
20 minutes
Copy a top-10 source-URL list for each question, three lists, not merged

3
40 minutes
Check each URL for a name and an entry point; don't send letters, don't ask about price

4
10 minutes
Count "reachable enough" and "to ask" for each list, write them onto the conclusion page one list at a time

Figure: shortcut 1 — four steps, 1 hour 40 minutes in total, to decide whether it's worth starting

The three lists in step 2 are called coarse-screen top 10, not the account-level Top 20 used as the acceptance denominator; merge them into one and step 4's three separate counts simply can't be done.

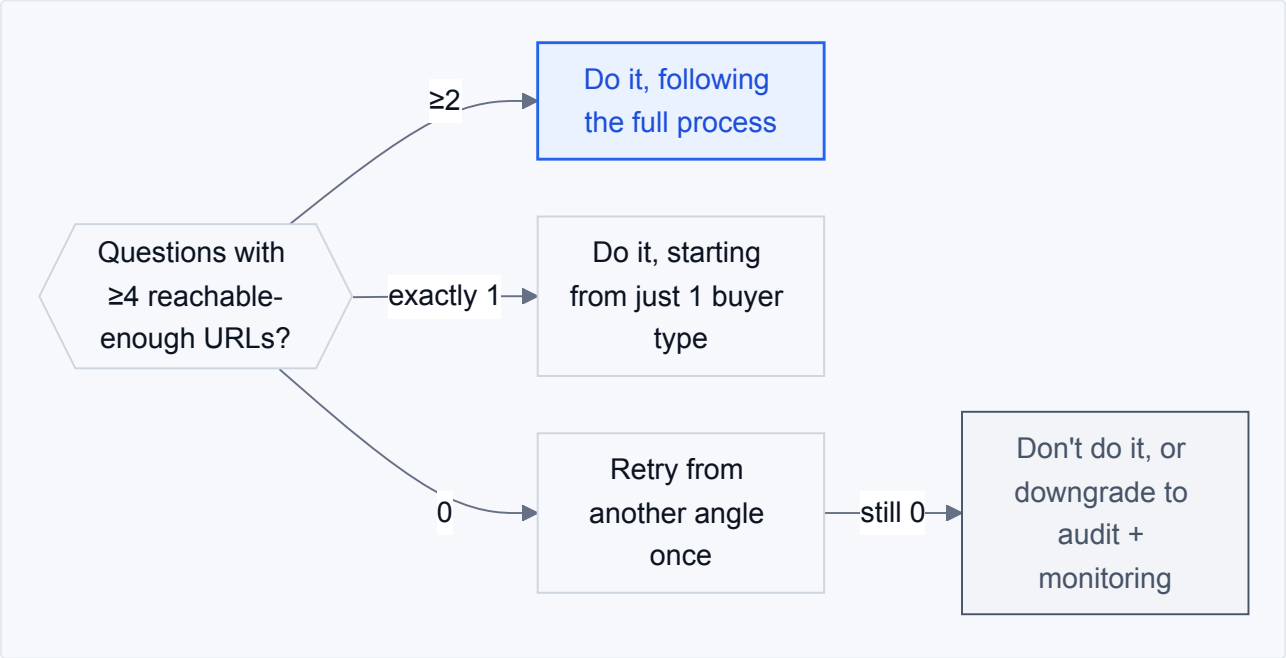


Figure: shortcut 1's conclusion — how the three questions are judged together

If you start from just 1 buyer type, add more types once the baseline has run. Retrying from another angle follows 4.7's "not reachable · dominated by government or big platforms" cell: turn the policy question into an execution question.

None of the numbers this step produces ever enter the contract or the acceptance record: the coarse screen is n = 1, single engine, by hand; a one-off, single-engine observation with n ≤ 3 must not be

used as a result or as acceptance evidence. Any conclusion-page reference to this step's numbers must carry, in the same sentence, "coarse screen, n = 1, single engine, <date>, not for external use".

Shortcut 2: a 48-hour quick start for picking targets

Take this route when things are urgent, or when you're starting with a small pilot first. **The output can only be used to pick targets — it can never be the acceptance baseline.**

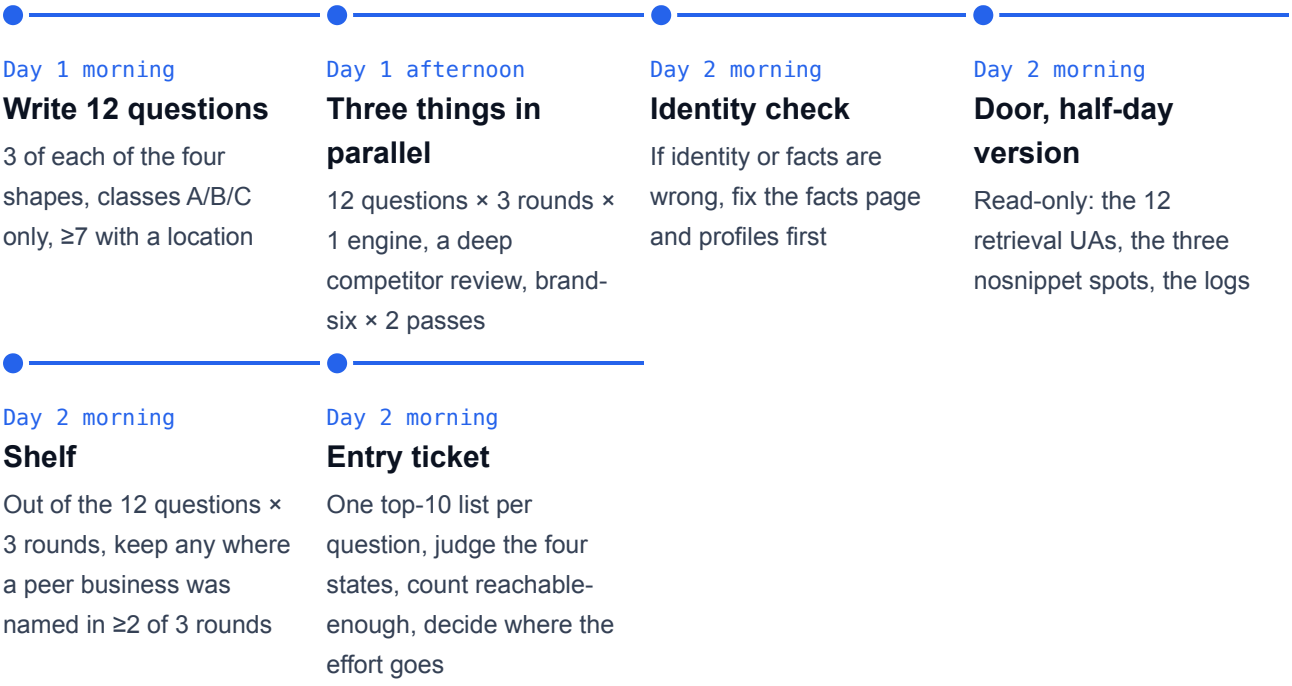


Figure: shortcut 2 — 48 hours to a baseline and a first work order

The four shapes are how much, who's best, A vs B, and which is nearby — never write an educational question. The half-day door version is read-only; don't touch anything. The deep competitor review takes the raw `domains_cited`; a person reads the brand-six answers for 15 minutes, used only to jump the queue and fix errors the same day — it is not a baseline, and it is not used to top up the full set either; the real baseline is run in full — 36 times — separately, in the freeze week (6 questions × 3 rounds × 2 engines, see 4.3). Day 2 morning's four signals take 2 hours in total; at the entry-ticket step, own site >40% of sources → work on your own site; lists and directories make up most of it → move the effort off-site.

The output is a one-page *Target selection and first work order*, with three coarse-screen top-10 tables and three "reachable enough" counts attached. Shortcut 2 is the baseline for the "only three things" path (→ General Edition 0.1 The book in one sentence, and the 90-day reading order), and it gives up a few things: n = 3, so a single question's conclusion is a direction, not a number; the dead pile can only say "no business named in the coarse screen"; there's no Chinese or minority-language coverage; there's no full baseline and no noise band — the full baseline and the noise band are filled in, in parallel, during the first week of work. **Shortcut 2's baseline must also be saved before the door is fixed** — this discipline doesn't relax just because you took the shortcut.

Checklist for this chapter

- [ ] **Mechanical gate: if an intent has no "before" sample, its page may not launch**
- [ ] The ruler is validated: a question known to get cited was run first; if the citation didn't show up, the run stopped
- [ ] The question pool is signed: all four source labels present, the 30-question ratios met, checked by a second person, signed in person by the decision-maker
- [ ] The dead pile was re-read: a second person re-read the original answer text before it was finalised, up to 3 questions lifted back into the battlefield pile
- [ ] Full rounds, all five run: battlefield pile + hold pile, five rounds each × two engines; a single screenshot on a single day does not count as a baseline
- [ ] The noise band was measured: same-week retest, the value written into `noise-band.md`
- [ ] The web leg ran first: run before the account-level Top 20 was frozen, unique domains forced in, `sources-top20.csv` has a `web_only` column
- [ ] Counts kept separate: listed as a source and written into the answer recorded separately, URL-level and answer-level separate, failed samples listed separately, no single-question small numbers
- [ ] The control agreement is signed: same-site pages not to touch, three competitors, the primary number to watch, and when to stop are all written down
- [ ] The shelf snapshot was reviewed page by page: anything naming a third party is marked internal, not put up for display
- [ ] The buyer type table's price per order, gross margin and monthly capacity cap were confirmed in person by the decision-maker
- [ ] The three-pile list has been reviewed, the three checks completed, and every entry in the four states has a judged-on date / basis / judged-by name
- [ ] This quarter's list passed all three combination rules, every row states its buyer type and carries its annual opportunity value and check-2 action, anything dropped from the pool has a record, and the decision-maker has signed; if it falls short, go back to the three-pile list and swap in a question — that is not starting over

Coarse-screen numbers never enter the contract or the acceptance record; any reference to them must carry, in the same sentence, "coarse screen, n = 1, single engine, <date>". Only once every item on this checklist is ticked off does writing pages unlock — the first page starts with the price guide page, see → General Edition 5.4 The 30-day writing order and the three mechanical gates.

Chapter 5

Writing pages (1): the 46 page types AI cites, and the rules for every page

The page-type map, the limits by how tightly you are regulated, and the rules every page follows

5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side

5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers

5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions

5.4 The 30-day writing order and the three mechanical gates

5.5 The page unit: one intent cluster, one page

5.6 Two legs, two kinds of sentence

5.7 Seven rules that hold for every page type

5.8 The conclusion block, the density gate and the source gate

5.9 Content form, FAQ, Chinese pages and other-language pages

5.10 Shared parts for all page types: six checks for numbers, no public price, billing units, dates and schema

5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side

What you'll do in this section: look at one master map to see which of the six families each of the 46 cited page types belongs to, and which leg mainly cites each family. Then use the index tables that follow to look up each type's id, evidence tier, which leg mainly cites it, typical length, and availability by side. Schedule your work by these 46 types, not by the content-format list.

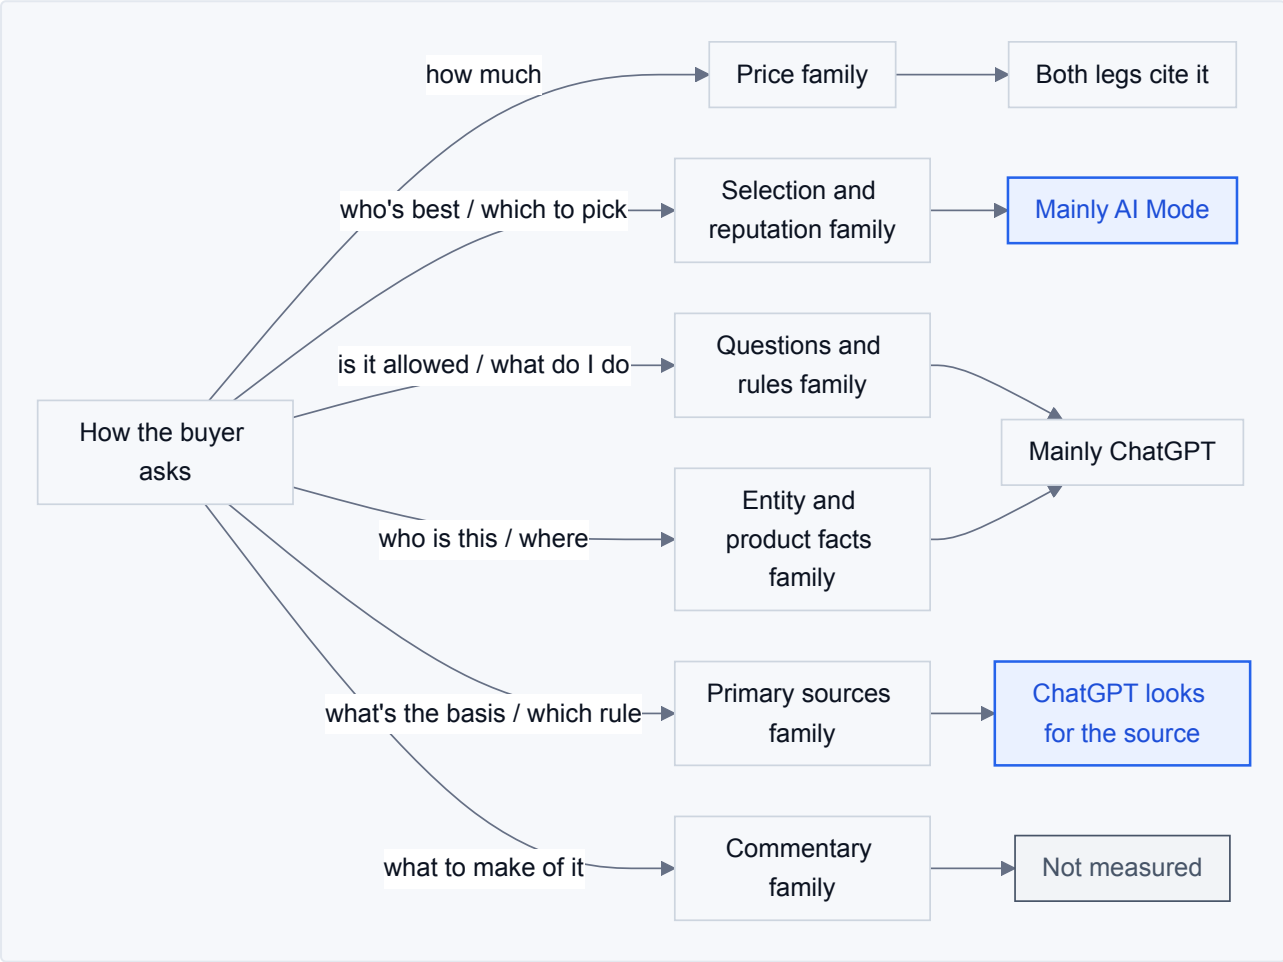


Figure: The buyer's question shape decides the page-type family; the family decides which leg mainly cites it

Two things the figure cannot show. First, each type belongs to exactly one family and is never double-counted across families. Second, "which leg mainly cites it" is a family-level tendency: there are still exceptions inside the same family, and the leg for each individual type is in the index tables below.

The nine tier-A types have been measured on both legs, so they are the best-evidenced. The figure below looks at the same batch of tier-A types to see which lean ChatGPT and which lean AI Mode:

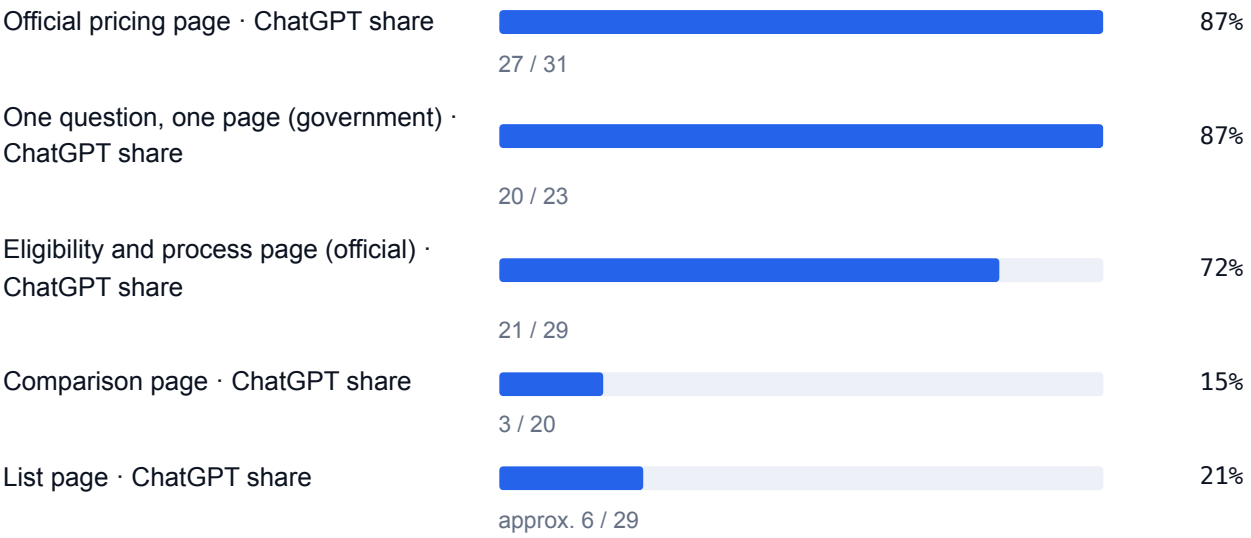


Figure: Within tier A, official pricing and the official versions of one-question-one-page and eligibility-and-process pages skew ChatGPT; comparison pages and lists skew AI Mode

Why the split works this way: ChatGPT looks for the source (official pricing, government Q&A, process conditions) and then writes its own conclusion; AI Mode looks for second-hand summaries (comparisons, lists). Remember the exception: the one-question-one-page and eligibility-and-process figures above are for the government or official versions; when commercial sites write the same page types themselves, those pages appeared only on AI Mode in our measurements, on a very small sample (evidence in Appendix A.3).

46-type index, grouped by the six families

How to read this: **A** = measured in round 1, on both legs, followed by which leg mainly cites it; **B** = measured in round 2, on the ChatGPT leg only; **C** = seen only in an external teardown, not measured by us. Length is English word count. Text after "switch E →" shows how the cell changes once switch E (peer comparison banned) is on. **There are only four availability values**: build = may self-build, adapt = may do after rewriting, cite only = downstream citation only or fill in your own row only, stop = do not build. Labels follow the rule text; a cell without a label does not mean you may relax it. On the unregulated side you may self-build by default; exceptions are noted under each table.

Price family (8 types)

id	page type	evidence · leg · length	strictly regulated side (strict side)	lightly regulated side (light side; "switch E →" = once switch E (peer comparison banned) is on)
pt01	Single-service price page	A · mainly ChatGPT · 62–1,609 words	Build (write only the fixed final price; Statute text)	Build (a range must state the span and what's excluded)

id	page type	evidence · leg · length	strictly regulated side (strict side)	lightly regulated side (light side; "switch E →" = once switch E (peer comparison banned) is on)
pt02	Price guide page	A · both legs, more AI Mode · 900–2,000 words	Adapt (no peer range; the official benchmark stands alone as its own section; label per the official-benchmark row in 5.2)	Build (a market range only as a stand-alone fact section with its source cited; Conservative line (not statute text)); switch E (peer comparison banned) → no peer range; feed the official-basis sentence and the procedure-variable sentence instead
pt03	Official pricing page (/pricing)	A · mainly ChatGPT · 1,691–8,265 words	Stop, do pt01 instead	Adapt (switch B (legally required fields) → make it an all-in price page for the course or product, matched item by item to the official approved name); switch E (peer comparison banned) → stop, do pt01 instead
pt17	Subsidy and limit rules page	B · approx. 1,600 words	Adapt (subsidies only as a stand-alone fact section, kept apart from your own price, no case-by-case confirmation tier; label per the official-benchmark row in 5.2)	Build (copy the subsidy rate from the official page and mark the check date)
pt27	Parameter and rate basis page	B · 150–503 words	Adapt (must be the final tax-inclusive price; no ranges)	Build
pt37	Category list page	B · 885 words	Stop (no listing peers, no category-by-category comparison); do the organisation's own service catalogue instead	Adapt (do the organisation's own catalogue; the name matches your registration); switch E (peer comparison banned) → stop
pt38	Calculator page	B · 2,500–3,000 words	Adapt (only a process estimate; the output must not be a price commitment)	Build (the raw billing unit + a conversion formula); switch E (peer comparison banned) → write "final figure to be confirmed on formal engagement"
pt42	Time-limited promotion page	B · length not retrieved (page blocked)	Stop (a discounted price with an expiry date; Statute text); do pt01 instead	Build (switch B (legally required fields) → a discount must state the full validity period and conditions)

Exceptions on the unregulated side: the official version of pt17 cannot be self-built on any of the three sides; your own version only holds in industries that have a subsidy or limit system.

Selection and reputation family (8 types)

id	page type	evidence · leg · length	strict side	light side ("switch E →" = once switch E (peer comparison banned) is on)
pt04	Comparison page (A vs B)	A · AI Mode · approx. 1,200 words (brand's own) / approx. 4,000 words (media)	Adapt (compare approaches only, not businesses)	Adapt (compare approaches and delivery formats only, not businesses; Conservative line (not statute text)); switch E (peer comparison banned) → compare procedures only, not businesses (Statute text)
pt05	List page (Best X for Y)	A · AI Mode · 2,100–9,000 words	Stop (Conservative line (not statute text) at the page-type level; ranking is itself a comparison, and it runs into the Statute text ban on comparisons — a sign-off cannot produce a publishable version)	Stop (Conservative line (not statute text); the client can relax it with written reasons); switch E (peer comparison banned) → stop
pt24	Third-party single-business review	B · 1,200–5,000 words	Stop (Conservative line (not statute text))	Stop (Conservative line (not statute text))
pt25	Sentiment, forum and news pages (a defensive item)	B · varies	Stop (cannot be self-built on any of the three sides; only add facts, never a reactive comparison)	Same as the strict side
pt28	Review aggregate page	B · 800–2,200 words	Stop (a rating is itself a testimonial; Statute text)	Adapt (testimonials carry a name, relationship and year); switch E (peer comparison banned) → stop
pt29	Self-built reputation and credentials page	B · 2,300–75,000 words	Stop (testimonials, star ratings; Statute text)	Adapt (testimonials carry a name, relationship and year); switch E (peer comparison banned) → adapt (write only provable expertise, years in practice, and the year of any third-party list)

id	page type	evidence · leg · length	strict side	light side ("switch E →" = once switch E (peer comparison banned) is on)
pt41	Verdict-first page	B · 2,436 words	Stop (the verdict is itself a comparison; Conservative line (not statute text))	Adapt (a scenario table only; never "best" or "best value"); switch E (peer comparison banned) → stop (you may still borrow the "boundary conditions get their own section" structure)
pt43	Buyer's selection framework	C · 2,100–3,700 words	Adapt (cut the business-by-business comparison; keep only the evaluation criteria and question list)	Build; switch E (peer comparison banned) → adapt (same approach as the strict side)

Exceptions on the unregulated side: pt25 cannot be self-built on any of the three sides; if you put yourself first on pt05, you must disclose it.

Questions and rules family (12 types)

id	page type	evidence · leg · length	strict side	light side ("switch E →" = once switch E (peer comparison banned) is on)
pt06	One question, one page	A · government version, mainly ChatGPT · 40–1,500 words	Build (the answer makes no case-by-case judgement; move personal-safety symptom questions to pt08)	Build; switch E (peer comparison banned) → cost questions carry no peer range
pt07	Eligibility and process page	A · mainly ChatGPT · 900–2,500 words	Build (general conditions are not written as a case-by-case judgement; duration is only the process's own timing)	Build (quote the statutory time limits clause by clause; no case-by-case judgement)
pt08	Remedy and second-opinion page	A · cited by both legs · 900–1,600 words	Build (the [Safety] block order for personal-safety items must not be rearranged)	Build; switch E (peer comparison banned) → do not criticise or rate a previous provider (Statute text)
pt10	Regulatory obligations and penalties page	B · 366–3,650 words	Cite only (make it a regulation summary page; schedule it as a trust asset, not a traffic page)	Cite only

id	page type	evidence · leg · length	strict side	light side ("switch E →" = once switch E (peer comparison banned) is on)
pt12	Definition page	B · 200–2,875 words	Build (anything touching personal safety needs a named reviewer and a date)	Build (tie each term to its statutory clause number; no success rates)
pt13	Step-by-step procedure page	B · 650–2,600 words	Build (your own version writes "how it works here")	Build
pt14	Preparation and bring-list page	B · 200–860 words	Build (the second-safest page type on this side)	Build
pt16	Schedule and deadline page	B · 650–700 words	Build (booking windows state only the process's own timing)	Build (dates must match the registered name of the course or product)
pt22	Collected FAQ page	B · 613–1,370 words	Build (the single most useful type for a regulated client)	Build; switch E (peer comparison banned) → no peer comparison inside the answers
pt26	Policy hub page	B · 441–6,653 words	Build (process and the official position only, no peer-by-peer comparison)	Build
pt34	Change notice / old-vs-new page	B · 1,700–2,400 words	Build (state the rule and the date only, not a rise or fall)	Build (explaining a regulatory change is a natural strength here)
pt40	Misconception page	B · up to 7,500 words	Build (one of the few page types on a regulated side that can write a "conclusion")	Build (tie fact sentences to statutory clause numbers; never a veiled dig at peers)

Exceptions on the unregulated side: pt10 is downstream citation only on all three sides.

Entity and product facts family (8 types)

id	page type	evidence · leg · length	strict side	light side ("switch E →" = once switch E (peer comparison banned) is on)
pt09	Entity anchor page(s) (organisation / person / PDP)	A · the organisation page mainly ChatGPT, the PDP ChatGPT only · organisation page varies, PDP 400–1,500 words	Build; testimonials, star ratings and before-and-after photos: stop (Statute text; no confirmation route)	Build (expertise must be backed item by item with evidence); switch E (peer comparison banned) → no statement of standing relative to peers
pt11	Store / branch page	B · 60–1,500 words	Build (the safest page type on this side)	Build (whether you offer a given service goes in the same line as the address)
pt21	Vendor legal terms page	B · 1,500–9,800 words	Adapt (its equivalent is the data handling and confidentiality policy page)	Adapt (same as the strict side)
pt30	Verification and lookup page	B · 267–716 words	Build (one of the safest trust assets; link only to government sources)	Build
pt31	Help centre / how-to page	B · 660–2,500 words	Adapt (its equivalent is the operations page: how to book, reschedule, get results; process only, no outcomes)	Build
pt33	Third-party directory listing	B · 128–800 words	Cite only (your self-supplied blurb is still bound by the price and advertising rules)	Cite only
pt35	Integration / marketplace listing	B · varies (JS-rendered)	Adapt (its equivalent is "which institutions and payment methods you work with"; name them, do not link out)	Adapt (same equivalent)
pt39	Trust centre / compliance proof page	B · 3,600–8,650 words	Build (a harder version: a licence and credential wall)	Build (a practising-qualification and insurance wall)

Exceptions on the unregulated side: pt33 is downstream citation only on all three sides.

Primary sources family (7 types)

id	page type	evidence · leg · length	strict side	light side ("switch E →" = once switch E (peer comparison banned) is on)
pt15	Legislation text page	B · 2,500–3,300 words	Cite only (a legislation summary page built on your own site was never cited in our measurements; schedule it as a trust asset, not a traffic page)	Cite only
pt18	Regulator guidance PDF	B · 20–62 pages	Build (the price section inside the PDF still follows the final-price rule)	Build
pt19	Statistics source page	B · 300–2,540 words	Adapt (data must pass the six checks and then the compliance officer; no success rates)	Build (no success rates)
pt20	Register / approved list page	B · 13–66-row table	Cite only (your own job is to fill in your own row)	Cite only
pt23	Official replies and speech records	B · 1,700–5,700 words	Cite only (make it an official-reply restatement page only; never imply the official body is endorsing your organisation)	Cite only
pt32	Device and product regulatory documents	B · 30–1,008 words	Cite only (your side-effects paragraph must be matched clause by clause to the official wording)	Cite only
pt36	Case-law page	B · 689–16,152 words	Stop (its equivalent is pt32)	Decided by the industry edition (holds only in industries with a body of case law; a summary only, never a record of wins, never a win rate)

Exceptions on the unregulated side: pt15, pt20, pt23 and pt32 are downstream citation only, or fill-your-own-row only, on all three sides; pt36 only holds in industries with a body of case law.

Commentary family (3 types)

id	page type	evidence · leg · length	strict side	light side ("switch E →" = once switch E (peer comparison banned) is on)
pt44	Concept pillar page	C · 4,500–5,000 words	Build (compare approaches and forms only, not businesses; must carry a named reviewer)	Build (no success rates, no peer comparison)
pt45	News and policy explainer	C · 1,250–1,400 words	Adapt (official figures stand alone in their own section; never use a new rule to solicit)	Build (the strongest type in this family)
pt46	Market observation page	C · 1,500–1,750 words	Stop (do not schedule it: no first-party market data, and it slides easily into promotion)	Stop (where it barely holds up, write it as pt45 instead)

Exceptions on the unregulated side: pt46 is really only suited to a service provider's own blog; we do not recommend scheduling it for clients.

A list of content formats is not a list of cited page types

The 30 content-format list	The 46 cited page types
Podcasts, short videos, livestreams, infographics, email	Not web pages; AI cannot cite them; no equivalent
Whitepaper	Roughly equivalent to the statistics source page (pt19)
Case study	Roughly equivalent to the verdict-first page (pt41)
Template	Roughly equivalent to the preparation and bring-list page (pt14)
Not on the list at all	The five most-cited types: regulatory obligations, store/branch, definitions, step-by-step procedures, preparation lists

Figure: The 30 formats are what you publish; the 46 types are which page AI copies from

"The 30 content formats" answers "what carrier are we publishing this month"; the 46 types answer "which page does AI copy a sentence from when it answers a buyer's question." The most-cited types are not on the format list because they are not "content" — they are fact carriers. **The number of page types is not a property of the industry; it is a function of how much question-space you have covered** — ask a fresh batch of questions you have not asked before ("what happens if you use X", "who do I call if X breaks", "what are the alternatives to X"), and this number will keep growing.

46 is a counted number: round 1 measured both legs only on purchase-decision questions ("how much", "who's best", "A vs B", "can this be done"), which turned up 9 types; round 2 switched to questions like "what is X", "how do I do this step by step", "what percentage", "what do I need to bring", "is this allowed", "which one is nearby", "what is this business like", "complete guide", which added 33 more; adding 5 from an external teardown and merging 1 duplicate brought the total to 46. Whether two things "count as the same type" is decided on exactly three criteria, and all three must be met — otherwise they are counted separately: ① the question shapes that trigger them belong to the same family; ② the copied sentence sits in the same position and form on the page; ③ the layout can be written from the same blueprint. The three disputed calls are in Appendix A.3.

A ruler bias: the 33 tier-B types were measured on the ChatGPT leg only (the AI Mode leg was not measured because the probe quota ran out), so the skew toward the Primary sources family is a bias caused by the ruler, not the whole picture; "not measured" does not mean "not cited" — page types that skew toward second-hand summaries (lists, comparisons, peer articles, video pages) could not be measured in round 2, and that is a known blind spot. We have no measurement of the 4 tier-C types (pt43–46); of the teardown subject's archetypes, the only one proven to work for that subject itself is the list, and the list is exactly the type the strict side can least use (evidence in Appendix A.4).

5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers

What you'll do in this section: identify the advertiser before you write a word, and stop and escalate if you cannot; before writing the price, the price hook in the title, promotions and freebies, or result numbers, check this section for what your side can write, what it bans, and what fills the same slot once something is banned — then check the row for any switch that is on for you. Together with → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions, this section is the whole book's only summary of "how to adapt for the regulated sides" — the page-type cards no longer rewrite all three sides; they leave one line pointing back here.

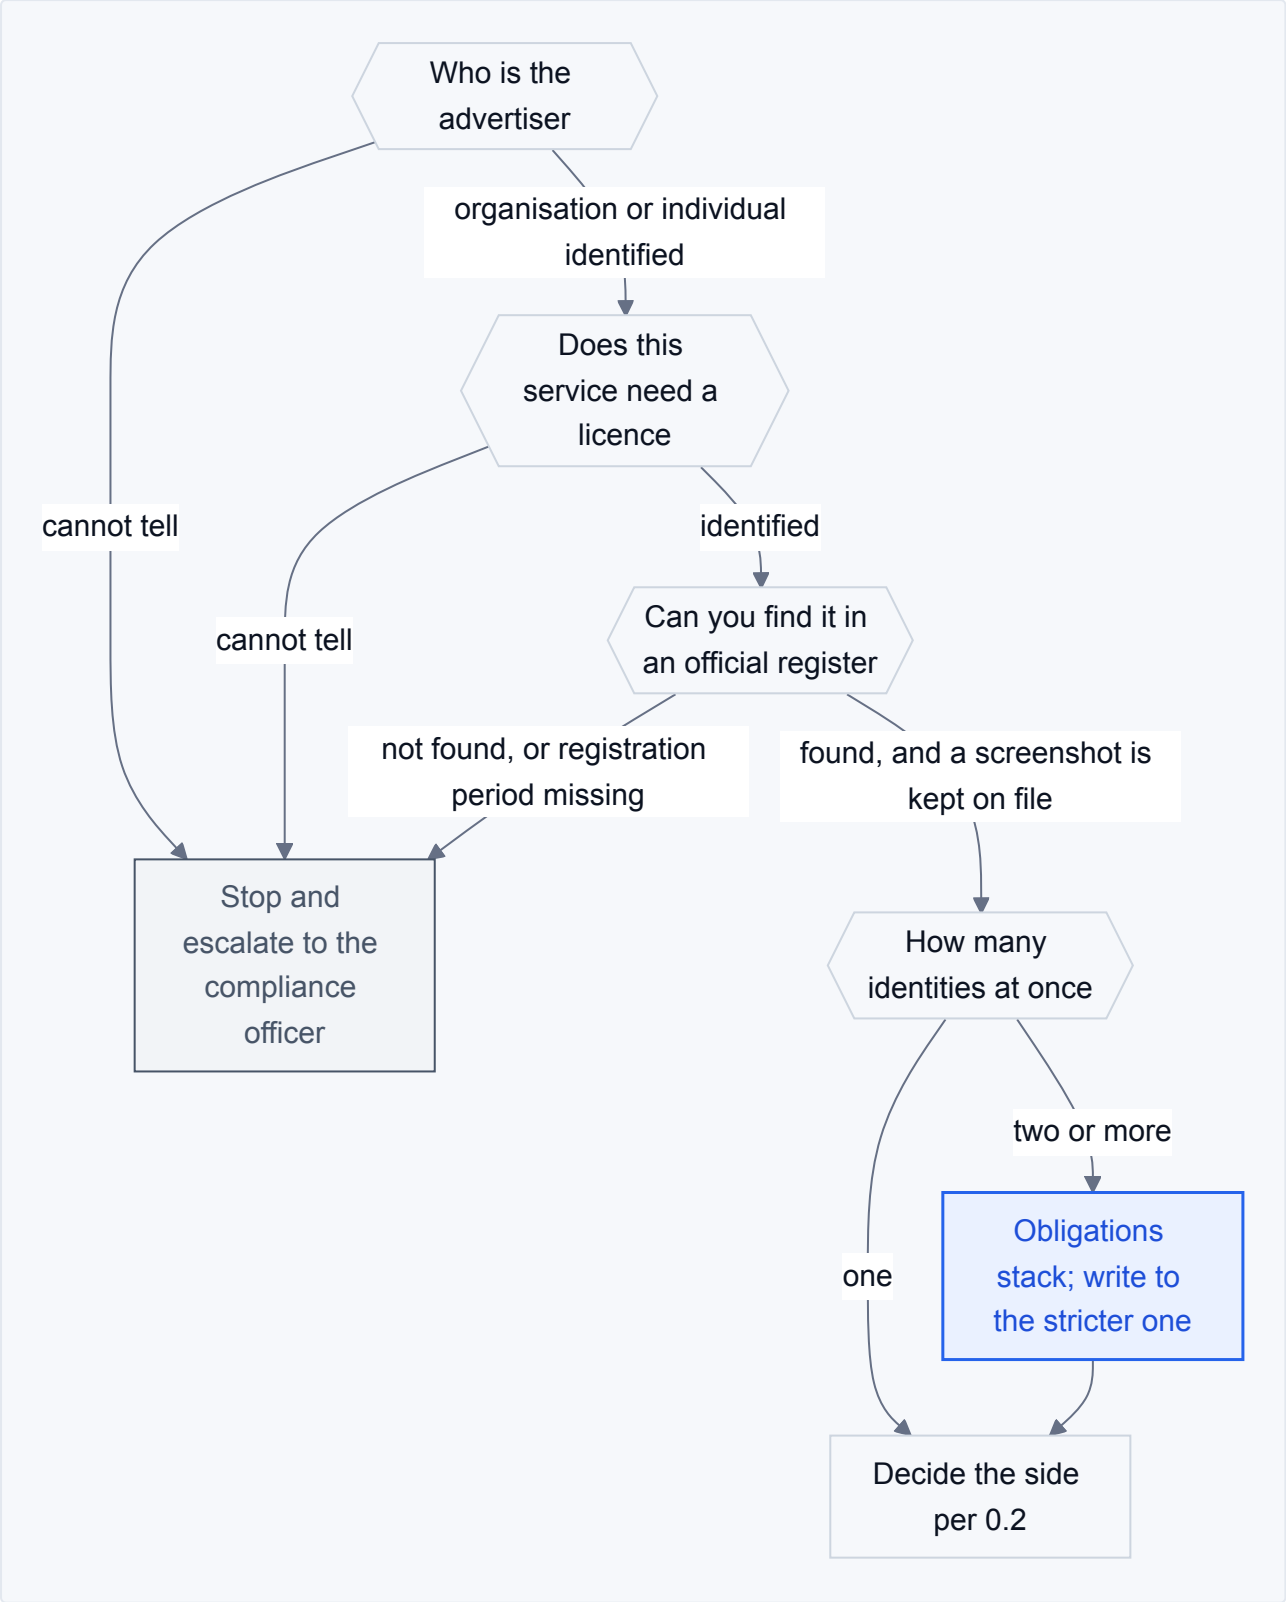


Figure: Identify the advertiser first; stop and escalate if you cannot

Identity is not the only reason to stop and escalate. The five stop conditions: **identity unclear; price cannot be a final price; a number fails the six checks in this section's result-numbers part; no client authorisation; comparison bases differ**. Hit any one of them, hand it to the compliance officer — never fill in a number yourself. Once you have identified the advertiser, decide the side per → General Edition 0.2 Decide which side you are on first. For how a stop-and-escalate item goes to sign-off, see →

General Edition 0.4 The three labels for compliance sentences, and the stop-and-escalate rule: a written sign-off resolves only the three stop-and-escalate situations; **it never turns something banned as Statute text into something publishable**.

Once you have decided the side, check whether any of the following switches also apply — a switch is not a new side; it is a condition added on top of the three sides:

switch	hard gate added once it's on
A (agency liability)	The party writing and publishing on the client's behalf shares liability with the client; do not publish with fields missing, and authorisation comes before work starts
B (legally required fields)	The legally required fields sit together in a fixed position, then repeat once more in the facts block; a mandatory disclosure item gets its own page; in some industries a missing field is presumed a violation outright — the fields are a hard gate before hand-over
C (regulated products)	The category you sell is itself governed by dedicated regulation; the copy is split by side according to the category
D (spans two sides)	If the business spans two sides, the whole site is built to the stricter side; it is not split by subdomain
E (peer comparison banned)	Peer price ranges and rankings are never written, whether named, unnamed, calculated by you or quoted from a third party
F (referral commissions banned)	You may only buy at a fixed rate-card price; you cannot pay per lead, per sale or by commission

Banned: writing this breaks the rules on the strict side		Replace: fill the same slot with this
Peer price ranges, market sentences	→	The official benchmark stands alone in its own section + a procedure-variable sentence
"From / as low as / original price / discount"	→	A fixed final price by tier, with the defining conditions spelled out
"Free replacement / free check-up"	→	"Included in the quoted price, once only"
The post-subsidy price in the same table as your own price	→	The subsidy stands alone in its own section, with the table number and update date
Success rate, rescue rate	→	Process duration, risks and the conditions it suits
"Most cases" / "extensive experience"	→	Organisation operating facts, backed by a ledger, free of adjectives

Figure: Every banned item has a legal replacement that fills the same slot

Prices: official benchmarks, subsidised prices, peer ranges, and whether market sentences go into feeds

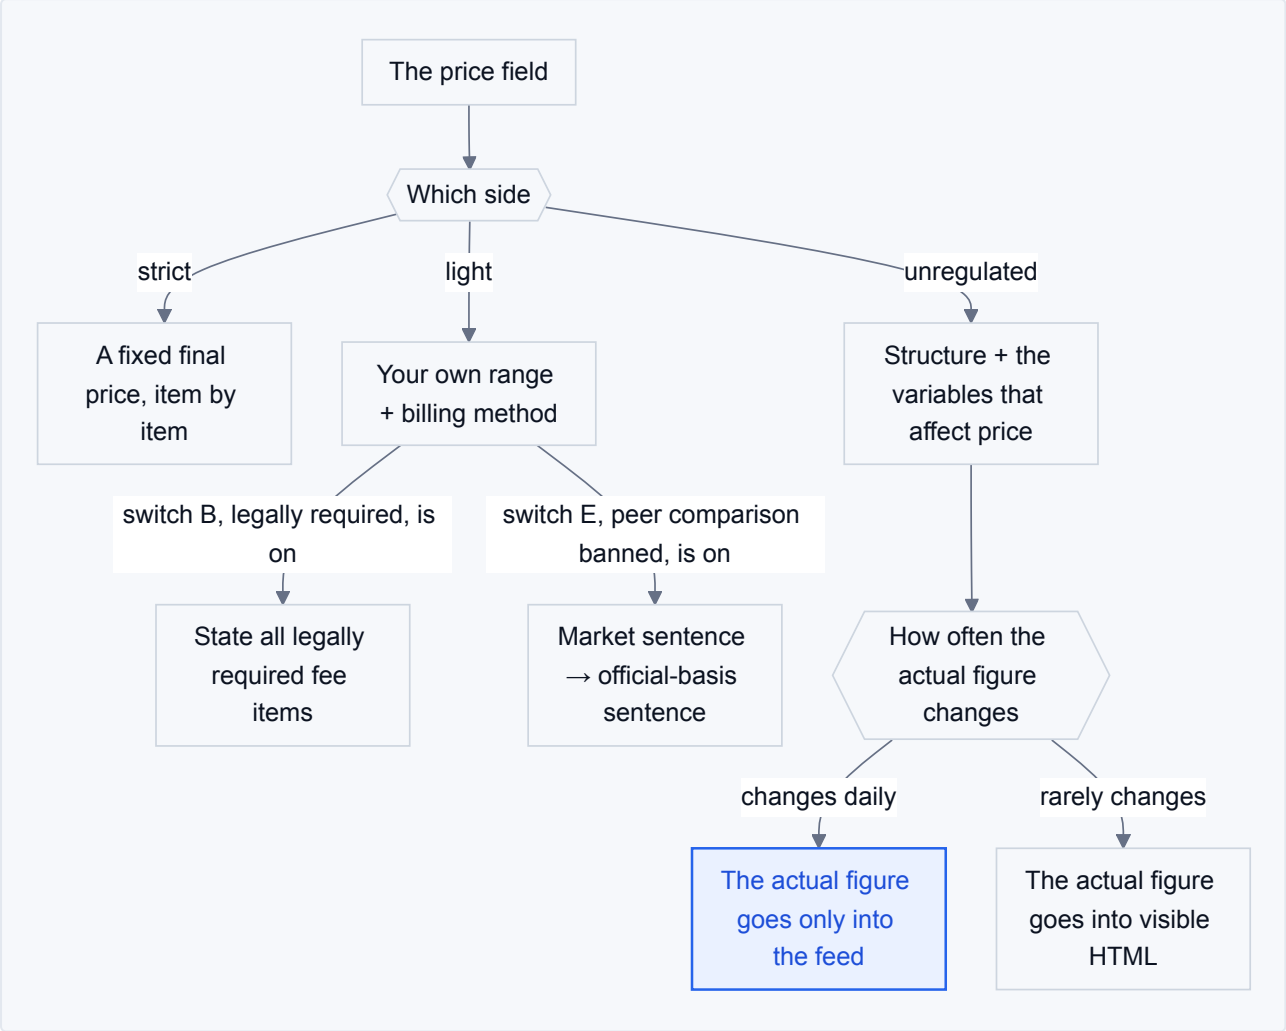


Figure: The price field is written by side, with one more split by half-life (how often the actual figure changes)

The detail on each side: on the strict side, write one final tax-inclusive figure for each specific service, and list what is mandatorily included; for case-by-case variation, use tiered, itemised pricing, with the defining condition spelled out for each tier — never just "mild / moderate / complex"; where a figure can only be set after an examination, write "confirmed after examination" and give the fixed price of the examination itself; you may state the effective date, what is included, and exclusions. On the light side, your own range must state the span, what is excluded, and the fees the official body charges; never write "lowest" or "cheapest"; once switch B (legally required fields) is on, write out the fee elements item by item (the name and form of the charge, whether it includes miscellaneous fees, how long the quote is valid, which terms apply, discount conditions and amounts, currency, tax, and whether a payment method is specified). On the unregulated side, write the structure: the billing unit, tier boundaries, the minimum starting price, the variables that affect price, the contract term, and how overage is calculated. When the actual figure changes daily, putting it in the body text means manufacturing your own factual errors, so it goes only into the feed; a product detail page still shows the real price on the first screen, and it must match the feed.

The official benchmark: on the strict side it can only stand alone in its own section, with none of your own prices appearing in that section, and it carries a table number and an update date; putting it in the same table alongside your own price counts as a comparative listing of prices, and no change of position can save it — there is no confirmation route (Statute text). Appearing in the same paragraph, an adjacent paragraph or the same Q&A, or writing "we are within the benchmark / below the benchmark", is Conservative line (not statute text) — the rule does not spell out these specific cases item by item, so the default is not to write it; relaxing it needs written reasons from the client's compliance officer, kept on file. On the unregulated side you may add an official-benchmark column to your price list. On the light side, write the fees the official body actually charges, copied verbatim from the official fee table, with the table name and date retrieved noted; if you cannot get the original table, write no figure at all — write "priced separately per the official fee table, see <link>" instead — never estimate, never carry over an old value.

The post-subsidy price: the strict side does not write "$X after subsidy" or "up to $X in subsidies" — anything that folds the subsidy into the price; the subsidy amount appears only in the official-benchmark section. The light side and the unregulated side may write it, but the post-subsidy amount stands on its own, noting the eligibility conditions + the official page's URL + the check date, and it is never written as if it applies to everyone.

Peer price ranges: never write these on the strict side or wherever switch E (peer comparison banned) is on — named, unnamed, calculated by you or quoted from a third party all count. The structure "give the market range first, then your own price" is itself banned (for what a market sentence is, see → General Edition 5.6 Two legs, two kinds of sentence); deleting the words "within the market range" does not count as fixing it — delete the whole sentence, and fill the gap per the next paragraph. Once switch E is on, the basis splits into two tiers: a named range, or one that can be traced to an identifiable business, is Statute text; an unnamed, anonymously aggregated range is Conservative line (not statute text) — but a range whose sample and basis cannot be checked is separately ruled out by the rule that information that cannot be checked is itself misleading (Statute text), so in practice it cannot be written either. On the light side, where switch E is not on: a market range can only be written as a stand-alone fact section with its source cited (the sample definition, source, date); it must not sit in the same table as your own price, in an adjacent paragraph, or in the same Q&A, and you do not write "we are within the range" (Conservative line (not statute text)). Only the unregulated side may put a market range and your own price in the same table.

What fills the empty slot: on the strict side, fill it with the stand-alone official-benchmark fact section and a variable sentence that carries no peer price. Where switch E (peer comparison banned) is on, fill it with two sentences: the **official-basis sentence** — swap "what peers charge" for "what the regulation / the regulator / the procedure itself sets", matching the market sentence's own position, word count and sentence form (one sentence in the conclusion block + one first sentence under "what affects the price", 40–60 words); and the **procedure-variable sentence** — its subject is the process or the treatment option, not peers, and it explains how your own price varies with the process.

Whether the market sentence goes into the feed: the market sentence and any range never go into the feed or into structured data — the feed and structured data hold only your own actual price; this

applies on all three sides (per platform merchant rules, not measured).

Off-site self-supplied blurbs: on the strict side, and wherever switch E (peer comparison banned) is on, the blurb you supply for a third-party directory is bound by the same price rules — no "From", no discount, no "best".

Price hooks in the title

The unregulated side	The strict side
"From $X" + a unit or condition	A fixed, final tax-inclusive price + unit + brand
You may also write a three-tier price	For several services: "Service A price · Service B price (incl. GST)"
Never write "cheapest" or "best price" with no source	Never write "From", "starting", "as low as", a range, a post-subsidy price, or a discount

Figure: The price hook in the title, at each extreme

The light side sits between the two extremes: you may write "From $X", but X must be followed immediately by the defining condition and any variable that could push the total up, or it amounts to leaving out a material fact; you may also write "$X nett + defining condition"; never compare prices with peers, never write a success rate. Once switch B (legally required fields) is on, the course or product name must match the officially approved name word for word, and the wording of your status may only use the word the law allows (such as "registered") — it must not be upgraded to "recognised / approved", and you do not write guarantee-type wording or score-improvement percentages. A post-subsidy price written into the title needs the eligibility conditions attached; never write it as if it applies to everyone. Titles with a price prefix and subsidy wording were cited in our measurements, but they break the rules on the strict side: ranking criteria do not apply to compliance (measured, 2 cases).

Promotions and freebies; rework and replacement policies

item	strict side	light side	unregulated side
Time-limited offers, countdowns	Stop: do not write a discounted price with an expiry date (Statute text)	Once switch B (legally required fields) is on, write out the full validity period and conditions; once switch E (peer comparison banned) is on, write only a fixed price or a tiered price, no time-limited discount	May write it; an expired price is taken down automatically

item	strict side	light side	unregulated side
Original price, discounts, price rises and falls	Stop: do not write an original price, a discount, or a rise or fall (Statute text)	Once switch B (legally required fields) is on, write out the discount conditions and amount in full	May write it; the original price must be a genuine price it was actually sold at
Free, complimentary	Stop: do not write "free", "complimentary", "at no extra charge" or "$0" (Statute text); write "included in the quoted price" instead	Once switch B (legally required fields) is on, write out the full terms in the same block; do not use "free" for an item that is really just deferred payment	May write it; it must be genuine and not misleading
Bundles, package prices	Adapt: write each bundle as one standing, fixed final price, one URL per bundle, never alongside a discount	May write it	May write it
Rework, review and replacement policies	Adapt: write only the time limit, the number of times, and what is included in the quoted price	May write it; write it as a fee term, not a guarantee of outcome	May write it, same as the light side

Promotions and referral incentives on the strict side are not themselves banned; they can only be disclosed in person at the point of payment (Statute text). This table applies equally to business profiles, off-site material and review replies. Once switch B (legally required fields) is on, the moment a promotion appears on the page, the full legal terms must be written in the same block (a complete description, the eligibility conditions, whether it can be declined, when it expires; a scholarship-type item also needs the amount, the payout schedule, the deadline, the conditions for payment, and the penalty). If there is no room, give a prominent link to the specific page that has them; a phone number or "ask for details" alone is not enough.

Result numbers: success rates, grades, employment, case counts, operating figures

Grades, case volumes, case counts, student numbers and ratings all count as "result numbers". Missing even one of the six checks means it does not go on the page: **definition** (what counts as one success), **denominator** (how many people or items in total), **period** (which year, which cohort), **sample selection** (everyone, or only those who stayed to the end), **calculation method** (where the data comes from), and **permission to publish** (client authorisation, personal data, whether industry rules allow it). Build the definitions table before you write, following → General Edition 5.10 Shared parts for all page types: six checks for numbers, no public price, billing units, dates and schema. Once all six are filled in, run the number through the by-side gate below.

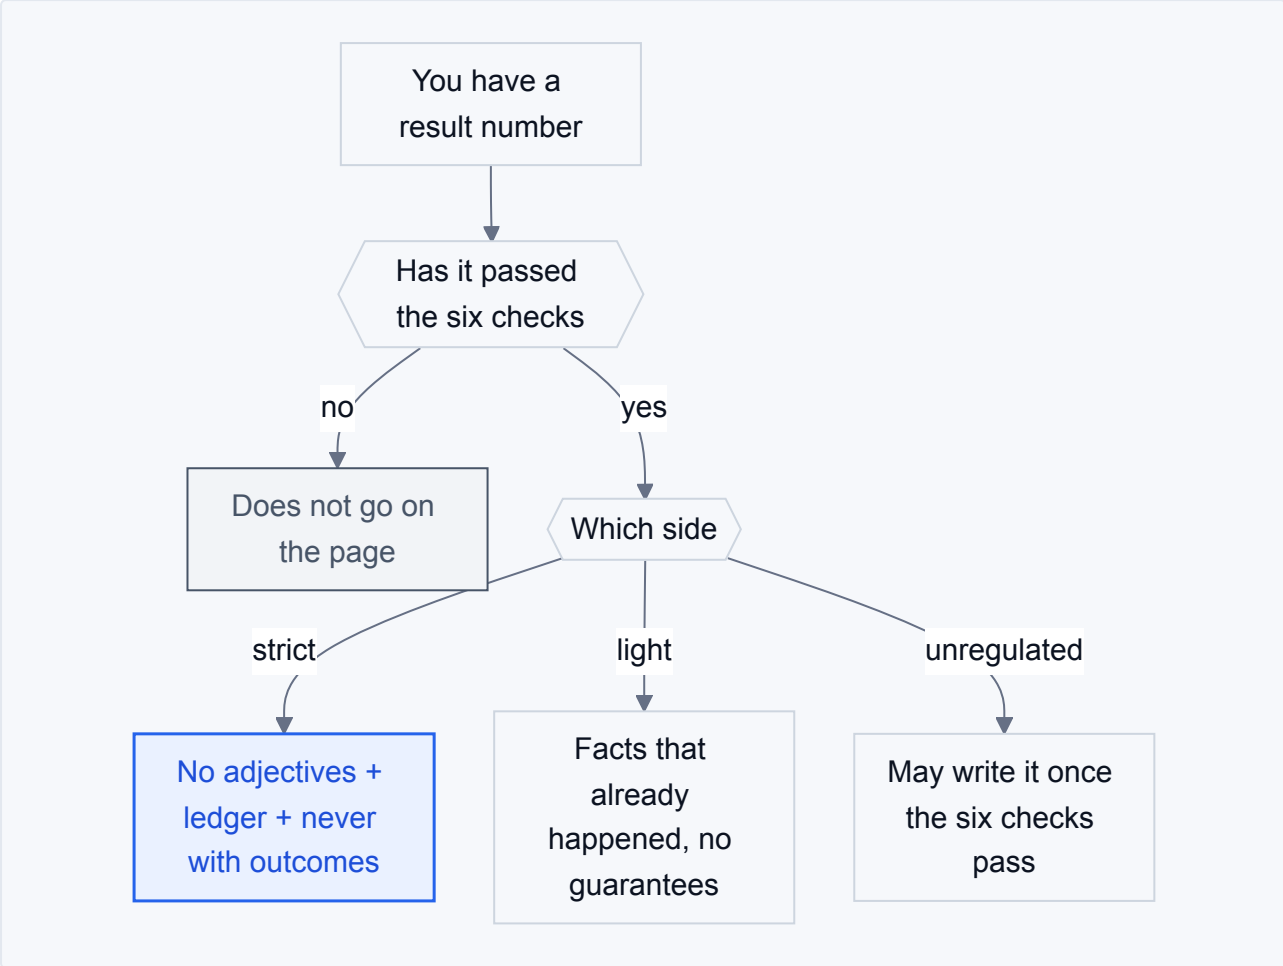


Figure: How result numbers are written on each side

Four gates on the strict side; missing any one means do not write it: ① no adjectives, no comparatives ("most cases", "extensive experience", "a large number of cases", "senior" are all banned, whether or not they are true); ② a ledger + sign-off: every number carries its basis definition, the statistical period, and the source system it came from, and the organisation's head signs off confirming it is accurate, with the export kept on file; ③ it is never placed next to outcomes, the conditions a service suits, or results; it appears only in the organisation-facts paragraph; ④ percentage price rises or falls are deleted entirely. This opening is narrow: what is banned is outcomes, results and every form of subjective praise; what is not banned is an operating fact with a clear basis, provable, free of adjectives (70% confidence, not measured; the signed compliance memo is the final basis). Numbers like case count and procedure volume are deleted by default — there is no confirmation process, and no tier that says "usable once confirmed".

The light side may only write facts that have already happened (how many were actually admitted in a given cohort, how many of N graduates were employed within a given time, the survey response rate), never a guaranteed result; once switch B (legally required fields) is on, there are also statutory banned phrases: do not write guaranteed admission, guaranteed employment, or guaranteed to obtain a licence, and do not write that "the qualification will be recognised by a given institution, employer or school". The regulator can demand the collection method and calculation, so the definitions table is built first, not patched in afterwards. In some industries the rules absolutely ban writing a success rate, and a

case volume may not be written or implied as one either. **The unregulated side** may write it once it passes the six checks, and it still must not mislead. One rule applies on all three sides: a result number cannot live only inside a chart — write it as a self-contained sentence that can be lifted whole, in the form "M of N, that is X%, over <time window>".

Rows for the price-family page types

id	page type	strict-side essentials	light-side and switch essentials
pt01	Single-service price page	The price list has no column for "results" (implying an outcome within a given time counts as soliciting)	No peer price mixed into the page
pt02	Price guide page	The official benchmark goes after the table, in its own section, in no column of the main table	Switch E (peer comparison banned): delete the peer range, swap in the official-basis sentence + the procedure-variable sentence
pt03	Official pricing page	No bundle cards, no recommended ordering — use pt01's row-format master table instead	Switch E (peer comparison banned): same as the strict side; switch B (legally required fields): make it an all-in price page for the course or product instead, each item matched word for word to the official approved name
pt17	Subsidy and limit rules page	The subsidy stands only as a stand-alone fact section, with no case-by-case confirmation tier; your own three-column table (item / limit / what you pay) does not hold up	Your own three-column table holds up on the light side and the unregulated side
pt27	Parameter and rate basis page	The parameter must be the final tax-inclusive price; write "which tier applies is set after examination" as one standalone sentence	Write the billing unit and the effective date
pt37	Category list page	No peer or category-by-category comparison — use the organisation's own service catalogue instead	Switch E (peer comparison banned): same as the strict side; switch B (legally required fields): the catalogue's name matches your registration
pt38	Calculator page	A process estimate only; the output must not be a price commitment	Switch E (peer comparison banned): write "final figure to be confirmed on formal engagement"; no comparison with peer prices
pt42	Time-limited promotion page	Adapt into a standing fixed-price page, one URL per bundle, drop the words "time-limited"	Switch E (peer comparison banned): write only fixed prices or tiered prices; switch B (legally required fields): write out the full discount terms

The unregulated side is bound only by consumer-protection law: the price on the page, in the feed, in structured data, and at checkout must all agree; the original price and any discount must be genuine; an expired price is taken down automatically. Where the identity rule is Original text not obtained, even a sign-off does not clear it: write the default wording from the book — it is not used as a basis for judgement, and you assume neither that the rule exists nor that it does not.

5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions

What you'll do in this section: before writing testimonials and reviews, comparisons, lists, titles, outbound links, or FAQ and captions, check this section's four diagrams and the by-side approach; the page-type rows give each affected type one pattern sentence, and each page-type card leaves only one line pointing back to → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side and these two sections, instead of rewriting the rules.

Testimonials and reviews

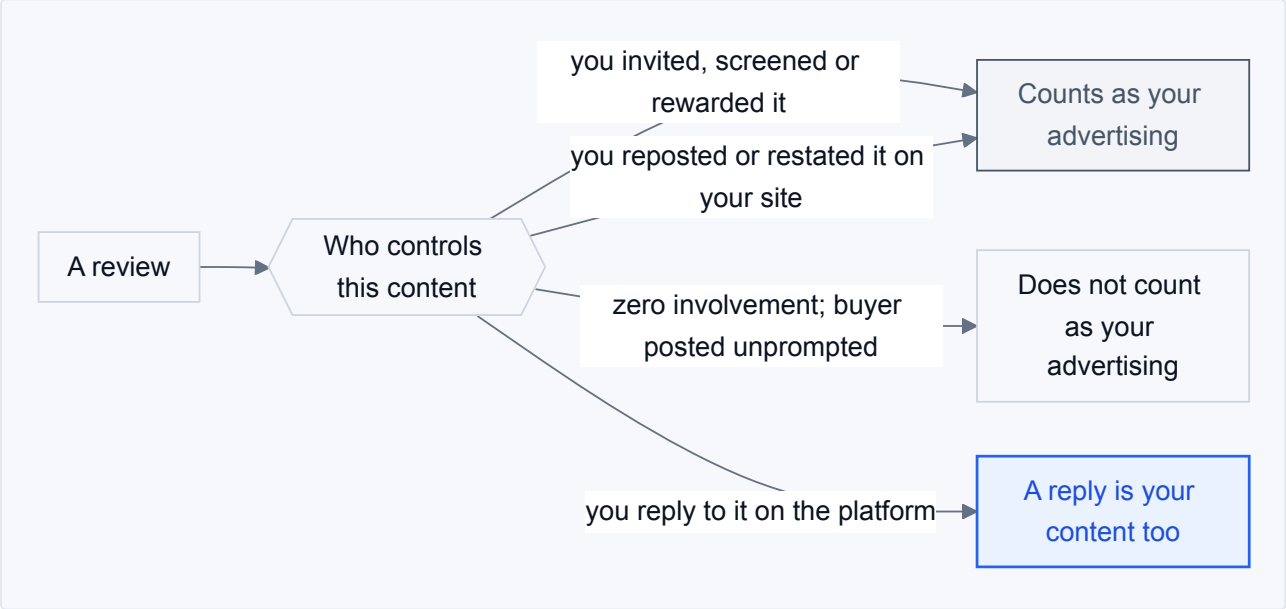


Figure: Whether a review counts as your advertising depends on who controls the content

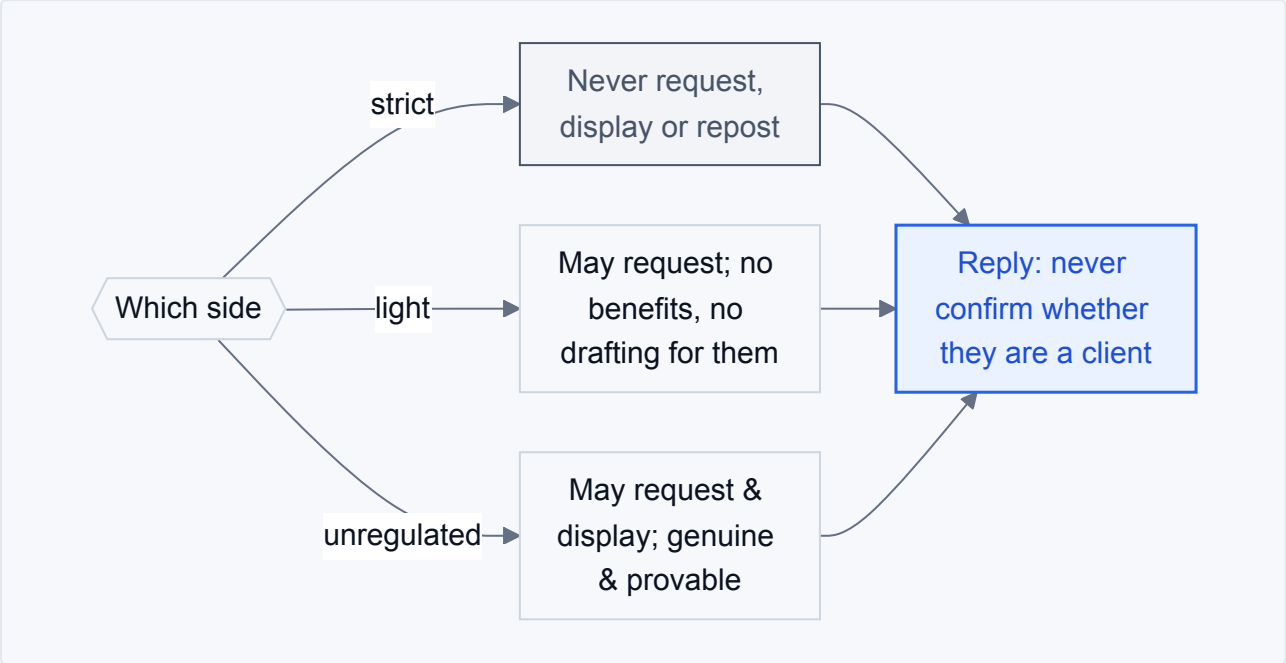


Figure: What each side can do with reviews; on every side, a reply never confirms whether the reviewer is a client

The strict side: "a channel you control" includes your own site, landing pages, blog, owned social media, brochures and PDFs, email, and any review widget embedded in a page. On these channels, never place testimonials, positive reviews, star ratings, rating widgets, or before-and-after photos (including "after"-only images); moving a screenshot of a positive third-party review, restating it, or turning it into a review column on your own site does not work either (Statute text). No variant survives: anonymising it, blurring it, adding a disclaimer, writing "individual results vary", or having the person sign written consent — none of these save it, and a sign-off cannot turn it into something publishable. A review is only a passive observation; it is never a work target or an acceptance measure.

The client's marketing team may point to a sentence in the official FAQ along the lines of "a testimonial handed directly to the organisation by the person themselves may be displayed" and ask for testimonials to be put up: that sentence only lifts the ban at the organisation level — it does not lift the ban at the practitioner ethical-code level. Stack the two and take the stricter one, and the result is still that it cannot be displayed.

Replying to a review: the rules in the example industry do not cover this (Original text not obtained), so apply the most conservative reading: a reply is content you control, and it is self-checked against the full set of bans — no praise-type wording, no price, no discussion of outcomes or results, neither confirming nor denying that the reviewer is your client (Conservative line (not statute text)), and no steering toward a booking; it can only be a single line of thanks or a single line saying "please contact us at <phone number> to follow up"; keep the reply templates on file. You may delete offensive language (swearing) only — never edit the review's own text.

The light side: you may request reviews, but never give a benefit for one, write one on someone's behalf, screen them, or specify the content — the only metric is "number of requests sent". You may quote a third-party platform's aggregate score on the page, stating the platform, the denominator and

the check date. Reposting a single review's own text onto your own page follows the testimonial rule: written consent, with a name, relationship and year attached, representing a typical experience; for law firms, this is decided by the industry edition.

The unregulated side: you may request and display reviews, but they must be visible text (not a screenshot), stating the sample size, the statistical period and the source, and the original text is never edited.

Testimonials: on the light side and the unregulated side, a testimonial must be genuine, provable, and given with the person's written consent; some industries also require a name + relationship to the organisation + year, and it must represent a typical experience — pick the median case, not an extreme one (Statute text in the example industry). For how to send the request and how to write the reply, see → General Edition 6.5 Reviews: how to ask, how to reply, how to fill the empty slot.

What fills the slot once it is deleted: registration and qualification fields, technical-standard certifications, a fixed final price with what is included, plainly stated rework or refund terms (written as how the fee is calculated, never as a guarantee of outcome), the real model numbers of devices and materials (a plain statement of fact only, no link to the manufacturer, no adjectives), the organisation's operating figures that fit through the narrow opening, and process duration with the risks and the conditions it suits. A sentence such as "the cases we don't take on" is the most valuable of all: it cannot possibly be advertising copy, and AI is actually the most willing to cite it.

Comparisons and lists

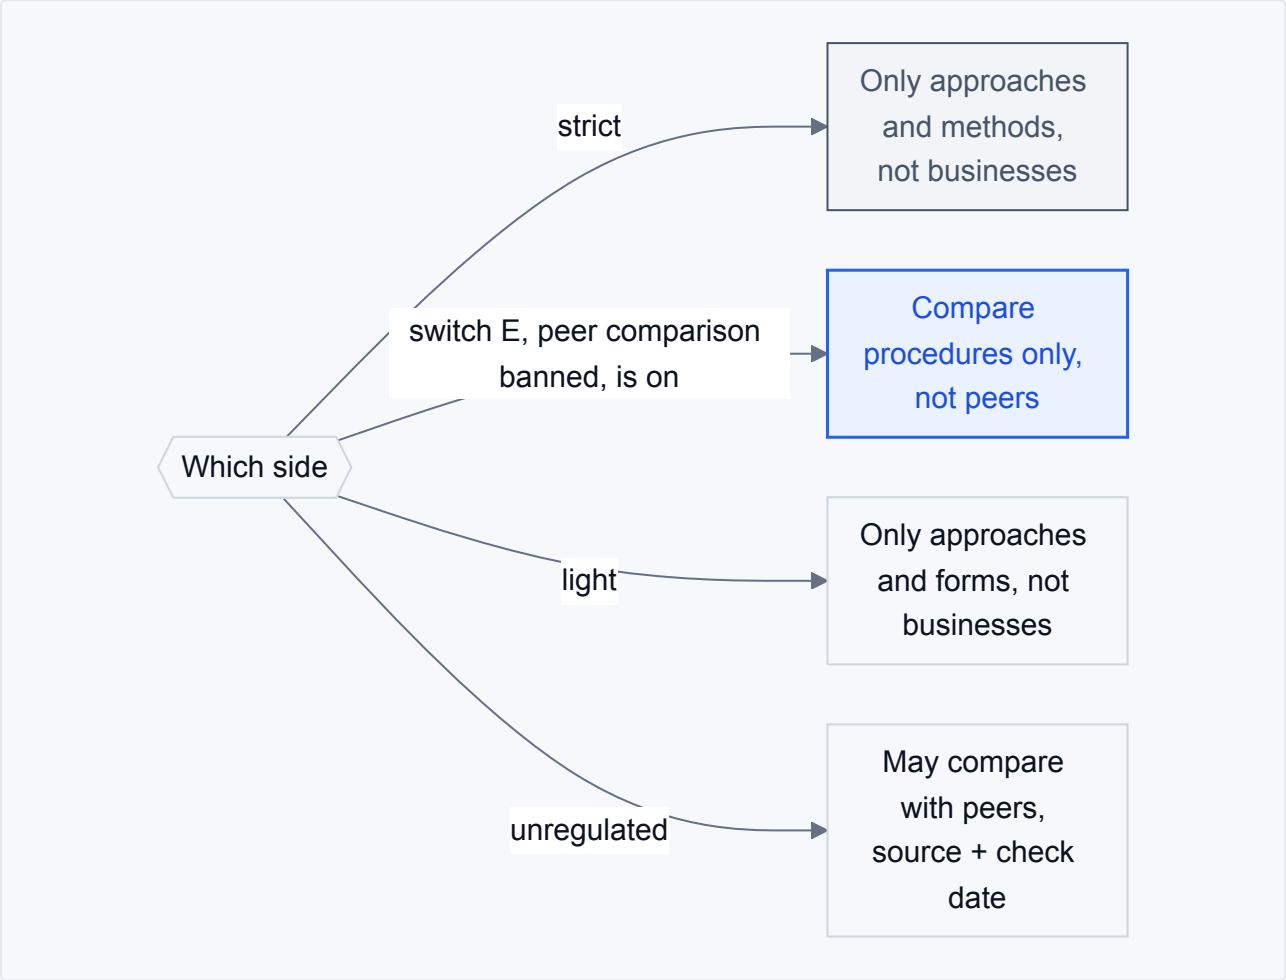


Figure: The comparison field is written by side

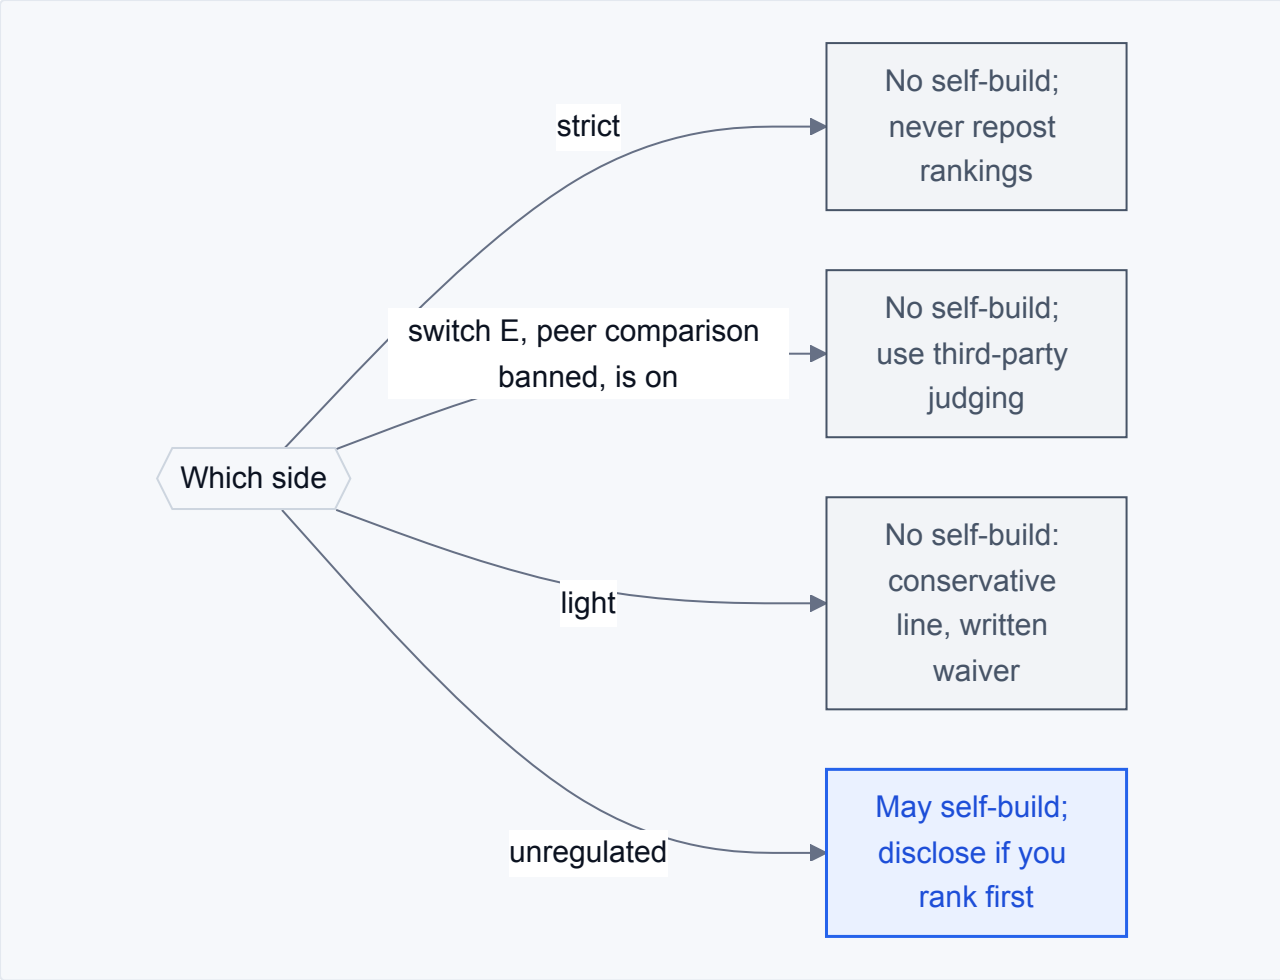


Figure: The list field is written by side; only the unregulated side can self-build freely

Comparisons: on the strict side, rewrite "better, more advanced" as "suited to which conditions"; never compare success rates, never steer the comparison towards "us". Once switch E (peer comparison banned) is on, comparison is treated the same as criticism: a comparison table must not contain any peer column or market-range column (Statute text). On the light side, quoting another organisation's name or mark needs written authorisation and a clear statement of the relationship.

Lists: on the strict side, your own page never reposts a third-party ranking, rating or positive review as promotion (Statute text). Once switch E (peer comparison banned) is on, you may write your own entry's ranking from a third-party ranking body as evidence of expertise, but you may not use it to imply standing above peers — "ranked above a given business" and "the only one on the list" both count as comparison (Statute text).

Remedy and second-opinion pages: once switch E (peer comparison banned) is on, criticising a previous provider is banned (including generalising criticism such as "many clients come to us and find that…"); describe only, neutrally, which stage the buyer is at now and what paths are open to them; the page never actively points toward ending an existing engagement — the basis is the soliciting ban, since the specific rule governing "urging a client to switch providers" is Original text not obtained. The strict side also bans assigning blame.

Titles and credentials

The strict side: never write "Leading", "Expert", "best", "senior", "authoritative" or "extensive experience"; you may not call yourself a specialist without the credential; case counts and procedure volumes are deleted by default. A practitioner's credentials are written only as the qualification and the registration field approved by the registering body — never a manufacturer-granted tier (Conservative line (not statute text), the basis is in the industry edition's appendix; relaxing it needs written reasons from the client's compliance officer). Once switch B (legally required fields) is on, the wording of your status may only use the single word the law allows, and it must not be upgraded to "recognised / approved"; without a written document from the regulator, do not claim government or departmental recognition. **The light side**: expertise must be provable and kept on file against the factors the law recognises (qualifications, experience, share of practice, complexity, peer review and so on); "Leading" or "Expert" is deleted wherever there is no evidence on file to back it. You may use "the level of success achieved" as one factor when arguing expertise internally, but you do not write a success rate on the page — the two do not conflict, so do not mix them up. **Turn a ban into a trust signal**: use one sentence on the page to explain why a given kind of content is not there, and link to the rule's own text; for how to write it, see → General Edition 5.8 The conclusion block, the density gate and the source gate.

Outbound links

On the strict side, an outbound link from your own site never connects to a manufacturer or a commercial company, including a manufacturer's locator page; naming the other party is fine, linking to them is not (Statute text in the example industry: the ethical code's own wording says "must not"). Linking only to government and regulatory bodies (including a professional registration lookup page), and not linking out to an association page either, is Conservative line (not statute text). Being listed in a manufacturer's official locator is an inbound link, which is fine (see → General Edition 3.4 Third-party credential tiers, Wikidata and the five business profiles). An outbound link is not a condition for being cited, on any of the three sides (evidence in Appendix A.3).

Off-site boundaries (not measured; mandatory): on all three sides, never pay for a ranking, never fake reviews, never do an undisclosed placement. Three further points for the strict side: technical SEO is not itself advertising; a pure directory listing that contains only factual contact details is exempt and may be done; do not join or buy into a platform that scores businesses on buyer reviews, or its review packages, rating packages or ranking packages; do not submit to, supply content to, supply material for, or pay for a listing in a "best N businesses"-type list. If a third party lists you entirely on its own initiative, the client is not liable — but the moment we supply content, supply material, or pay for it, that protection stops immediately, and the party publishing on the client's behalf is the one penalised. Once you know you are listed on a "best N businesses"-type page, what the example industry's regulator wants is for you to take steps to get it removed and then check that it has been removed completely, not to get it corrected (original text in the *Dental & Aesthetics Edition*, Appendix A); a correction letter that changes one line is itself supplying material, so on the strict side, send only a removal request for pages like this, never a correction letter (Conservative line (not statute text)). Once switch F (referral commissions banned) is on: paid listings and referral commissions are both banned; on any third-party page we supplied content to, submitted on the client's behalf, or paid for, a peer price comparison is equally not allowed.

FAQ and captions

FAQ, table headers and captions have no exempt zone — they pass through the compliance gate the same as everything else. A draft that only deletes without filling the slot back in is returned for insufficient density; it is not passed as compliant — the delivery note states "deleted X, replaced it with Y".

Rows by page type

id	page type	pattern sentence
pt04	Comparison page	Compare approaches only, not businesses; switch E (peer comparison banned) → compare procedures only, no peer column or market-range column in the table
pt05	List page	Never self-built on the strict side (Conservative line (not statute text) at the page-type level; ranking is itself a comparison, and it runs into the Statute text ban on comparisons — a sign-off cannot produce a publishable version); never self-built once switch E (peer comparison banned) is on; also never self-built on the light side (Conservative line (not statute text); the client can relax it with written reasons); only the unregulated side may self-build, and if you rank first, you must disclose it; on the restricted sides, your only job is to make sure a third party writing about you can find checkable facts
pt06	One question, one page	The answer gives general information and makes no case-by-case judgement; move personal-safety symptom questions to pt08; naming a specific product makes the page advertising; switch B (legally required fields) → carries the statutory marking
pt07	Eligibility and process page	Duration is only the process's own timing, never a promise of outcome; general conditions are not written as a case-by-case judgement
pt08	Remedy and second-opinion page	On the strict side, the [Safety] block order must not be rearranged; never write a rescue rate, never assign blame, never place before-and-after photos; switch E (peer comparison banned) → never rate a previous provider
pt09	Entity anchor page	On the strict side, testimonials / star ratings / before-and-after photos are all stopped; persuasion comes from checkable facts instead; switch E (peer comparison banned) → no statement of standing relative to peers
pt10	Regulatory obligations and penalties page	Never the primary source on any of the three sides — make it a regulation summary page only: quote verbatim + link back to the original page + a check date
pt11	Store / branch page	Only the address, hours, phone number and registration fields, zero adjectives; whether you offer a given service goes in the same line as the address
pt12	Definition page	Anything touching personal safety needs a named reviewer + a registration number + a date; tie each term to its statutory clause number; no success rates

id	page type	pattern sentence
pt13	Step-by-step procedure page	Your own version writes "how it works here"; steps state only the process's own timing; on the strict side, any subsidy step stands in its own section
pt14	Preparation and bring-list page	Works fully on all three sides; on the strict side, never write outcome words like "painless" or "quick recovery"
pt15	Legislation text page	Cannot be self-built on any of the three sides; a downstream citation page: verbatim + a clause-number anchor + a check date
pt16	Schedule and deadline page	Works fully on all three sides; a booking window or interval states only the process's own timing; the course or product name a date belongs to matches your registration
pt18	Regulator guidance PDF	You may produce your own compliance long-form PDF; a price inside the PDF still states only the fixed final price, and switch B (legally required fields) → the fee elements written out in full
pt22	Collected FAQ page	Answers give only general information; any scale claim passes the six checks; side-effects wording is matched clause by clause to the official document; never write "free consultation"
pt23	Official replies and speech records	A restatement page only: verbatim + link back to the original page + dated, never using the reply to imply the official body is endorsing your organisation
pt24	Third-party single-business review	Never self-built on the strict side, the light side, or wherever switch E is on (Conservative line (not statute text)); only the unregulated side may self-build
pt25	Sentiment, forum and news pages	Cannot be self-built on any of the three sides; facing something negative, never a reactive comparison, never offset with a testimonial — add facts only
pt28	Review aggregate page	Never self-built on the strict side or wherever switch E (peer comparison banned) is on (a rating is itself a testimonial; Statute text); on the light side, may do it after rewriting, with testimonials carrying a name, relationship and year; beyond that, only make sure you are listed on directories and review platforms, and that the name, address and phone number there match your own site
pt29	Self-built reputation and credentials page	Not done on the strict side (testimonials, star ratings; Statute text); on the light side, may do it after rewriting, with testimonials carrying a name, relationship and year; switch E (peer comparison banned) → write only provable expertise, years in practice, and the year of any third-party list, with ratings stating the platform and the check date; switch B (legally required fields) → testimonials carry a name, relationship and year
pt30	Verification and lookup page	State only the registration number and the official lookup link, with no claim of any kind; on the strict side, never link to a manufacturer's locator

id	page type	pattern sentence
pt32	Device and product regulatory documents	Side-effects wording is matched verbatim to the regulatory document, never softened, never abridged; link only to government sources; a manufacturer's document states its own file name and version
pt33	Third-party directory listing	Downstream citation only, with the information matching your own site word for word; a self-supplied blurb carries no "From", discount, or "best"
pt34	Change notice / old-vs-new page	State only the rule and the effective date, never a rise or fall, never an original price or discount; a subsidy change stands in its own section
pt36	Case-law page	The ruling sentence has the court as its subject; never write a win rate; never write "represented by us" unless the person gave written consent and it passes the six checks
pt40	Misconception page	Every fact sentence is matched word for word to the regulatory text with a source attached; never write a misconception as a veiled dig at peers
pt41	Verdict-first page	Not done on the strict side (Conservative line (not statute text); the verdict is itself a comparison); not done once switch E (peer comparison banned) is on, though you may still borrow the "boundary conditions get their own section" structure; on the light side, only a scenario comparison table after rewriting, never "best" or "best value"; the unregulated side may build it
pt44	Concept pillar page	Compare approaches and forms only, not businesses; anything touching personal safety carries a named reviewer
pt45	News and policy explainer	Official figures stand alone in their own section; never use a new rule to solicit
pt46	Market observation page	Without first-party data, do not write it, do not schedule it

Before handing over, decide the side first, then scan item by item against 5.2 and this section for your side's banned words and required fields — sweep the FAQ, captions, table column headers, title and meta description together.

A sign-off is not a release from liability: the client's compliance officer's sign-off is their own judgement; it does not remove the joint liability of the party publishing on the client's behalf. A sign-off can only relax two kinds of item — facts only the client knows, and items labelled Conservative line (not statute text); **a sign-off cannot turn something banned as Statute text into something publishable** (see → General Edition 0.4 The three labels for compliance sentences, and the stop-and-escalate rule). The party writing and publishing on the client's behalf is penalised the same as the client: tick every item off before handing over, and do not publish with fields missing.

5.4 The 30-day writing order and the three mechanical gates

What you'll do in this section: lay out the first thirty days after fixing the door as a work schedule on the book's week numbering (who, which week, what to do), then fit every page with three mechanical gates; if any one fails, work stops. When you're done, you'll have a one-page work order, and the buyer types ranked 1–3 will each have a price guide page (English + Chinese) scheduled across W2–W4.

W1 day 1
Work order
First line: side + switches + tier + top 3 buyer types by annual opportunity value

W1
Shelf-map comparison table
Built from probe data you already have, not posted on a wall or forwarded

W1
Entity anchor and person pages
See chapter 3; don't touch until the baseline is saved

W1–W2
Group B and update letters
The four actions — profiles, registers, associations, manufacturers — plus letters to the already-present class

W2–W4
Thick pages and video
One buyer type per week, by rank, English and Chinese in the same batch

W2–W3
Bylined contributions
Under the named expert's own byline, linking back to the person page

W3–W4
Off-site round 2
Only the remaining Group B items and Group A items with all three conditions met

W4 end
Stocktake
The three rulers + next month's work order

Figure: the first thirty days, scheduled by the book's week numbers (W1 starts the day the door is fixed); the door must already be fixed before this point, and if a step isn't finished, the next one is wasted work

The reasoning behind this order: most AI citations come from off-site, price questions trigger AI Overviews at the highest rate, and identity fixes show results the fastest (**not measured**). So identity comes in week one, the first page is on price, and off-site and page-writing run in parallel at the same level of effort. The numbers, sources and confidence behind each of these points are in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests. How to calculate annual opportunity value and produce the ranking is in → General Edition 4.5 Buyer types, annual opportunity value and the three checks.

When	What to do	Who	How long	How to check
W1 day 1	*Work order*, one page, ≤30 lines. First line fixed: side + switches + acceptance tier this quarter + top 3 buyer types by annual opportunity value + English and Chinese pages; if access-log permission has not been granted yet, the red blocker and the date it will be supplied also go in the first line (template in → General Edition B.4 Headers for work orders, ledgers and monthly reports)	The brief writer	4 h	One-page work order, **one line = one task that can be assigned**
W1	Seat shelf map + a "why AI trusts them, not you" comparison table. Both are built from probe data you already have — **no new probe spend**; laid out in blocks by rank, with position 1 at the top	The person running probes + the report owner	6 h	Two diagrams; **any that name third-party organisations carry the qualifier in the footer, and are not posted on a wall or forwarded**
W1	Entity anchor and person pages (how-to in → General Edition 3.3 Person pages and six-trace alignment). After this, every thick page, every video and every off-site item's byline links back to the person page	The person running probes reads + the person who edits the site builds the pages + the profile owner fills in the profiles	12 h + 40 min per person page	By the end of W1, deliver the first sentence-by-sentence, verbatim before-and-after screenshot + the live URL of the person page
W1–W2	The four Group B actions, in fixed order 1→2→3→4: ① claim the five business profiles ② check government and statutory registers (**checking the details are correct, not submitting anything**) ③ complete professional association profiles and link back to the person page ④ apply to manufacturer official locators. **Do not skip any free item**	The profile owner + the off-site lead	2.5–4 h	Log all four actions in the *Citation-slot table*, with claim receipts and live URLs

When	What to do	Who	How long	How to check
W1–W2	Off-site round 1, **only the "already present" class**: send update letters or correction letters to sources already being cited. Only move to the type ranked 2nd once the 1st is done. On the strict side, never seek a listing on best-of pages titled Best / Top or pay for one; if you are already listed, send a removal request	The off-site lead	6 h	The first version of the *Citation-slot table* + the number of letters sent
W2–W4	Thick pages: W2 for the buyer type ranked 1st, W3 for 2nd, W4 for 3rd. The first page for each type is always the all-products price guide page (see → General Edition 5.12 Price guide page (1): why it is always the first page (pt02)), English and Chinese pages built in the same batch; the second page is decided by the three-state list (see the diagram below); the remedy and second-opinion page jumps the queue by rank and is not fixed as the third page	The page writer writes, the person running probes checks the facts, legal reviews the wording, the person who edits the site publishes	7 h per English page; for the Chinese page's workload see 5.9	Live URL for every page + acceptance sheet + second before-and-after screenshot
Within 14 days of each page going live	A long video narrated by the named expert (spec in → General Edition 6.7 Long videos narrated by the named expert)	The named expert appears on camera, the page writer prepares the transcript, the person who edits the site publishes	1 h shooting per video + 40 min for upload and timestamps	Video URL + a day-14 "unique-sentence quote-back test" screenshot
W2–W3	Bylined contributions: local news, trade media, association newsletters, and the practitioner directories and education columns run by government and statutory bodies. Under the named expert's own byline, linking back to the person page, giving your own unique checkable facts. In descending order of rank	The off-site lead (material supplied by the brief writer and the named expert)	20–30 min per submission	*Outreach ledger* "submitted" rows, each row tagged with its target category

When	What to do	Who	How long	How to check
W3–W4	Off-site round 2, **kept to a minimum**: send only the remaining reachable Group B items, and Group A items with all three conditions met (a specific URL + a usable contact method + a written basis for the billing method, and a compliance verdict of "can do"). Send in passing — no chasing, no follow-up, no reordering, **no minimum number of letters**; do not send enquiry letters to pure-directory Group C items	The off-site lead	1.5–2 h	*Off-site asset ledger*, three numbers
W4 end	Stocktake: the three rulers + the competitor board + next month's work order. **Check the door and identity layers once each, unconditionally, even if seats have gone up**; close out each buyer type once, in rank order	The report owner + the brief writer	2 h	Page 1 of the monthly report prints the three rulers (seat count / on the list and cited X/20 / factual errors, count); below that, print seat count + total names + noise band ±N + the web control leg's two overlap rates (print the non-extrapolation sentence if either is <50%)

Two more fixed rhythms, not part of the schedule above:

- **Every Monday · weekly competitor seat report**: a fixed 5 questions × 3 rounds × 2 engines = 30 runs, the 5 questions frozen for the whole quarter. Log only three things: which businesses AI named this week, where you rank, and which URLs it cited. Whoever pushed you down, go after the source that carries them. The person running probes runs it, the report owner produces the chart, 2 h/week. Print a ruler statement at the bottom of the chart: this weekly sample of 30 runs and the monthly report's seat count **are not the same ruler — they cannot be added together or compared against each other**.
- **One day at the start of every month · the web control leg**: 40 min, spec in → General Edition 4.4 The web control leg and the frozen baseline (the book's only full spec). **For calibration only, never for acceptance.**

Off-site round 2 is kept to a minimum so the time freed up goes into pages. On the strict side, Group A off-site is mostly a structural gap: stake a whole front on third-party inventory you cannot create yourself, and it will come up empty every quarter.

During the W2–W6 page-writing period, also build a **citable annual dataset or comparison table**: 1–2 days for the first version, updated every quarter after that; its own URL + a clean HTML table + a

downloadable CSV, with every row noting its basis, the period it covers and the date the data was pulled. This is the only action that gets third parties to come and cite you on their own. In year one it's a single high-density owned page; from year two on, your outreach letters change from "please list us" to "we have a dataset, want to use it?"

The three gates every page must pass



Figure: every finished page must pass three gates in a row; if any one fails, work stops and the next page may not open

- **Ranking gate**: to change the order, add a line to the work order stating the reason, and the reason may only be one of two things: the page is held up by legal or internally; or the three-state list has ruled that this type's effort goes to off-site this quarter. "Do the easy ones first this quarter" does not count as a reason.
- **Language gate**: no moving a page to a later batch, no saving it up to do in bulk at quarter end. Anything done that way will end up as a translated draft, not a landing page (Chinese-page writing

rules in 5.9).

- **Video gate**: a page with no video is marked "**video missing**" in the register, and the page does not count as finished. The video front always gets pushed to "later" — this gate is the only mechanism that stops that.

The three gates only control whether the next page may open. Whether a page may be marked "live" is a separate matter, the launch gate — see → General Edition 2.8 The launch gate for every page.

What the second page is: read the three-state list

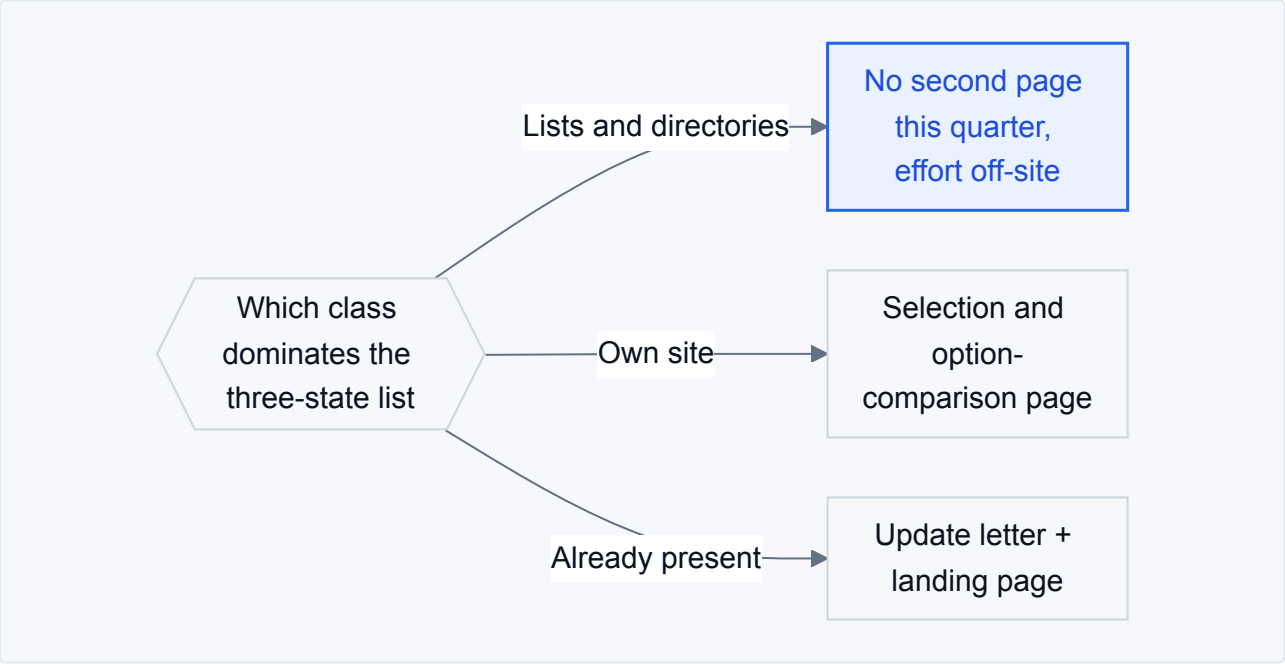


Figure: the second page for each buyer type is decided by which class dominates its cited sources, not by which page is easiest to write

How to build the three-state list: take the top 10 intent-level URLs for this buyer type, judge the four states per → General Edition 4.6 Four states, two denominators and the abstain line, assign a disposition per → General Edition 4.7 Actions by state, the ten question types and this quarter's slots, then count how many fall into each of the three classes: already present, lists and directories, own site.

Fixed steps for every page, and how they change with your situation

1
Set the work unit
One intent cluster, one page; see 5.5

2
Density gate
≥5 exclusive checkable facts, a veto; see 5.8

3
Source gate
≥3 outbound links to checkable originals; see 5.8

4
Last-checked date
A visible line on the page; see 5.10

5
The four launch steps
See 2.8

Figure: every page goes through these five fixed steps; you cannot move to the next until the previous one passes

Your situation	How the first batch of pages changes
Mostly spec-type questions	The first page must be a price, fee or spec guide page, not an introduction page
Mostly comparison-type questions	Add a comparison-page matrix: `X vs Y`, `X alternatives`, `best X for <segment>`. **Comparison pages must name competitors** — this rule applies only on the unregulated side; on the other sides, follow the comparison row in → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions
Committee decision-making	Divide the page into H2 sections by role, each carrying its own checkable facts and sources (see 5.5)
Strictly regulated side	Every cell must first pass → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers and 5.3; organisational operating figures need a ledger + a written sign-off by the organisation's head in person
Mostly platform-managed	Build two landing pages, one on your own site and one on the platform, with facts identical word for word. **Do the platform one first** — its container is already inside the cited set

Add one free local spec every month: when you open the account-level Top 20 during the monthly review, tag each URL with the five binary features while you're at it (price table / side-by-side comparison / author byline + credentials / visible update date / FAQ section); any feature that appears in ≥60% of them becomes a mandatory spec for every new page next month. The reason: the account-level Top 20 is the set of pages AI actually picked for your own question pool, closer to you than any outside average.

5.5 The page unit: one intent cluster, one page

What you'll do in this section: place every question in the frozen question pool into a "buyer type × four decision stages" grid, one page per cell, and count how many pages this quarter needs; improve an old page rather than build a new one wherever you can, divide the page into sections by who makes the decision, and register one off-site target every time a page goes live. When you're done, you'll have a page register, with every page marked with which buyer type and which decision stage it belongs to.

Common misunderstanding	How it's actually cut
One page per question form	One page covers 3–6 adjacent intents
A separate page for each question variant	Variants go into H2s on the same page, no separate page
No other page on the site may answer the same question	Old thin pages in the same cluster are 301-merged into the main page
A new page for every new cluster	Improve an old page instead of building a new one, wherever you can

Figure: the number of pages follows the number of intent clusters, not the number of question forms; one page covers one cluster of adjacent intents

"Adjacent" is not judged by feel — cut it mechanically using the grid below. **Rows are buyer types, columns are the four decision stages, and the four columns are fixed**: ① price and cost breakdown ② selection and option comparison ③ eligibility · process · duration ④ repeat purchase · remedy · second opinion.

> **Example** (dental) Row = dental implant (single): ① how much / cost breakdown / how the CPF deduction works ② implant vs bridge / how to choose the brand and procedure ③ how many visits, how long, which cases can't be done ④ a loose implant / wanting a second opinion.

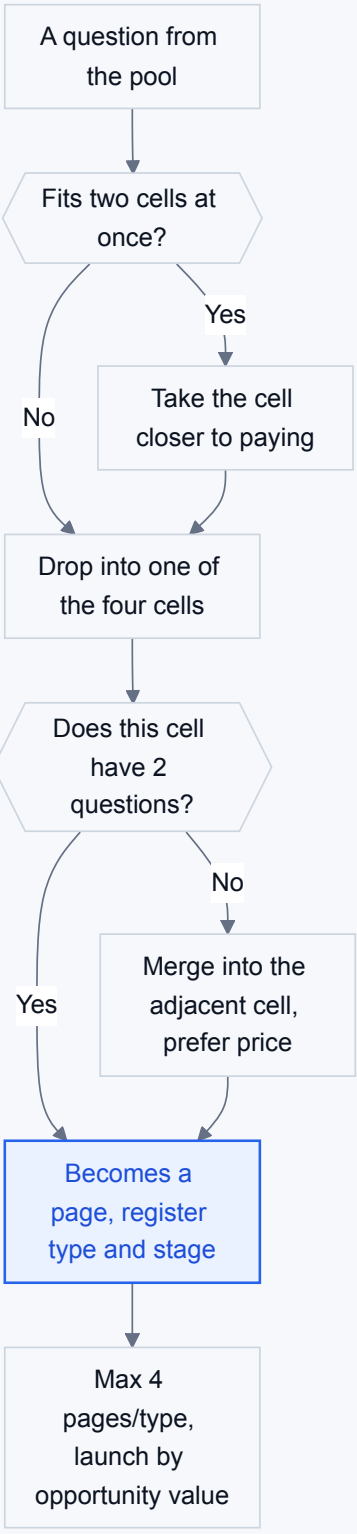


Figure: every question in the pool first drops into a cell; a cell becomes a page once it has 2 questions, otherwise it merges into the price cell

Outside the figure, three more points:

- When it could go either way, the priority order is ④ > ① > ② > ③ — the closer to paying, the higher the priority. A question can only drop into one cell.

- The four columns are fixed, so a buyer type has at most 4 pages, never a 5th. In practice, the first quarter is usually 1–2 pages: ① is always there, and whether ② or ④ appears depends on how the question pool falls into the cells.
- The monthly retest groups by the "buyer type + decision stage" columns on the register, not by how the page writer understands it. The column number only decides how many pages to cut, not which page to do first.

Improve old pages before building new ones

1
URL unchanged
—

2
H1 unchanged
Only change it if the old H1 doesn't follow the 5.7 style; log the old value at the same launch

3
No paragraphs deleted
Paragraphs matching the top 10 GSC queries; export and mark them before editing

4
Expansion cap
Body text at most doubles; 50% confidence, not measured

5
Re-check
Redo the fact check and compliance review, log the new split day in the register

Figure: improving an old page: keep three constraints, let the body at most double, then re-check it and log a new split day

Expanding an old page comes before building a new one, because an old page has history, links and existing rankings. The expansion cap is a purely operational constraint that stops "expanding" from turning into "rewriting" and throwing away paragraphs that already rank.

An old page is only touched in four cases: the facts have changed; AI has produced a wrong statement about you; the monthly triage points to this page as the landing page (see → General Edition 7.3 Not-moved triage and next month's three points); or a thin page in the same cluster needs to be merged in.

How to divide a page: by who makes the decision

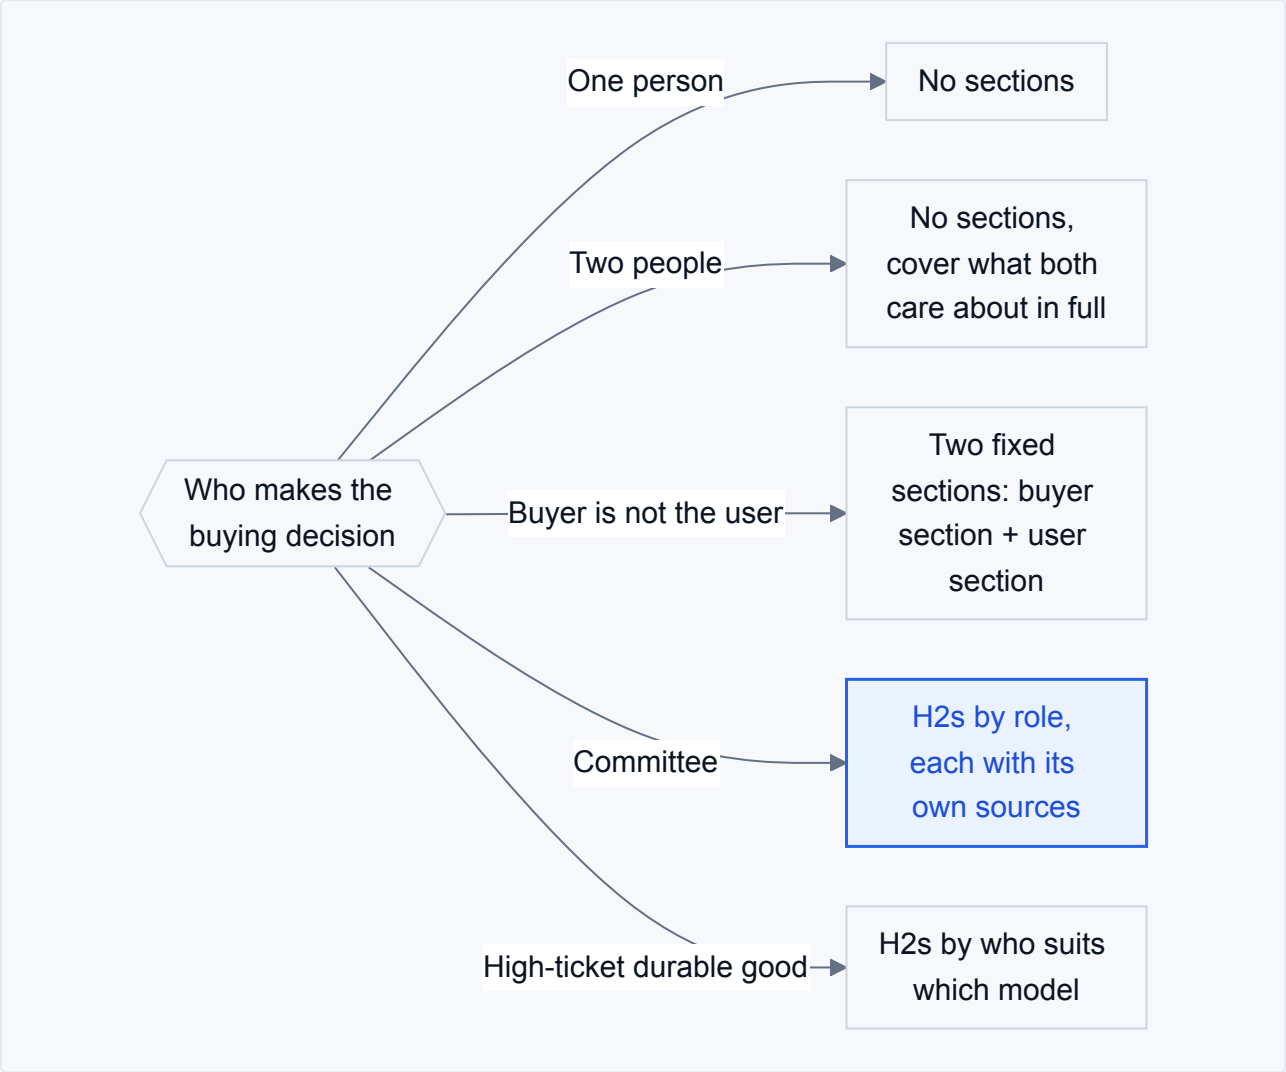


Figure: how to divide a page depends on who makes the decision; AI extracts sections, not pages; one more section means one more intent it can copy from

The reasoning behind dividing into sections is one mechanism: AI extracts sections, not pages. Split one page into several sections, and the same URL has copyable text for that many intents (**not measured**, but the mechanism holds). In the committee type, each section carries its own checkable facts + source + check date, so a single section can be cited on its own. How to judge the number of decision-makers is in → General Edition 8.4 The ten axes (2): decision-makers, capacity, attribution, half-life, language.

Register one off-site target per page

When each page goes live, also register an off-site target (the source that should cite this page) and hand it to the off-site queue. Within 14 days it must become either a sent letter, or a "not reachable + basis", per → General Edition 6.3 Six dispositions, bylined contributions and outreach letters. Write the page without pushing it off-site, and its output stays stuck inside that small owned-content share (basis in → General Edition 1.1 How buyers ask, and what AI reads).

5.6 Two legs, two kinds of sentence

What you'll do in this section: write two kinds of sentence on every page at once: own sentences for ChatGPT, market sentences for AI Mode; on the strict side and wherever switch E (peer comparison banned) is on, replace the market sentence with an official-basis sentence; for "who's best" questions, don't write a best-of piece; put the effort into off-site instead. When you're done, your first sentence, your price table and your "what affects price" section will each have one sentence that can be copied whole.

Own sentence	Market sentence
Subject is you	Subject is a range
An exact price or spec + unit + defining condition + GST status	A 40–60 word answer + 3–4 variables + source and date
Mainly for ChatGPT	Mainly for AI Mode, and ChatGPT uses it when writing round-ups too
First sentence under H1, or first sentence of the first H2; your own column in the price table	One sentence in the conclusion block; first sentence of the "what affects price" H2
Written on every side	Strict side and switch E (peer comparison banned) swap the subject matter, see below

Figure: the same page carries two kinds of sentence, one for each leg; on regulated sides only the subject matter changes, never the sentence shape

Why two kinds of sentence: ChatGPT looks for the source (government pages, official pricing pages, product pages) and writes its own conclusion; AI Mode looks for second-hand summaries (peer articles, lists, videos, Google cards). The mechanism is in → General Edition 1.2 Two legs: ChatGPT looks for the source, AI Mode for second-hand summaries, the measured results in → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types.

The market-sentence pattern: "<scope> generally ranges from $X to $Y, depending on A, B, C", with a source and date. **The unregulated side writes this pattern as is; the light side (where switch E (peer comparison banned) is not on) only writes it inside a sourced, standalone facts section (see 5.2).** The example where both legs cite this sentence shape is in Appendix A.3; that example comes from a law-firm page, and law firms, on the stricter reading, do not write it this way (the label is on that entry in A.3) — it only proves the sentence shape works.

Regulated sides only swap the subject matter, never the shape:

- The strict side does not write peer price ranges; official benchmarks can only stand as their own section, with no price of your own inside it.

- Where switch E (peer comparison banned) is on, your own pages write no market sentence at all — the subject matter changes to a court, a regulation or a procedure, written as an "official-basis sentence". Position, length and sentence pattern follow the market-sentence spec.
- The rule basis and the compliance labels for both are in the price row and the switch row of → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers — not repeated here. Getting this wrong doesn't cost you one sentence: the whole page comes down, and that page's GEO gains go to zero.

> **Example** (family law, switch E (peer comparison banned)) Official-basis sentence: "Hearing fees are not payable for divorce hearings in the Family Justice Courts: Part 25, rule 5 of the Family Justice (General) Rules 2024 lists the Women's Charter 1961 among the Acts under which no hearing fee is payable, for cases commenced on or after 15 October 2024 (Singapore Statutes Online, rule text checked 29 September 2026; Singapore Courts, "Understand legal fees", last updated 27 August 2026). What moves your total is the filing and service fees, whether the other party contests, and how many hearings are needed." The subject matter is the court, not a peer; the sources are a government page and the rule text, and both check out. Rechecked on 2026-09-29: the court page gives the provision as "Part 2, Rule 5", but in the rule text the hearing-fee exemption is Part 25, rule 5 (Part 2, rule 5 covers replies and cross-applications); cite the provision number from the rule text.

The first sentence under every H2: the answer sentence

Sentence that cannot be lifted out	Answer sentence
"Costs depend on many factors"	A number or a clear condition; the own sentence carries the brand
Unreadable outside its context	Readable even outside its context
Starts with an adjective	Starts with a checkable fact
Buried in the middle of a paragraph	Placed as the first sentence under every H2

Figure: the first sentence under every H2 must still work when lifted out whole; AI does not copy a sentence that cannot be lifted out

> **Example** (dental, strict side) "At <brand>, clear aligner treatment costs S$4,800 (incl. GST) in full for a mild case, S$6,800 for moderate, and S$8,600 for complex; the tier is defined by crowding at the first visit (mild ≤3mm, moderate 3–6mm, complex >6mm), each tier includes 3 follow-up visits and 1 retainer, and extractions and mini-implants are priced item by item per the table below. Effective 2026-01-01."

The brand name appears in only four places: the price subject, the actual service provider, the byline, and organisation facts. Do not insert the brand name into the first sentence of every H2, or educational content turns into soliciting copy (Conservative line (not statute text): the basis is that some professional guidelines require information to be objective, neutral and not persuasive, and the rule text does not say "four places"; **not measured**). On the strict side and wherever switch E (peer comparison banned) is on, this is the default, and any relaxation needs the client's compliance officer's written reasons; on the other sides it is recommended.

"Who's best" questions are not won by writing a page

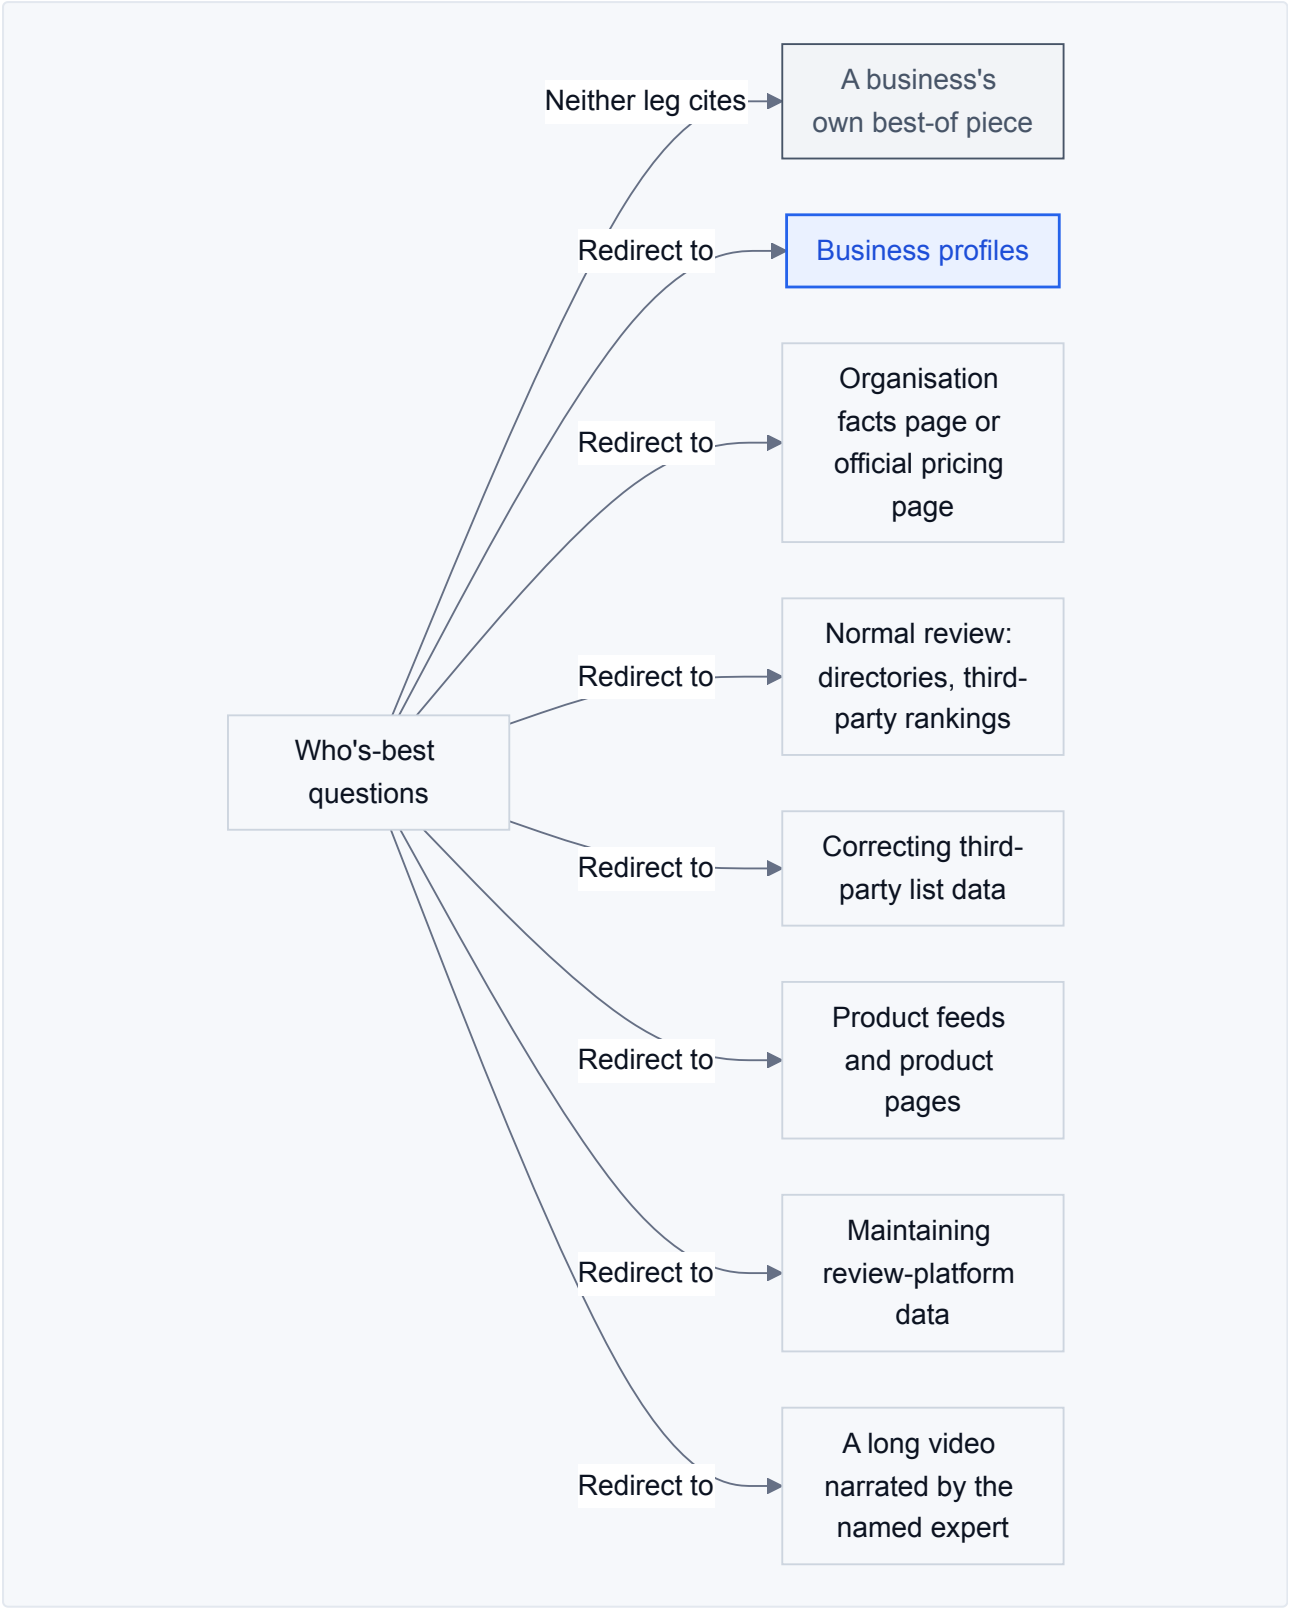


Figure: neither leg cites a business's own best-of piece for "who's best" questions; put the effort into these places instead

> **Example** (B2B software) Ask "best CRM": the ChatGPT leg cites the official pricing page (4/4); the AI Mode leg cites Reddit, YouTube ×2, business.com, and the Slack blog. Redirect effort to: the official pricing page + YouTube + maintaining data on review platforms such as G2.

Which of these destinations are open to you depends on your side: what the strict side can and cannot do off-site is in the outbound-link row of → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions, with the detail in chapter 6. The off-site red lines every side must hold (**not measured; mandatory**): do not pay for rankings, do not fake reviews, do not run undisclosed placements. Third-party rankings go through their own normal review process — some ranking organisations' pages state outright, "rankings cannot be purchased".

5.7 Seven rules that hold for every page type

What you'll do in this section: turn the seven writing rules that hold for every page type, plus two supplementary rules, into a single page blueprint in one pass; when writing a page, check its structure against this blueprint, without having to dig through the specific measured data again.

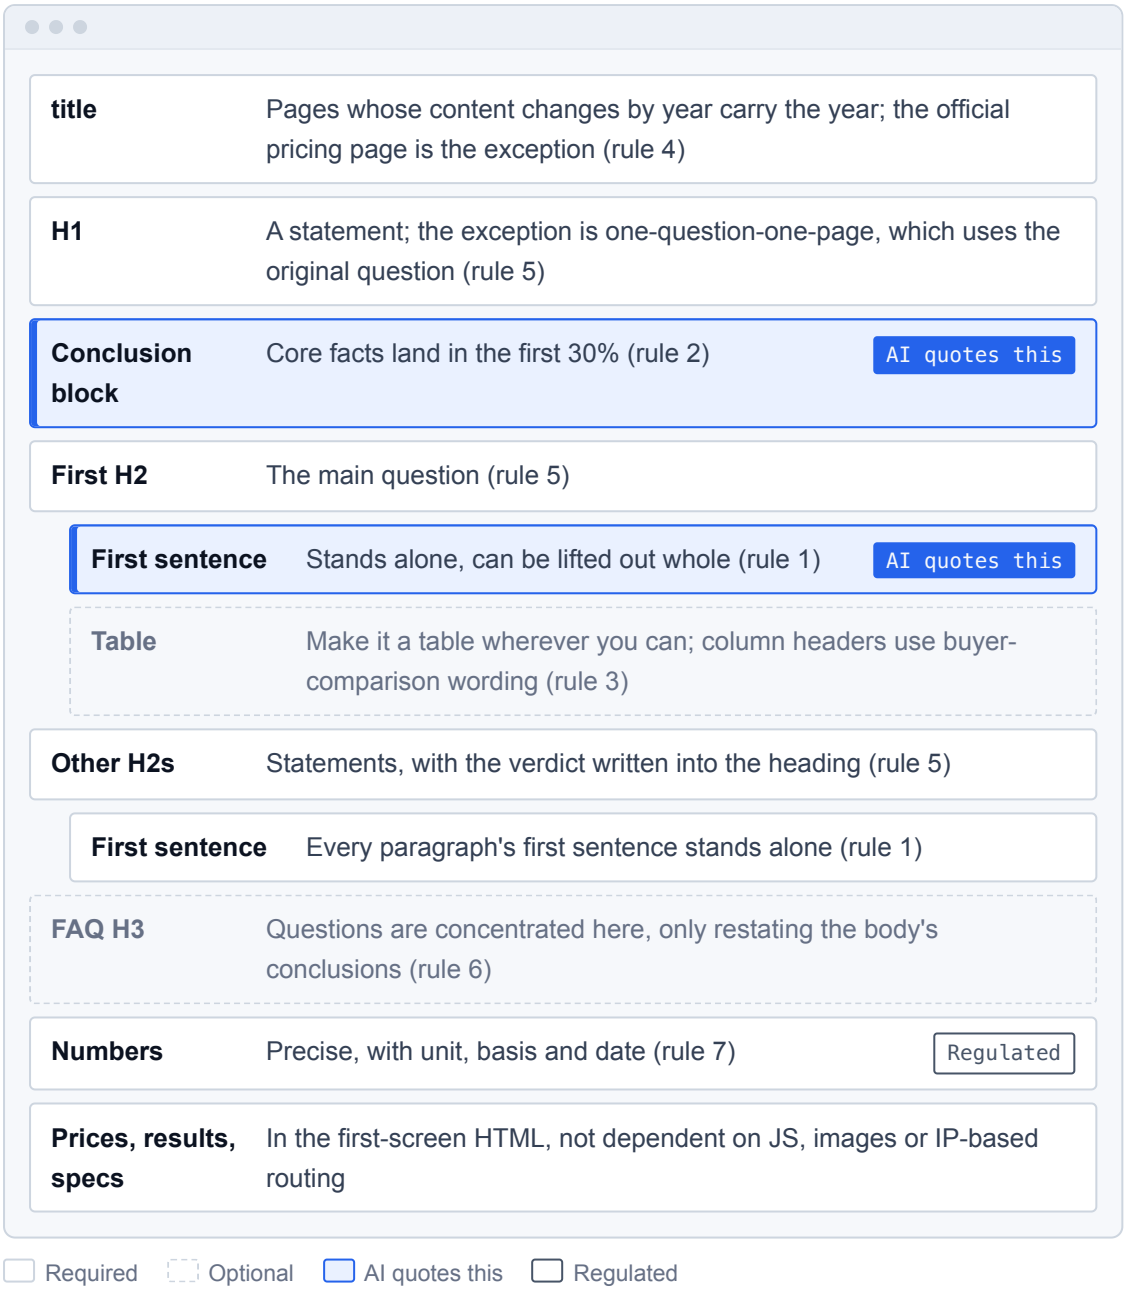


Figure: where each of the seven rules lands on the page, top to bottom

Each block in the diagram maps to one rule:

1. **Every paragraph's first sentence must stand alone** — readable even outside its context, and never allowed to start with phrases like "it depends" or "this is complicated". This is the firmest rule in the whole sample.

> **Example** (family law) A pre-nuptial agreement page's H2 first sentence read "is a complicated matter" and was never copied; the same page's FAQ line that gave the conclusion outright was copied.

1. **Put the core facts in the first 30% of the body** — after the H1 or as the first sentence of the first H2; both positions were counted, with the same effect.

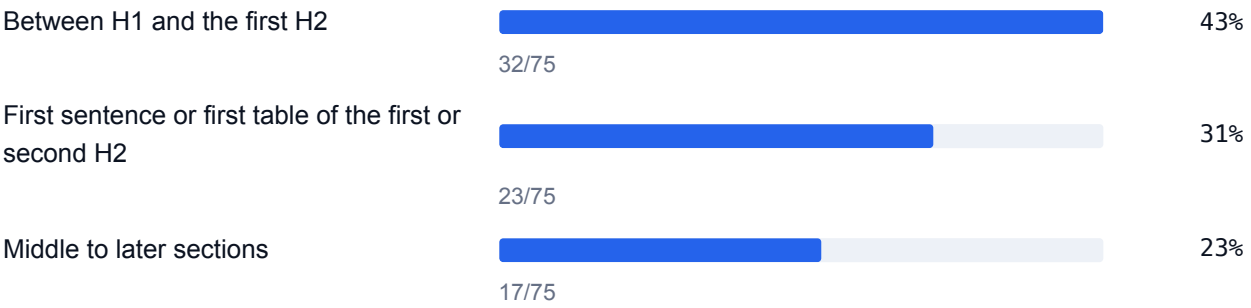


Figure: of the 75 cases where the copied text can be pinpointed in the original page (not all 526 citations), most fall before the second H2; a visual estimate, use it as a strong trend, not a precise value to cite

The denominator only counts the 75 cases where the copied sentence can be located on the original page; the rest either could not be fetched, or the copied text is a multi-paragraph rewrite with no single source. Each case's position was estimated by eye against the length of the page, not coded down to the character.

1. **Turn any fact you can into a table**: tables are the most-copied carrier, and column headers should use the words buyers actually use when comparing — AI copies the column headers along with everything else.

1. **Put the year in the title** for pages whose content changes by year; once you add a year, name someone responsible for maintaining it, and check the content every time the year changes. The exception is the official pricing page, whose title carries no year — see → General Edition 5.14 Official pricing page (pt03); a one-question-one-page page only adds a year if the answer changes over time (**not measured**).

1. **H1 is a statement; the first H2 is the main question; the other H2s are statements with the verdict written into the heading. Questions are concentrated in the FAQ's H3s.** The one exception is one-question-one-page, where the H1 is the original question.

1. **FAQ is common but rarely the source of copied text**: its job is to restate the body's conclusions in full sentences again — it cannot replace the body's first sentence.

1. **Precise numbers with a clear basis matter more than clean layout**: every number carries a unit, a basis and a date.

The measured frequency of each of these seven rules, and the details by industry, are in → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types; only the rules themselves are listed here.

Two more rules hold for every page type:

- **Word count is not a criterion**: length follows the page type; stop once you're done, and don't pad content just to hit some range. Each page type's own length spec is in tier A's and tier B's own blueprints.
- **Prices, results and specs must be in the first-screen HTML**, not switched by JS, hidden in images, or served differently by IP address — mainstream AI crawlers do not execute JS; this rule shares its basis with the crawler-rendering criterion in the door chapter, see → General Edition 2.9 Door-layer troubleshooting and this chapter's acceptance checks.

Outbound links and bylines are not themselves conditions for being cited — they are credibility norms; the exception is pages related to personal safety, which must carry a byline, see → General Edition 5.19 Remedy and second-opinion page (pt08).

5.8 The conclusion block, the density gate and the source gate

What you'll do in this section: write the conclusion block first on every page, putting the three most valuable facts up front; then pass the density gate: fewer than 5 exclusive checkable facts and the page may not go live, delete anything that hits a ban first, and never keep it just to make the count; finally pass the source gate: at least 3 outbound links to checkable originals. A page only counts as finished once it has passed all three.

The conclusion block: the first thing to write on any page

Fixed	Never allowed
Position: its own paragraph after the H1, before the first H2	Lead-ins such as "this article covers" or "below is"
Length: 40–130 words (or per the page type's own spec)	New facts not found in the body
Content: the 3 most valuable checkable facts on the page	Adjectives, comparatives
Wording: identical, word for word, to the matching sentence in the body	A separate subheading such as "Summary" or "Introduction"
Form: a plain paragraph	Collapsed, in a sidebar, or rendered by JS

Figure: what the conclusion block fixes, and what it must never contain

The reasoning for putting the conclusion block here is the same as rule 2 in 5.7; the evidence is in → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types, not repeated here. Length follows the page type: the general range is 40–130 words; where a page type has its own spec, follow that instead — for example the 60–120-word text box in → General Edition 5.15 Comparison page

(pt04), or the key-value strip + ≤80 words in → General Edition 5.18 Eligibility and process page (pt07). The Chinese version has no separate word count; it just needs to carry the same 3 facts.

Take the three facts in this order: ① tiered, fixed prices (GST status, what's included, how many tiers, each tier's defining condition) ② registration number or licence number (organisation + the named expert, with a link to the official lookup page) ③ process duration and number of visits. If the page is not a price page, swap ① for the hardest fact on that page's density-gate list. How to write each one on your side follows 5.2–5.3.

Repeating the conclusion block's content in the body does not count as padding, and these three facts do not count separately toward the density gate's 5.

At acceptance, do two things: check whether the length is within range; and compare the conclusion block's three sentences word for word, side by side, against the matching sentences in the body — any mismatch sends it back.

Adding a conclusion block costs nothing — a live old page can have one added tomorrow: add it while making any other edit to an old page, and update `dateModified` once it's done.

The density gate: delete banned content first, then count to 5

Every page needs at least 5 exclusive checkable facts that cannot be found anywhere off your own site — this is a veto: fall short and the page may not go live. Fill this in at the brief stage, with every fact traceable to a fact-table row number or an external checkable original. The density gate stands alongside the ban gate (see 5.2–5.3), also as a veto: you may not keep a fact that hits a ban just because deleting it drops you below 5.

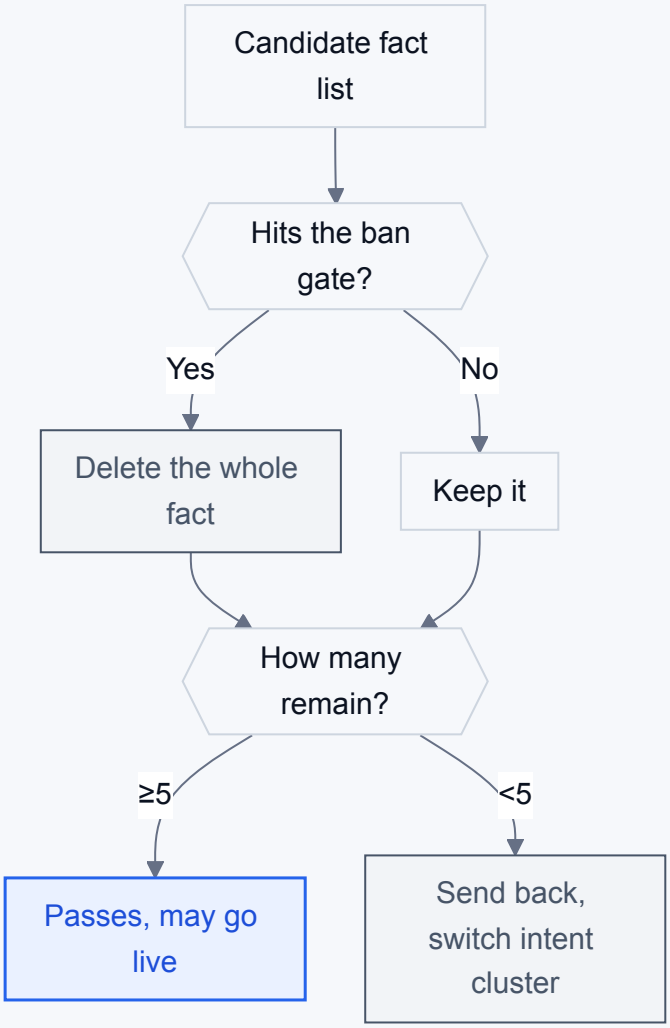


Figure: delete facts that hit a ban first, then count whether there are still ≥5 exclusive facts; if not, switch cluster — never relax the rule

The table below lists 8 categories of qualifying facts; whether and how fully each can be written on your side follows → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers and → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions — this table only gives the writing points themselves.

#	Qualifying fact category	Writing points
1	Price	See 5.2 for the rules by side; the strict side writes no price prefix and no range
2	Licence / registration number	Number + a link to look it up
3	Device / material model	The real brand and model + a link to the manufacturer's page (no link on the strict side)

#	Qualifying fact category	Writing points
4	Repeat-purchase / replacement / remedy policy	Write only the term, the number of times, and what's already included; the strict side does not use words like "free"
5	Process duration and number of visits	An exact number + unit
6	Scope and cap of subsidy / insurance coverage	State clearly what is and isn't covered
7	Byline and credentials	Name + registration number + one line of checkable credentials
8	Operating figures (allowed under restrictions)	See the four gates in the diagram below

Category 8, operating figures, is allowed only through four gates; miss any one and do not write it. The first gate is no adjectives and no comparatives; the diagram below shows the rewrites:

Banned: unsourced adjectives		Replace with: numbers with a clear basis
"Extensive experience", "senior"	→	Year practice began + registration number
"Most cases", "a large number of cases"	→	Strict side: delete; other sides: annual completed volume (basis + data source)
"Foremost", "top-tier"	→	Number of devices on hand / team size
Assertion only, cannot be evidenced	→	Statistical period + data source + sign-off kept on file

Figure: operating figures must be written as evidenced facts, not adjectives

Three more: ② it must be supported by a ledger, with the basis, the statistical period and the data source stated clearly, and the organisation's head signs to confirm it's accurate and keeps the exported record; ③ it must not appear in the same sentence as a claim about efficacy, indications or outcomes, and it belongs only in the organisation-facts section; ④ delete price rise-or-fall figures entirely, since they count as a comparative listing of prices (see the price row in 5.2). For how the strict side writes this, defer to the result-numbers section of 5.2.

This exception is narrow: what's banned is efficacy, outcomes and any subjective praise; what's not banned is an operating fact with a clear basis, evidenced, and free of adjectives — cross that line and it becomes praise (70% confidence, **not measured**; the signed compliance memo is the final word, and where it conflicts with this rule, the memo governs).

> **Example** (dental, strict side) "The clinic has practised at <address> since 2011; the registration numbers of its four dentists are listed on the team page" — an adjective-free, evidenced operating fact about the organisation, not appearing in any sentence describing outcomes.

If the count falls short of 5, try these three steps in order:

1
Get the ledgers
Price lists, purchase orders, scheduling-system exports

2
Narrow the scope
Down to a sub-cluster where you really have exclusive facts

3
Switch intent cluster
Go back to target selection and pick again

Figure: if density falls short, try these three steps in order; never pad with generic educational content

The source gate: at least 3 outbound links to checkable originals per page

Every page needs at least 3 outbound links, pointing to originals that can be verified — not just any related page.

Outbound link type	Example	Why it's worth linking
Original regulation text	Advertising rules, fee guidelines	An authoritative outbound link + turns restricted content into a trust signal
Association / regulator register	Professional registration lookup page	Helps AI match your registration number to you
Device / material manufacturer page	Model spec page	Third-party endorsement of an exclusive fact (does not apply on the strict side)
Government fee / subsidy page	Official fee or subsidy terms page	The field most often copied in price questions

The strict side does not link to commercial companies; the outbound-link scope is in the outbound-link row of → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions; the specific list for your industry belongs in the industry edition.

For the original-regulation-text link type, you can turn a ban into a trust signal along the way: add one sentence on the page explaining why some kind of content isn't here (because it's regulated), and link to the original rule text. The industry's original text is in its own industry edition.

5.9 Content form, FAQ, Chinese pages and other-language pages

What you'll do in this section: decide your side first, then decide whether the content form is professional or consumer: these are two separate judgements, made one after the other, never bundled together; split FAQ into an allowed rule and a banned rule and run them separately; treat the Chinese page as the Chinese version of each buyer type, built in the same batch as the English page, not taking up the second-page slot set by the three-state list, and never missing its three fixed points; for other languages, search first, then decide whether to build, and don't mistake a wrong target for an open field.

Decide your side first, then your content form

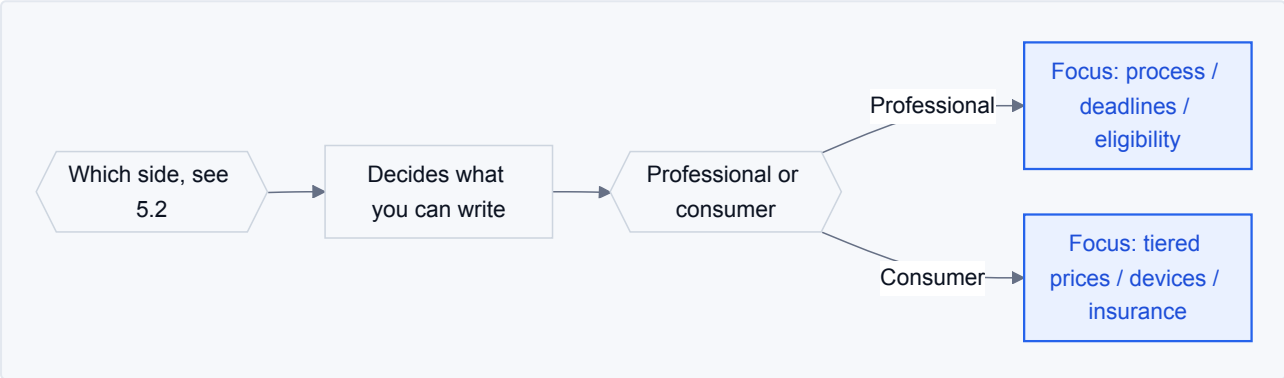


Figure: two separate judgements — decide your side first to fix what you can write, then decide the content form to fix where the emphasis goes

	Professional	Consumer
H2 focus	Process steps, statutory deadlines, eligibility and jurisdiction, fee structure and billing method, common misconceptions	Tiered prices, what's included and excluded, devices and materials, follow-up and replacement policy, subsidies and insurance, waiting times
Main facts	Licence, process duration, byline credentials	Price, device model, subsidy/insurance, operating figures

For why the two forms are written separately, the basis and confidence are in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests. The writing bans for each form are in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers and → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions; here we only give the H2 focus and the main facts. For off-site work, see chapter 6.

Reviews only enter the answer for recommendation-type questions; for price, scenario, process and second-opinion questions, reviews mostly don't enter the answer. So thick pages get no review action assigned to them: no aggregate rating goes on the page, and no review content is excerpted. Evidence in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests. How review requests and replies are handled is in → General Edition 6.5 Reviews: how to ask, how to reply, how to fill the empty slot.

Split FAQ into two rules; do not ban both

Allowed and encouraged	Still banned
Q&A-style H2 (the H2 itself is a question)	Stacking FAQPage schema for rich results
Q&A-style subsections (a question subheading + a direct-answer paragraph)	Self-rating markup, aggregateRating
3–5 short Q&As at the foot of the page, 40–80 words each	Using it as a word-count filler; every one must pass the density gate

Figure: run FAQ as two separate rules; do not lump them together and ban both

A Q&A section is not an exempt zone: the short Q&As at the foot of the page, and the answer under every Q&A-style H2, pass through the same ban gate as the body — check 5.2–5.3 for your side. The strict side must also swap out the questions themselves: don't write questions asking about efficacy, how it feels, or how you compare to others; swap them for questions about process, suitability and what the fee includes.

> **Example** (dental, strict side) The line-crossing question "how soon will I see results" becomes "how many visits does the whole course take, and how long is each one".

Chinese pages are built alongside English pages

A Chinese page is not a language add-on: it's the Chinese version of each buyer type, built in the same batch as the English page, and it does not use up the second-page slot the three-state list sets. By default, the work order produces both an English and a Chinese page for each buyer type, scheduled and accepted together; without the Chinese page, that buyer type is not done. A Chinese page counts as its own page: its own URL, its own H1, hreflang pointing to each other, entered in the register as its own landing page. On the schedule, the Chinese page is built in the same batch as its matching English page — never pushed to a later batch, never saved up to do in bulk at quarter end. Count it as 0.8 of a page's workload — the only thing you save is checking the numbers against the same batch of facts (matching row numbers, not finding evidence again); rewriting the questions, the compliance review, the anchor self-check and the launch acceptance all still have to be done in full.

How a translated draft is done	How a Chinese page is done
H1/H2 are a direct translation of the English titles	H1/H2 are rewritten in the Chinese buyer's own question wording
Answer sentences are translated over, following the English sentence pattern	Every H2's first sentence is written independently, carrying the brand + a number or condition
H2s are back-translated from the English H2s	H2s are taken from the Chinese originals in the question pool
Facts are rewritten and re-sourced from scratch	The same batch of fact-table row numbers is reused, and the numbers are not changed at all

Figure: a Chinese page is not a translated draft; the question forms change, the facts don't

Fix these three points	What happens if you miss them
H1 contains "新加坡" (Singapore)	Without the country fixed, the answer borrows pricing from another country's site
The first paragraph has an S$ amount, immediately followed by GST status (incl./excl.)	Without the currency fixed, another country's currency price takes over
The body has at least one real MRT station or area name	A generic "新加坡" (Singapore) doesn't count; the geographic anchor is what fixes location

Figure: three fixed points on a Chinese page; miss one and answers get anchored to pricing from another country's site

Why this is a hard gate: Chinese-language content defaults to being anchored to other countries' sites — in one live search of a Chinese question, the answer copied a foreign-currency price straight from another country's site. The three points are checked visually, page by page, by whoever accepts the launch; miss any one and the page may not go live: it goes straight back, and this check is never left to sampling.

Other-language pages (Chinese pages and pages in other minority languages alike) must not mix two languages on one page, must not be machine-translated and published directly, and must pass compliance review; on minority-language pages, points ① and ③ apply the same way, and point ② is written in the matching currency's basis.

For Chinese "who's best" questions, a few businesses already hold a spot; the threshold isn't squeezing onto the list, it's whether you have a Chinese landing page at all. Chinese-language price, scenario and second-opinion questions are, for the most part, an open field — this is the only battlefield where you

can get a position from zero within 90 days, and the threshold is just one landing page. This judgement has not yet been retested; the evidence, sample and confidence are in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests; before launch, re-check it against 3 questions sampled from your own question pool.

Other languages: search first, then decide

Three patterns, and they apply to every language:

- Search first, then decide whether to build. If everything you find points to another country's home market, that question is a wrong target, not an open field — swap it for a Chinese-language question instead of investing in a page in the minority language.
- Until the relevant advertising law of the destination country has been checked, run no paid placements there; add one disclaimer-and-scope sentence to the other-language page.
- Build the banned-laudatory-terms list once, turn it into an equivalent list for each target language, and use it for both on-site and off-site; the full industry-specific word list goes in each industry edition's appendix.

5.10 Shared parts for all page types: six checks for numbers, no public price, billing units, dates and schema

What you'll do in this section: run result-type numbers through six checks before they go on a page, and build a definitions table; when there's no public price, fill in the four-sentence template; when converting billing units, put the original unit in the first column; write the three kinds of date separately, and only touch `dateModified` for a substantive change; do schema only where it applies, filling in only the values visible on the page.

Six checks before a number goes on a page

1
Definition
What counts as one success or one case; counter-example: a percentage with no stated standard

2
Denominator
Total number of people or items; counter-example: "100%" with no total given

3
Period
Which year, which cohorts; counter-example: a cumulative figure with no period stated

4
Sample selection
The whole group or only part of it; counter-example: counting only those who stayed to the end

5
Calculation method
How it's calculated, where the data comes from; counter-example: tallied from verbal reports

6
Permission to publish
Authorisation, personal data, industry rules; counter-example: a name written without consent

Figure: a number must pass six checks before it goes on a page; miss any one and it may not go on; each item is followed by a common counter-example

Build the definitions table before you write, not after: some regulators can ask you to show how the data was collected and calculated. The six checks apply to every side; once they pass, go on to the specific rules for your side (see 5.2–5.3).

> **Example** (education) A ready-made compliant wording given by a regulator: "100% students obtained at least 5 passes at GCE 'O' levels (2008 – 2010)" — the year range, the exam's full name and the proportion are all in one sentence.

No denominator: reads like a slogan	With denominator: AI can check it and copy it
"90% pass rate", no headcount given	"<year> <programme> N participants, M passed (X%)"
Only the best figures are written down	Both participant count and pass count are given
The next crawl could well disprove it	Statistical period + definition of basis + data source

Figure: a percentage with no denominator is just an adjective; add the denominator and it becomes a citable fact

Adding the denominator is not a rule requirement, it's a self-imposed spec, and the reason is that it directly decides whether this sentence can be copied. If legal insists on deleting the denominator, give way, but rewrite that line as "Sample: <year> <programme>, all N participants".

What to write when there is no public price

Where there is no public price (a "get a quote" button), fill in these four sentences below it (this wording is not measured):

1. Billing unit: per <unit>
2. Minimum conditions: minimum quantity, shortest term, whether only annual payment is accepted
3. Variables that affect price: 3–5 of them
4. Date statement: "As of <day-month-year>, <Plan> has no fixed public price; pricing is quoted per <unit> based on <variables>."

Do not pass off a guessed figure, or a stale figure from a procurement platform, as the current price.

> **Example** (B2B software) For a piece of software's Enterprise tier, ChatGPT wrote "Custom quote" exactly as given, while copying the real prices of its other tiers as is (measured, 1 case).

Converting billing units

The original billing unit must be kept and placed in the table's first column; when converting to a common unit (such as a price per hour), give the conversion formula its own column; incomparable extras such as enrolment, materials and deposits go in their own column and are not folded into the unit price (**not measured**, though there is 1 measured case as supporting evidence).

Three kinds of date, and schema

Date	How to write it, and when to move it
Price check date	Mark "as of <date>" next to any price that changes often, recheck every 30 days
Regulation check date	Marked separately from the price check date, never merged into one date
Date of substantive change to the body (dateModified)	Move it only for a substantive change, matching the visible "Updated" line

Figure: write the three kinds of date separately; move only the one that changed, never all of them at once

These three date rules are **not measured**. For pages that get a new version each quarter, update `dateModified` and the visible "Updated" line together when the new version goes up, and state in the page which lines changed this time.

Schema is not a proven lever for being named: after adding JSON-LD, AI Overviews (AIO) citations changed by −4.6%, the only significant result, and a negative one; the other two platforms showed no significant effect — schema is a disambiguation tool, not a volume tool, evidence in → General Edition A.1 Evidence for mechanics, the door and identity. Use it where it applies, filling in only the values visible on the page, with the Product / Offer numbers matching the visible price one for one:

Item	Spec
@type	Matches the page's content type
Organisation entity identifier	`sameAs` points to the regulator's registration-lookup page + the business profile page; `identifier` carries the business registration number (valid only for Google's Knowledge Graph; none of the five major systems' live crawls read JSON-LD, so the same fields must be written into the visible HTML first)
Dates	Publication and modification dates are real, matching the visible dates in the body
Price field	Filled in per your side's rules (spec by side in 5.2); the strict side fills in only the fixed final price, with no minimum-maximum range

Do not build FAQPage for rich results (FAQ rich results are reported to have been retired on 2026-05-07, **not independently verified**). There's no need for `name` / `description` to match the title and meta description word for word — matching in meaning is enough.

Chapter 5

Writing pages (2): the 9 most-cited page types, block by block

Full blueprints for price pages, price guides, comparisons, lists, one-question pages and more

5.11 Single-service price page (pt01)

5.12 Price guide page (1): why it is always the first page (pt02)

5.13 Price guide page (2): blueprint, column headers, sentences (pt02)

5.14 Official pricing page (pt03)

5.15 Comparison page (pt04)

5.16 List page (pt05)

5.17 One question, one page (pt06)

5.18 Eligibility and process page (pt07)

5.19 Remedy and second-opinion page (pt08)

5.20 Entity anchor pages: three subtypes, and the product detail page (pt09)

Nine page types, each measured on both legs. One card per type: Intent / which leg / title and H1 / blueprint / standard column headers / sample sentences / most common mistake. Compliance gets only two lines: availability by side and cell-by-cell wording are both in 5.1–5.3, so here we add only the one line specific to this page type. Cited samples, n =, where the copied sentence sits, and length statistics are all in → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types; the cards do not repeat them.

5.11 Single-service price page (pt01)

What you'll do in this section: build one page for every service buyers have already settled on, all from one template: the own-price sentence in the first screen, followed immediately by the price table. These are the two spots ChatGPT copies most often. When you're done, you'll have a copyable single-service price template; change the service name and you can open another page.

Intent: the buyer has already chosen a service and is asking "how much does this place charge". Which leg: **mainly ChatGPT**. This is the page type the book's "own sentence" comes from (evidence: see → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types). One URL per service, the same template (Intro → Evidence block → Decision → FAQ → CTA; empty fields hide automatically).

title: `<Service> Cost in Singapore (<Year>) | <Brand>`; H1: a statement, `<Service> Price in Singapore (<Year>)`, not a question.

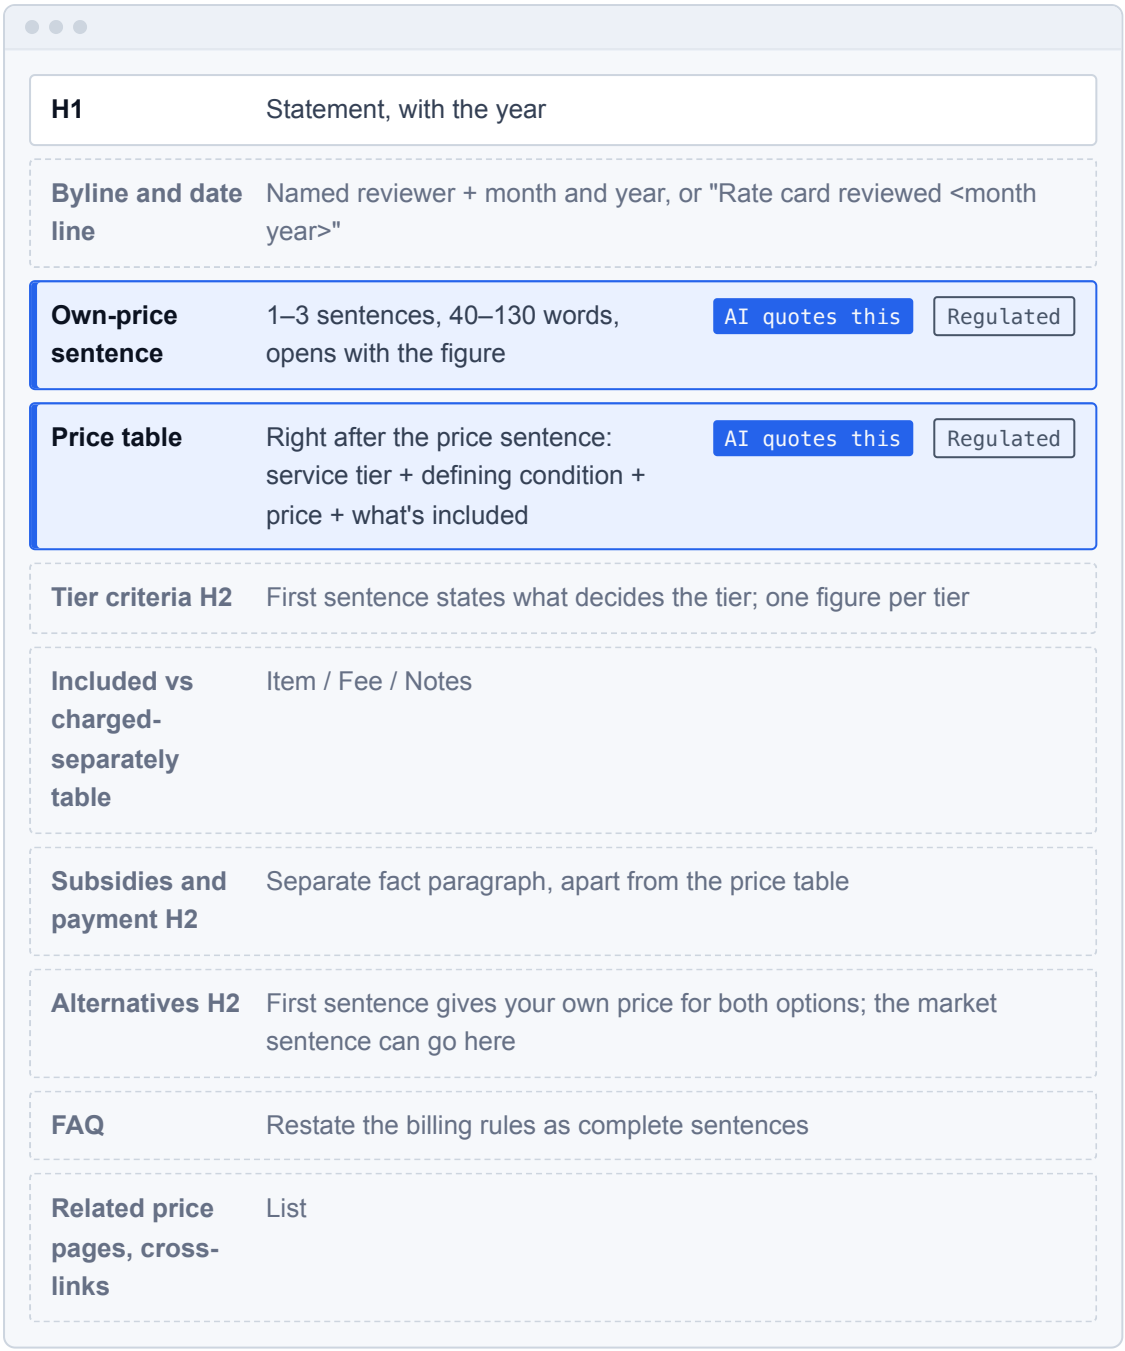


Figure: Single-service price page; AI most often copies the own-price sentence and the price table right after it

Standard column headers: keep only one set — `Item / Fee / Notes`, or by tier `Tier / Includes / Price`; every other industry-specific column header goes in the industry edition.

Sample sentences: "<Service> costs S$400 for <tier 1> or S$900 for <tier 2> at <Brand>, with <add-on> at S$200." On the strict side, every amount is written as the final tax-inclusive price, marked incl. GST. The market sentence is optional for this type; put it in the alternatives block or the FAQ. On the strict side, rewrite it as a variable sentence carrying no peer prices.

The minimum viable version keeps only H1, the own-price sentence and the price table, with each row giving "service tier + defining condition + nett price + what's included". No lower bound is set on length (evidence: see → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types).

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: on the strict side, the price table must not carry any column that implies results or a timeframe (such as a results cell in a spec strip). Keep only prices and objective process parameters.

Most common mistake: the price sits only in an image, a booking widget or a JS toggle, so it cannot be found in the source code. Before handing over, search the source with Ctrl+F for the price figures; if you can't find them, send it back.

5.12 Price guide page (1): why it is always the first page (pt02)

What you'll do in this section: stop arguing about what the first page should be. For every buyer type, the first page is always an all-products price guide page: every item you currently sell goes into one master table, the official public benchmark is handled by side, the H1 carries the year and is revised every quarter, and a Chinese version ships at the same time. When you're done, you'll have a build order that is not up for debate, and three hard specs, none of which can be missing.

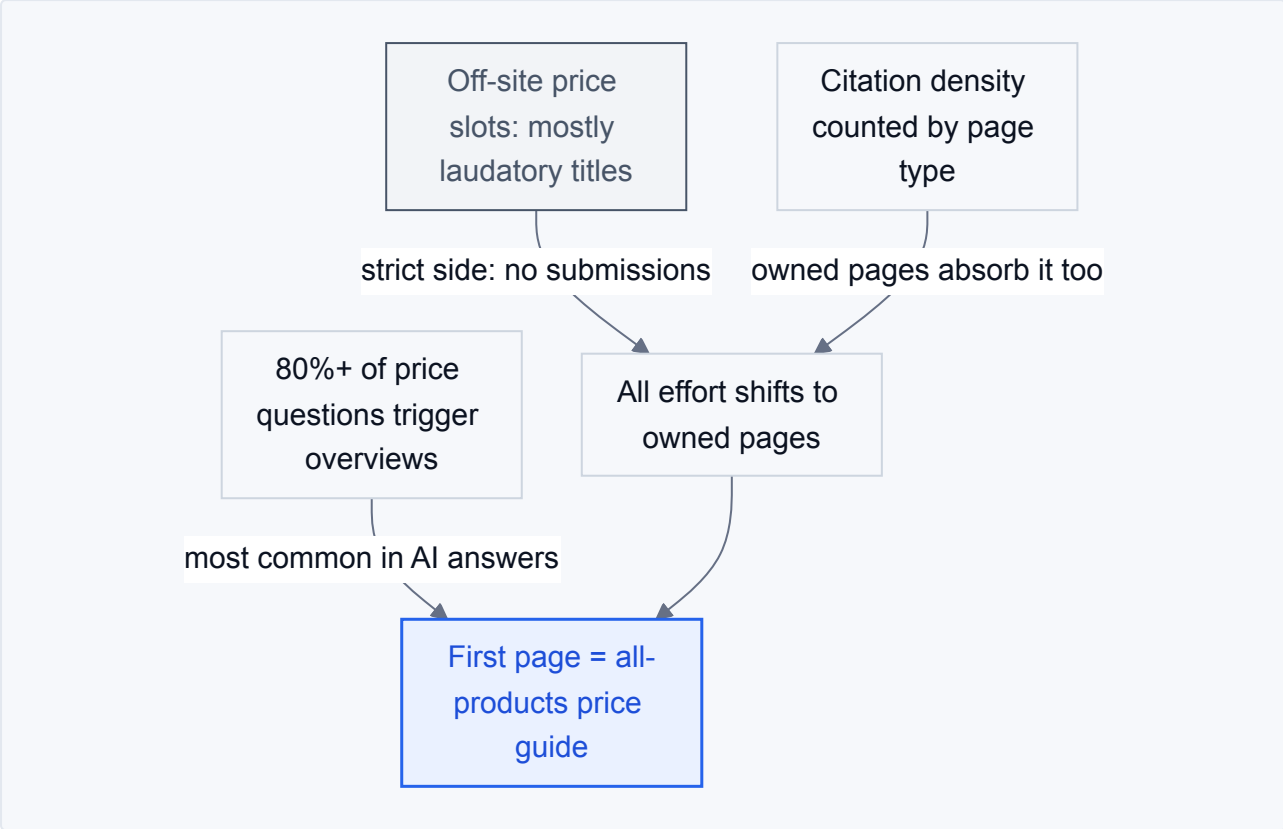


Figure: why the first page is always this one: price questions trigger AI answers most often, the off-site slot is locked by compliance, so all the effort goes to owned pages

The off-site slot is not empty because the market lacks such pages. Third-party price comparison pages do exist, with side-by-side price tables and a way to submit, but almost all their titles carry Best / Top, and on the strict side a title like that rules out any submission or paid listing; if you are already listed, all you can do is ask the publisher to remove you. Owned pages can absorb this density also because, locally, the price guide pages that actually get cited were written by practitioners themselves in the first place. The off-site slot is downgraded to an opportunistic slot. Evidence: see → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests.

Spec	How to check before launch
Every item you currently sell goes into the same master table, each row linked to its single-service page	Row count = the number of items on sale in the signed-off price list; state the reason for any missing row
The official public benchmark is placed by side; write "Not published" where there is no benchmark	The acceptance reviewer opens each official link row by row; if a figure doesn't match, the page does not launch
H1 carries the year, revised every quarter, with "Updated" refreshed at the same time	The page register logs the "next revision date"; if not revised by quarter end, the page has expired

Figure: three specs; miss one and it doesn't count as this page; each has an on-the-spot check before launch

Spec one: build the thick single-service pages as usual, but the entry page is this one master table. A missing row can only be a service you do not offer or have stopped offering, and it must never be dropped silently. Spec three: when you revise the page, update `dateModified` at the same time and state on the page which rows changed this time. A page not revised by quarter end is downgraded to an expired asset.

Spec two follows the stricter side: on the unregulated side you may add a public-benchmark column to the master table; on the strict side, the official benchmark can only be a section of its own, giving the table number and update date, with none of your own prices anywhere in that section and no "we are within the benchmark"; where switch E (peer comparison banned) is on, write no peer figures at all. How to judge the three sides and the switches: see → General Edition 0.2 Decide which side you are on first.

The entry question: ask it by side

On the strict side the question is "are you willing to publish a fixed final price and tiered, itemised prices"; on the light side the question is "are you willing to publish a fee range and your billing method". **Don't use "willing to publish a range" as the entry bar for everyone.** Writing a range is itself not compliant on the strict side, so using it as the entry test would wrongly rule out a client who is in fact compliant and simply cannot write prices as a range.

The Chinese version ships at the same time; it is not an extra

By default this page ships in both English and Chinese at the same time: separate URL, separate H1, and hreflang tags pointing to each other. The reason: most Chinese-language price questions are still unclaimed, so this is the page most likely to go from 0 seats to at least one within 90 days; for the open-field share, see → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests; for how to write the Chinese page, see → General Edition 5.9 Content form, FAQ, Chinese pages and other-language pages.

This page's density gate starts from three fact categories: price, subsidies and insurance, and process duration; for the density gate itself, see → General Edition 5.8 The conclusion block, the density gate and the source gate.

5.13 Price guide page (2): blueprint, column headers, sentences (pt02)

What you'll do in this section: build the price guide page from the blueprint. The conclusion block gives an own sentence plus a market sentence, the first H2 is the main price question with the price grid table right below it, and the official public benchmark goes into its own section, placed by side. When you're done, you'll have a price guide page skeleton you can copy directly, and every figure in the conclusion block can be found in the table.

Intent: the buyer hasn't chosen a business yet and is asking about pricing across the whole category. Which leg: **both legs cite it, AI Mode more**; when ChatGPT treats it as an overview, it cites the official benchmark alongside it. Length comes from the table, not the word count; there's no lower bound on the number of figures, and outbound links are not a condition for being cited (evidence: see → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types).

title: `<Category> Cost in Singapore (<Year>): <Price hook> | <Brand>`; H1: `<Category> Cost in Singapore (<Year>)` or `<Category> Rates in Singapore (<Year>)`, not a question.

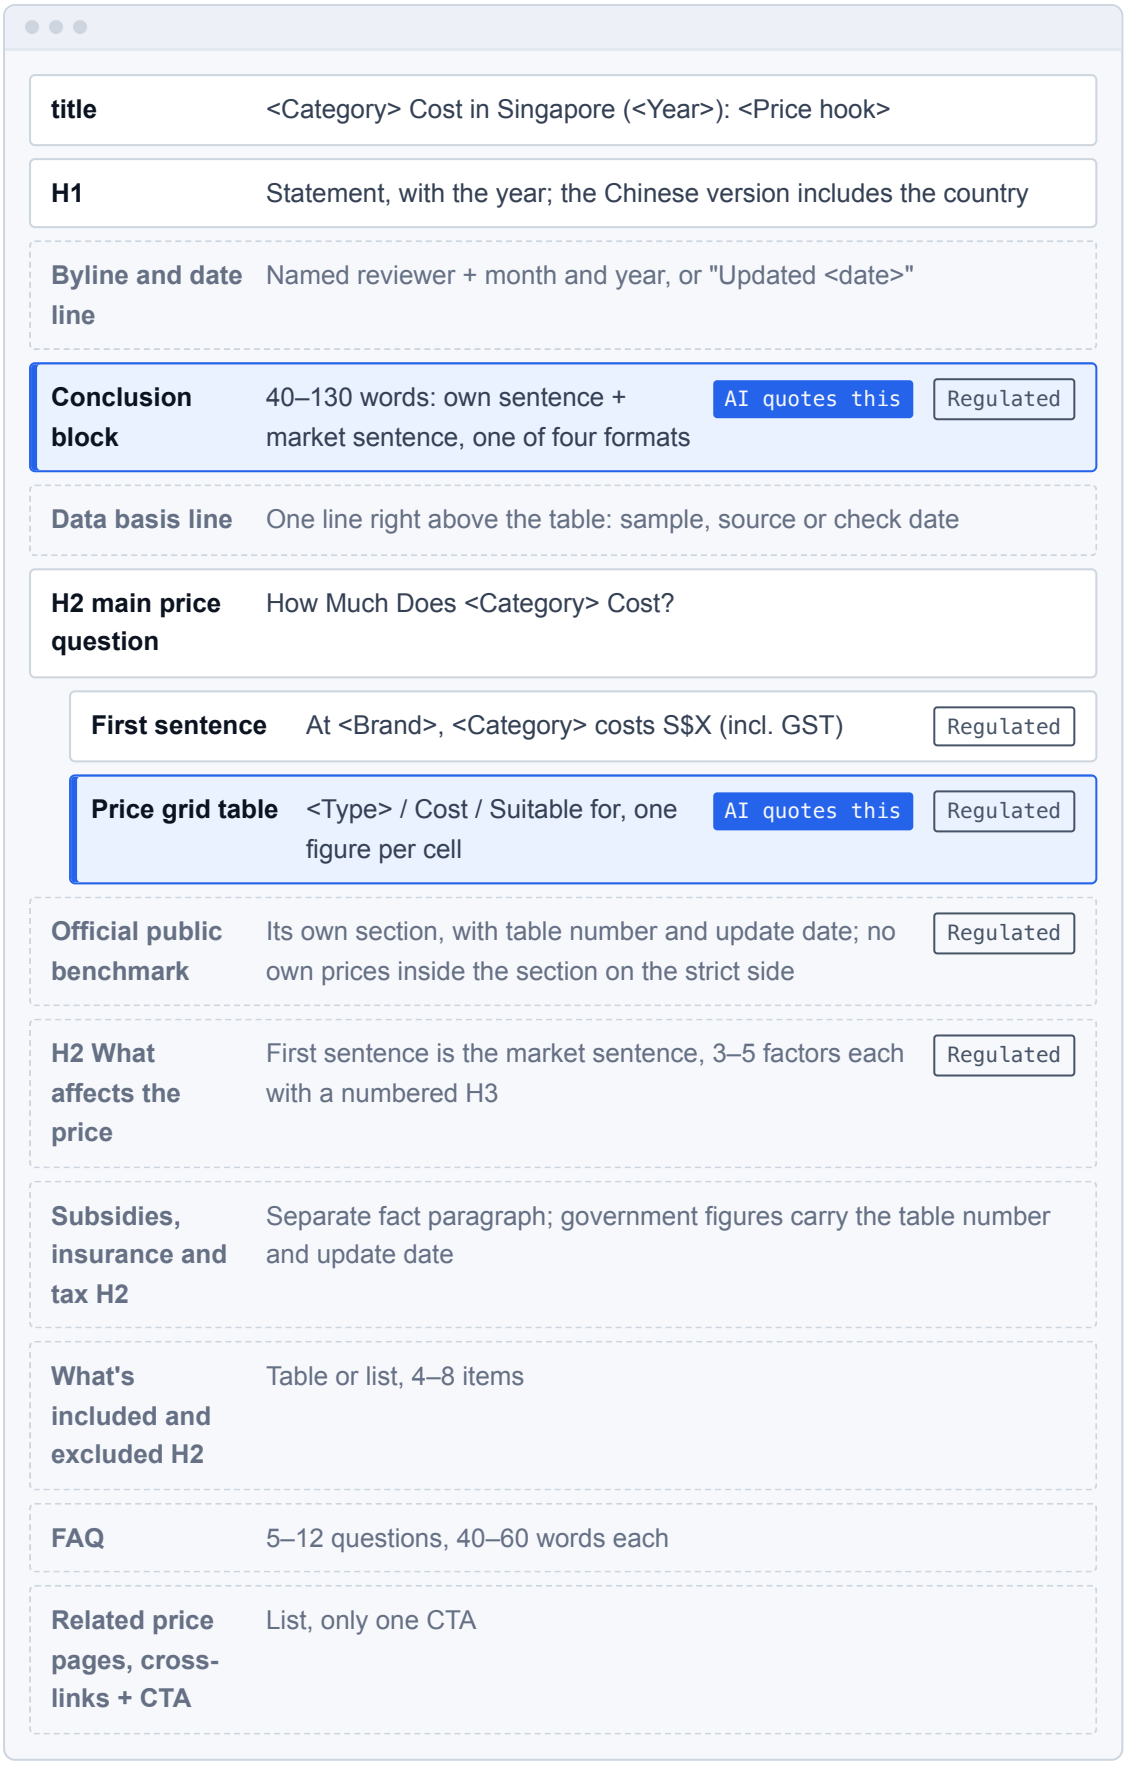


Figure: price guide page block order; AI most often copies the conclusion block and the price grid table under the first H2

The conclusion block can be one of four: a labelled text box, a full price list, a KPI stat card, or a price card. For the three fixed points on the Chinese version, see → General Edition 5.9 Content form, FAQ, Chinese pages and other-language pages; the blueprint only notes "includes the country" on the H1 row and does not draw a separate diagram.

Standard column headers: keep only one set — `<Type> / Cost / Suitable for`, one figure per cell, with the unit and basis. "Market range and own price in the same table" is for the unregulated side only. Write subsidy worked examples separately from the price table, and state that a subsidy is a way of paying, not a price cut.

Sample sentences: own sentence: "At <Brand>, <Service> costs S$X (incl. GST) for <defining condition>"; market sentence: "<Scope> generally ranges from $X to $Y, depending on A, B, C". The unregulated side writes this pattern as it is; the light side (switch E is not on) writes it only in a separate fact section with its source, never in the same table as your own price and never next to it (Conservative line (not statute text); see 5.2). On the strict side, and where switch E (peer comparison banned) is on, replace the market sentence with an official-basis sentence: same position, word count and sentence pattern as the market sentence, but describing a regulation or a procedure instead; for how to write it, see → General Edition 5.6 Two legs, two kinds of sentence.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: besides the conclusion block, the first sentence of the "What affects the price" H2 also counts as a market-sentence slot. On the strict side, and where switch E (peer comparison banned) is on, it too must become an official-basis sentence, or a variable sentence with no peer prices.

Most common mistake: the first screen only sets the scene without giving a figure, or the first-screen summary doesn't match the table's numbers. The fix: take every figure in the conclusion block from the table, and check them one by one before handing over; when the top card and the table disagree, AI uses the table's figures (evidence: see → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types).

5.14 Official pricing page (pt03)

What you'll do in this section: for a product with a public price list, fill the /pricing page completely. The price cards sit right under the H1, and none of the four elements is missing: the price unit, the annual and the monthly price, the free-tier cap, and the minimum number of seats. When you're done, you'll have a price source that comparison pages, list pages and third-party pages all link back to.

Intent: the buyer asks "how much is <Brand> per month"; this type also picks up "best <category>" and "A vs B", because when ChatGPT compares products it comes back to this page to check the price. Which leg: **mainly ChatGPT** (evidence: see → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types). Every product with a public price list must have this page, filled in completely; comparison pages, list pages and third-party pages all link their prices back to it.

title: `<Product> Pricing Plans | Starting from $X/<unit>/month`, no year: this is the exception to rule 4 of → General Edition 5.7 Seven rules that hold for every page type; H1: `<Product> Pricing` or `Plans for <Product>`, not a question.

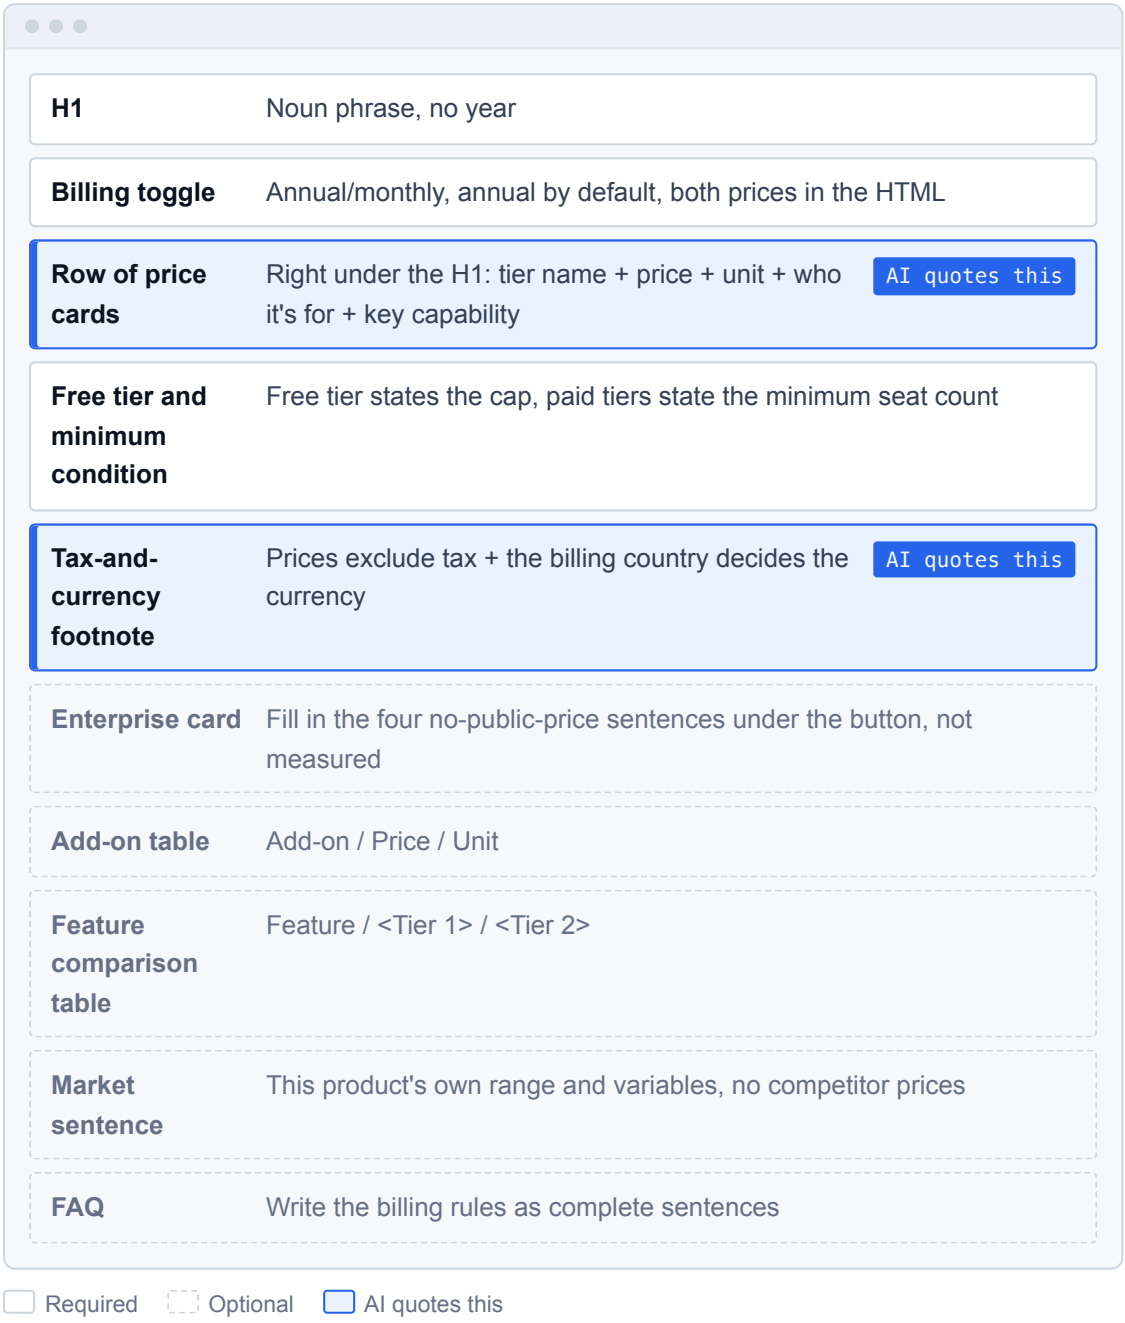


Figure: official pricing page; AI most often copies the price cards and the tax-and-currency footnote

If any one of the four price-card elements is missing, ChatGPT goes to a third-party page to fill the gap. All four must be in the first-screen HTML, not waiting for a click or JS to appear. There is no length requirement, and no byline is needed; add one extra line, `Prices last checked <date>` (not measured). For how to write the Enterprise card's four sentences, see → General Edition 5.10 Shared parts for all page types: six checks for numbers, no public price, billing units, dates and schema.

Standard column headers: add-on table `Add-on / Price / Unit`; feature comparison table `Feature / <Tier 1> $X <unit>/month / <Tier 2> $Y <unit>/month`.

Sample sentences: keep only one own sentence: "Basic costs $X per seat per month, billed annually; paid plans start from N seats, and prices do not include tax." The market sentence gives only this product's own price range and variables, never a competitor's price (sample, not measured): "Most small teams on <Product> pay $X–$Y per seat per month billed annually, depending on plan, number of seats and add-on usage."

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: none. Availability and wording are all in 5.1 and 5.2: on the strict side, and with switch E (peer comparison banned) on, don't build plan cards or a recommended ranking; build → General Edition 5.11 Single-service price page (pt01) instead; where switch B (legally required fields) is on, build an all-in price page.

Most common mistake: only a slogan and "Contact sales" sit under the H1, and the price cards need a click or JS to appear. A strong domain's empty-shell pricing page can still get cited. That is a strong-domain exemption, and a weak domain cannot rely on it.

5.15 Comparison page (pt04)

What you'll do in this section: when the buyer is asking "A 还是 B" ("A or B"), organise the comparison page by scenario. The verdict goes into the conclusion block and the first sentence of the first H2, the comparison table sits in the first 25%, each criterion H2's heading names the winner, and you write a separate section for scenarios where "the other one fits better"; on regulated sides, compare only options, never organisations. When you're done, you'll have a comparison page with the verdict up front and a structure organised by scenario, not candidate by candidate.

Intent: "A vs B which is better", "A or B for <scenario>", "A 还是 B" (A or B). Which leg: **mainly AI Mode**; for the citation ratio between the two legs, see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. When ChatGPT meets a "vs" question, it mostly goes to both sides' official pricing pages, spec pages and help centres and builds its own table, so every figure about your own product on the comparison page must come from the same source and carry the same

value as your official pricing page and product page (the same-source point is not measured). Two cases where you should not use this type: on a regulated side, wanting to compare organisations (not allowed; compare options only); or comparing only your own models (that is a category list page).

title: `<A> vs <B> (<Year>): Which Is Better for <Scenario>?` ; H1 defaults to a statement, `<A> vs <B> for <Scenario> (<Year>)` ; the first H2 is the main question, `Which Is Better for <Scenario>?` .

title / H1	title is a question, H1 is a statement, both carry the year	
Comparison-basis line	Check date + both sides' version or tier	
Conclusion block	60–120 words, or a by-scenario verdict table	AI quotes this · Regulated
H2 main question	First sentence gives the verdict with its condition, followed right away by the scenario table	AI quotes this
Criteria H2	Numbered H3s, each first sentence states the threshold that flips the verdict	
Master comparison table	In the first 25%, one checkable value per cell	Regulated
Per-criterion H2 × N	The verdict goes into the heading; the first sentence repeats it with a figure	
Price comparison H2	Same tier, same currency, same billing cycle	
Scenarios where the other one fits better H2	≥2 specific scenarios, each paired with one public fact about the other side	
Not-a-fit conditions H2	Skip it if…	
Data sources and dates	Source + check date for every figure	
FAQ	H3 questions, first sentence answers directly	
Related comparison pages, cross-links	3–5 entries	

Required · Optional · AI quotes this · Regulated

Figure: comparison page; AI most often copies the conclusion block and the verdict in the first sentence of the first H2

On pages where the verdict sits in the middle, only that middle sentence gets copied (evidence: see → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types), so organise the blocks by scenario, not candidate by candidate. The four formats each have their place: "[Competitor] Alternative" and "You vs [Competitor]" are limited to the unregulated side; the former keeps the "who should switch / who should not" and migration section; "Alternatives" in the plural is really a list, and belongs under → General Edition 5.16 List page (pt05); naming yourself inside an "A vs B" page is allowed on the light side, but must come after the comparison and verdict, while regulated sides compare only options and never bring in "us". For comparisons of physical goods, AI Mode prefers to cite video, so make the video first and the web page its companion.

Banned: write this and it's disparaging		Rewrite: put this in the same spot instead
Writing only your own strengths, never the other side's	→	Make "scenarios where the other one fits better" its own H2, with ≥2 specific scenarios
Writing the other side's figures with no source at all	→	Every cell carries a public source and check date; where there's none, write "not published"
"better / leading / more powerful" with no metric behind it	→	Every such adjective must be followed immediately by a checkable metric
Not making both sides' basis clear	→	Write the version and date in the table header
Leaving the other side's price up after it's out of date	→	If an expired price isn't refreshed, delete the cell; label prices as of <date>

Figure: five criteria for not disparaging, which even unregulated-side comparison pages must pass

Self-check for criterion 1: delete the "the other one fits better" section. Does the verdict still hold? If it does, you haven't told the truth. Criterion 2: for any cell with no public value, write "Not published as of <date>"; estimates never go into the table. Criteria 3–5 are not measured; recheck criterion 5 every 30 days.

Standard column headers: scenario table `If your situation is… / Choose / Deciding fact / Source`; keep only one set for the master comparison table, as a cross-industry example: `Criterion / <A> / <B> / Better for`. One value per cell, with a unit; the same row uses the same tier, the same date, the same currency, the same test conditions; write one line for the data basis above the table.

Sample sentences (figures are placeholders): own sentence: "<Own product> Starter costs S$X per user per month billed yearly, with a free tier capped at N users (pricing page, checked <date>)." Market sentence: "For a team of under 10, both tools generally cost $X–$Y per user per month, depending on

seat count, annual vs monthly billing, and whether add-ons are needed." The unregulated side writes the market sentence as it is; the light side (switch E (peer comparison banned) is not on) writes it only in a separate fact section with its source (Conservative line (not statute text); see 5.2); the strict side, and anyone with switch E on, do not write it.

How to split a comparison page from a concept pillar page: a title containing "哪个更好" ("which is better") belongs to this type; a title containing "有什么区别 / 什么是" ("what's the difference / what is") belongs to the concept pillar page; see → General Edition 5.30 The four tier-C types: unverified page types.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: the comparison table must not have a peer column or a market-range column (on the strict side, or with switch E (peer comparison banned) on). This is a problem of table structure, not of wording.

Most common mistake: opening with background and pushing the verdict to the middle, so only that middle sentence gets cited; on a regulated side, turning an option comparison into an organisation comparison: that is a compliance breach, not a layout issue.

5.16 List page (pt05)

What you'll do in this section: first check → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side and → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions to judge whether your side may self-build a list; if it may, write the methodology or the check date into the title or the first screen, write each item's H2 as "<Name>: Best for <Scenario>", make every figure in the at-a-glance table match the item card figure for figure, and disclose it whenever your own business ranks first. When you're done, you'll have a list-page template built to the structure of lists that actually get cited.

Intent: "Best <Category> for <Scenario>", "<Competitor> alternatives", "Top N <Category>". Which leg: **mainly AI Mode**; for the citation ratio, see 5.1; everything the ChatGPT leg cites is a third-party list; it did not cite a vendor's own list even once (evidence: see → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types). A "who's best" question about choosing an organisation does not use this type; see → General Edition 0.3 Which page types are limited for you, and what to build instead.

The condition for getting into ChatGPT: every list ChatGPT cites has the methodology or the check date written into its title or first screen. A list where your own business ranks first can still be cited by AI Mode, but the conflict-of-interest disclosure line must be there.

title: `<N> Best <Category> for <Scenario> in Singapore (<Year>)`; to get into ChatGPT, also write the methodology or the check date into the title, such as `… (Tested on <N> Models)` or `… Prices Checked <Date>`; H1 is a noun phrase that makes a statement, not a question; the first H2 is `Which <Category> Is Best for <Scenario>?`, whose first sentence names the winner with its condition.

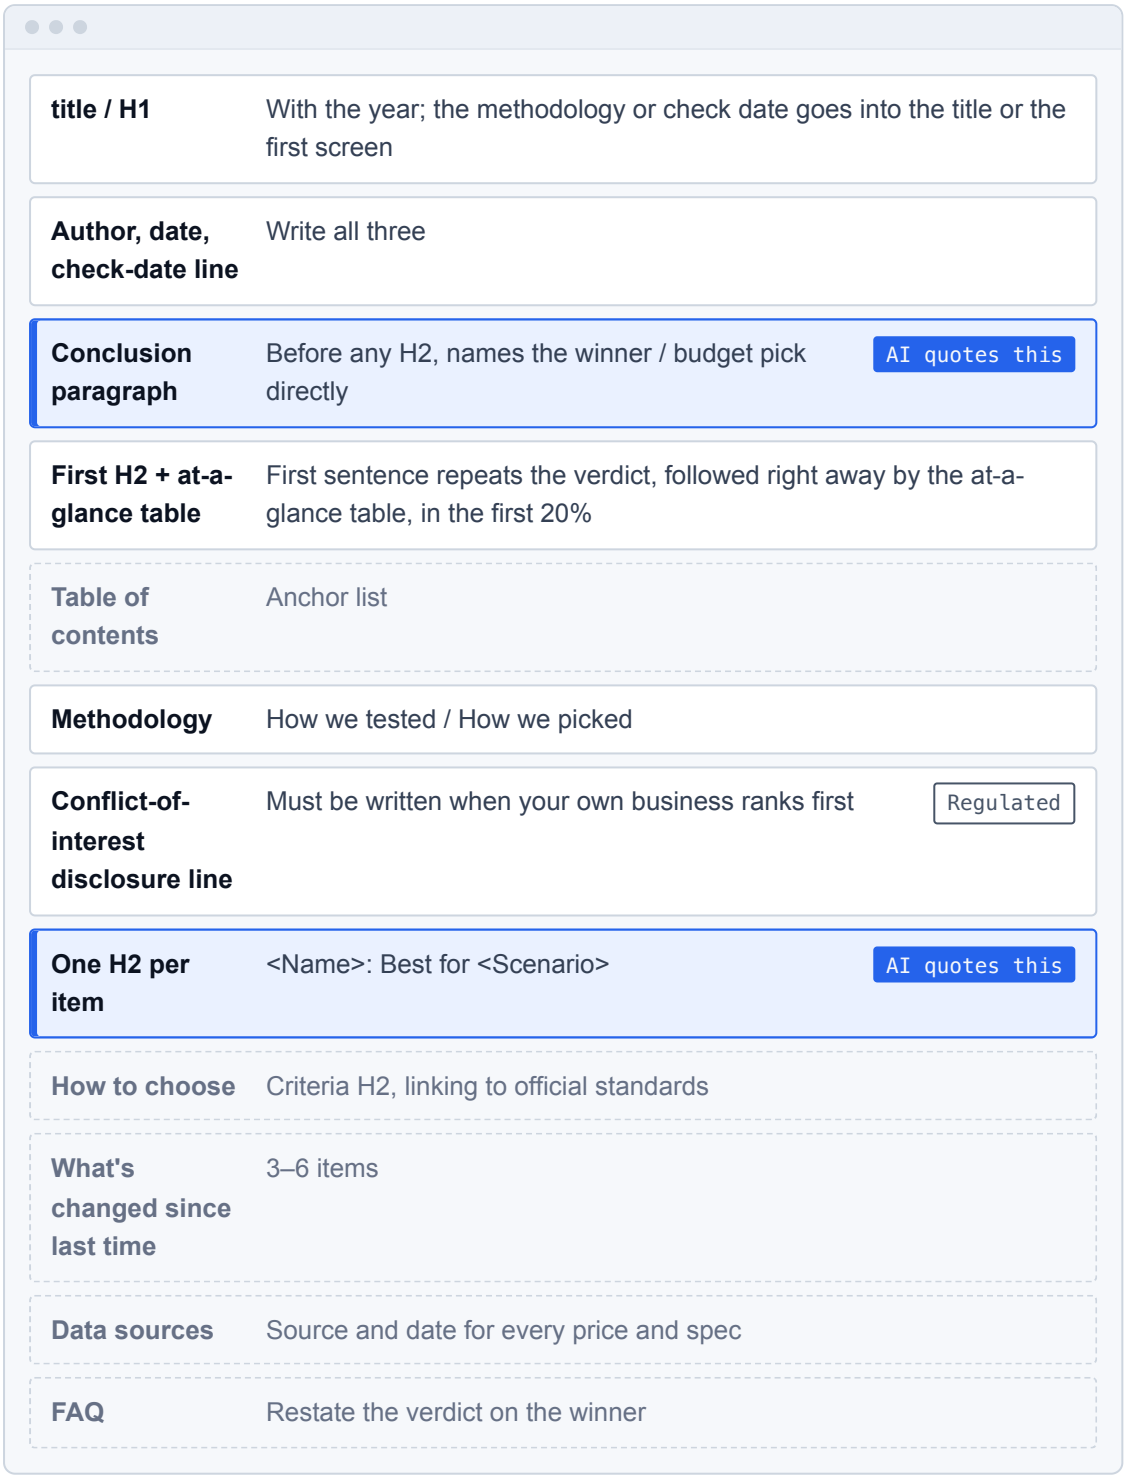


Figure: list page; AI most often copies the conclusion paragraph that names the picks, and the "<Name>: Best for <Scenario>" H2s

For an item whose heading gives only a name, the line right below it must be `Best for: …`; the order inside the card is: first-sentence verdict → spec strip → 2–3 pros and 2–3 cons → `Skip it if: …` → the actual price + check date + source. Every figure in the at-a-glance table must match the item card figure for figure, and the price column header carries the check date.

Standard column headers: keep only one set as a cross-industry example — `<Subject> / Best For / <Key metric> / Starting Price / Free Tier or Trial`.

Sample sentences (figures are placeholders): specific sentence: "The <Model> was S$X at <Retailer> on <Date>, with a <N>-year local warranty." Market sentence: "<Category> suited to <Scenario> generally cost S$X–S$Y, depending on <Variable A>, <Variable B>, and <Variable C>."

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: none. Availability and the wording for each side are all in 5.1 and in the list-page figure in 5.3; this card only points there and does not rewrite it side by side.

Most common mistake: each item has only a name and a description, with no "Best for" verdict sentence and no actual price. On a weak domain, a list written like this gets skipped.

5.17 One question, one page (pt06)

What you'll do in this section: for single-point questions with a definite answer (can I, should I, what's the cap, how long does it take, who's eligible), give each question its own URL, with the H1 using the buyer's original question word for word, the answer paragraph right under the H1, the first sentence giving Yes/No or a figure, and any conditions in the 2nd and 3rd sentences. When you're done, you'll have a batch of short-answer pages that can be checked sentence by sentence against the original government page.

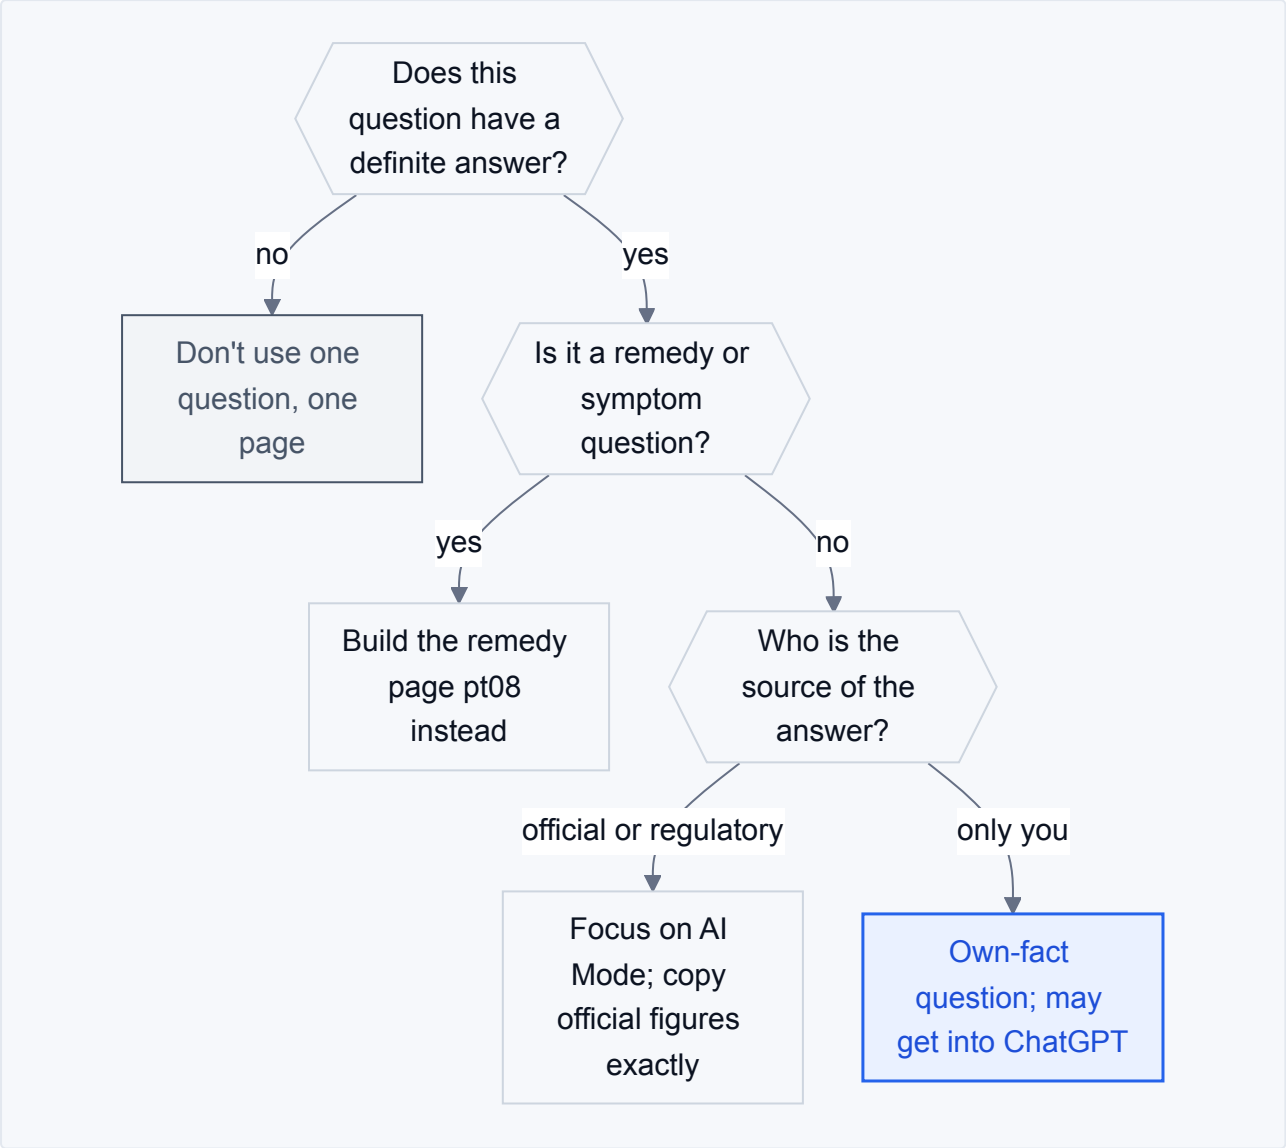


Figure: which questions use one question, one page; for those that do, the answer's source decides which leg to focus on

Which leg: **government versions get into ChatGPT, commercial versions get into AI Mode** (evidence: see → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types). For market or regulatory fact questions, the ChatGPT leg is occupied by government pages, so commercial sites should focus on AI Mode; an own-fact question can only get into the ChatGPT leg when you are the sole source (not measured). For questions with no definite answer, ChatGPT cites only government and media; for remedy and symptom questions, write the safety triage first [Safety]; see → General Edition 5.19 Remedy and second-opinion page (pt08).

H1: the buyer's original question, word for word, with a local qualifier. This is the only one of the nine types whose H1 is a question; title: `<Original question>? <Short-answer hook>`, and the year is added only when the answer changes over time and someone maintains it (not measured); the URL slug is the question itself, and one URL answers only one question.

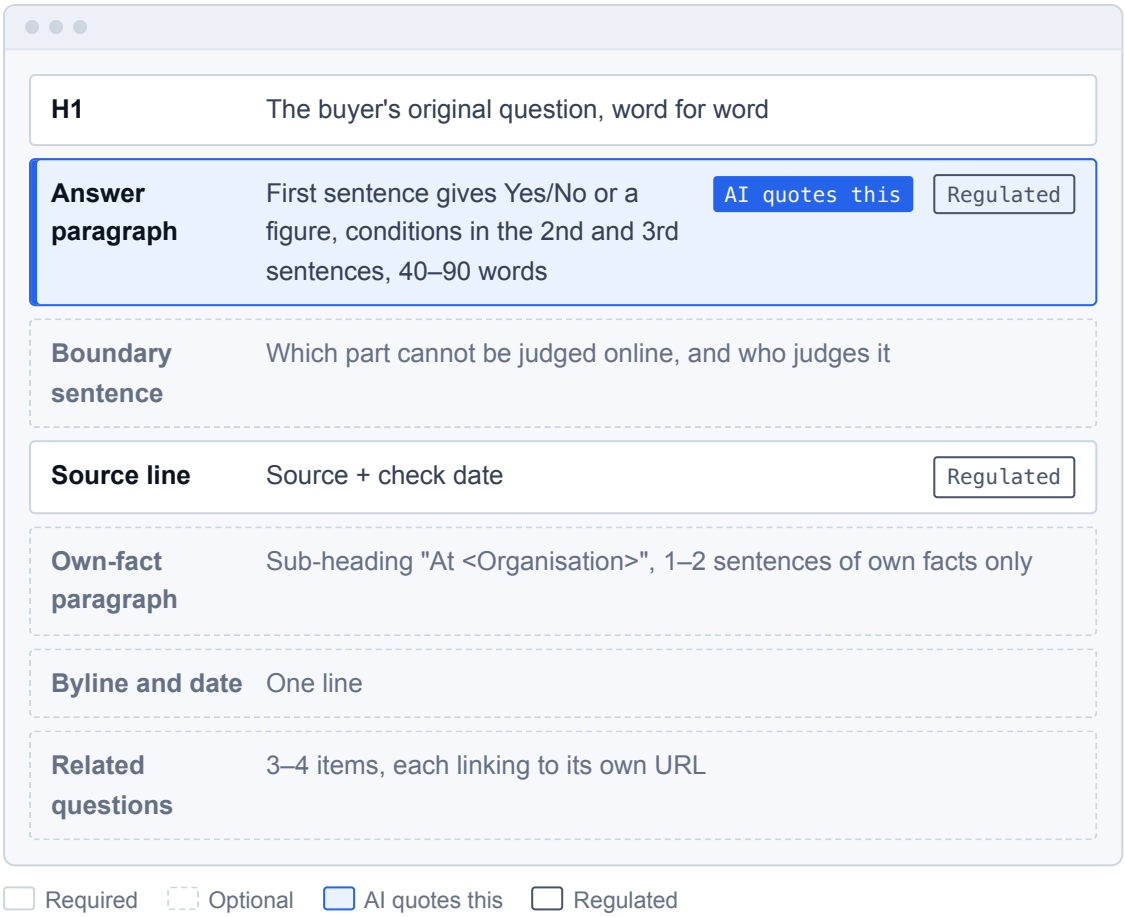


Figure: one question, one page, short version; what AI copies is the answer paragraph right under the H1

A regulatory fact question with many conditions uses the long version: the answer paragraph comes before the first H2, followed by 4–8 verdict-statement H2s, an optional conditions table, and related questions each linking to their own URL; the short version is 40–300 words, the long version no more than 1,500 words. Most one-question pages need no table; a table is only added when a single answer carries several conditions.

Standard column headers: for eligibility questions, use the question itself as the row name; for limit questions, use `Item / Limit / Condition / Source`.

Sample sentences (figures are placeholders): regulatory sentence: "Yes. Under <Rule>, <Item> is allowed up to $X; the remainder is generally paid <method>, and the total depends on <variable>." The figures must match the official page word for word; note the check date. Own sentence: "<Own business> accepts <that item>. For <defining condition>, the final fee is S$X including GST and includes <mandatory inclusion>; <the variable part> is priced after assessment." Write the billing rules as complete sentences: whether tax is included, the bundled price for the first N people, the per-unit price after that, and the local currency.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify

the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.

- Specific to this page type: the answer gives only general information, never an individual diagnosis; questions involving symptoms or complications move to the remedy and second-opinion page (5.19) [Safety].

Most common mistake: turning one question, one page into one page with ten questions (really an FAQ collection), or opening the answer with "it depends".

5.18 Eligibility and process page (pt07)

What you'll do in this section: when answering "what conditions are needed, how many visits, how long it takes, how to apply, what to bring", use a task-style statement for the H1, put a five-cell key-value spec strip in the conclusion block, and have the first H2's first sentence give the verdict followed right away by the steps table. When you're done, you'll have a process page that feeds both legs at once: a source for ChatGPT on your own process and duration, and, for AI Mode, a sentence of the form "generally X–Y, depending on A/B/C".

Intent: what conditions are needed, how many visits, how long it takes, how to apply, what to bring; if you're only answering a single point, use one question, one page (5.17) instead. Which leg: **mainly ChatGPT**, where government process pages are cited the most; commercial sites' process articles only turn up in AI Mode (evidence: see → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types).

H1: a task-style statement, `How to <Task> in Singapore`, or `<Service> in Singapore: Steps, Timeline and Eligibility`, not a question; title = H1 + organisation name, and the year is added only when the rule changes by year and someone maintains it (not measured).

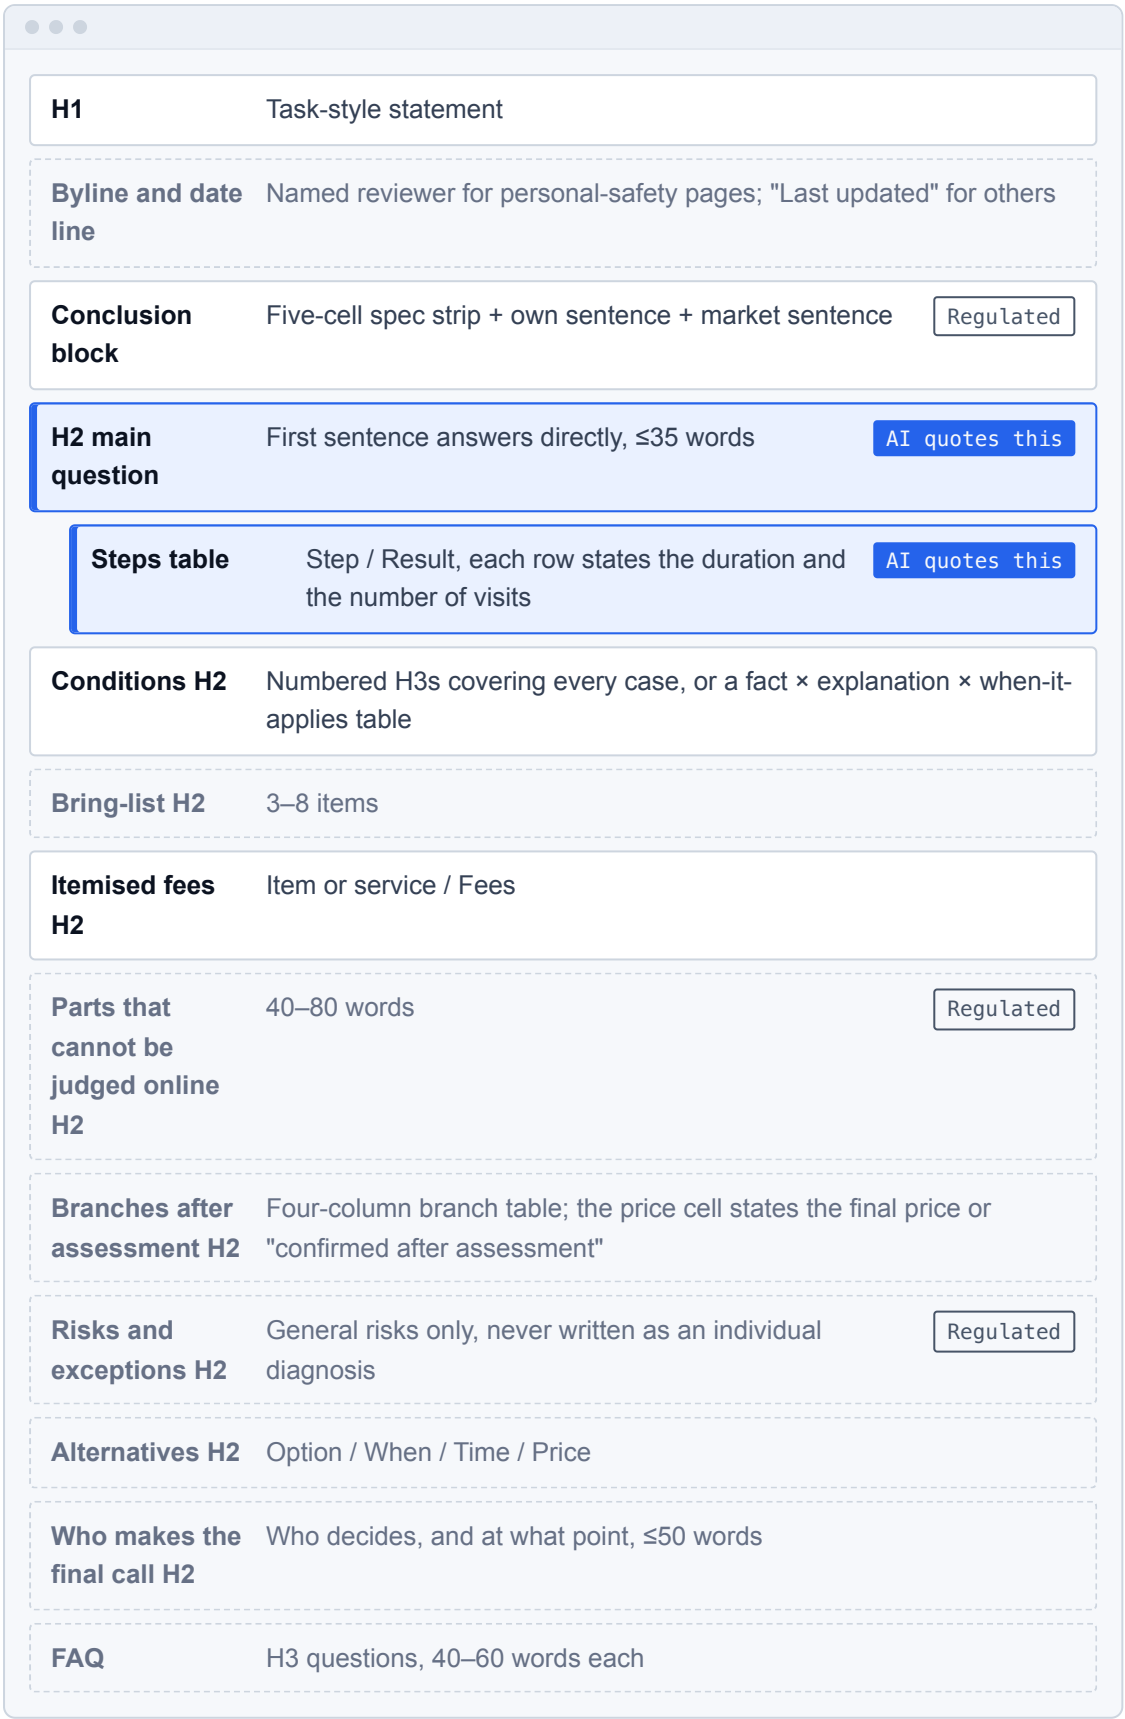


Figure: eligibility and process page; AI most often copies the verdict in the first H2's first sentence and the steps table

The five cells of the conclusion block are `Total time / Visits or steps / Who it is for / Who it is not for / Who decides`, followed by one own sentence and one market sentence. A government page's conclusion block states only the scope and gives no answer. That is a privilege of government pages; commercial pages should not copy it. When there are many conditions, use the questions as table row names (Who is this for? / How much?), with columns split by group or format, one checkable figure per cell. End the page with a Need help section or 3–4 related-page links.

Standard column headers: `Item or service / Fees`; `Step / Result`; branch table `If the assessment finds / Next step / Added time / Price (final, incl. GST) or "confirmed after assessment"`.

Sample sentences: own sentence: "At <Own business>, <Task> takes <N> visits and <duration range> from <starting point> to <end point>; whether <variable> is needed is decided by <professional> after <assessment>." Market sentence: "<Task> generally takes <duration range>, depending on <Variable A>, <Variable B>, and <Variable C>." Measured: asked how long a particular process takes, ChatGPT cited a government page and wrote "typically takes about 4 to 5 months from filing to Final Judgment", and AI Mode wrote almost the same sentence. A range plus a condition is the sentence shape both legs copy.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type [Safety]: on pages related to personal safety, write who is suitable as "judged by a professional after an assessment", never as a self-check list for the reader. A self-check list amounts to letting readers diagnose themselves.

Most common mistake: conditions written only as "it depends".

5.19 Remedy and second-opinion page (pt08)

What you'll do in this section: when the buyer is asking "what do I do if it went wrong, who do I complain to, I want a second opinion", first pick a skeleton by risk type. For personal-safety cases use pt08-safety, which leads with what to do now and the red flags; for rights-deadline cases use pt08-deadline, which leads with the urgent deadlines. When you're done, you'll have one of the two skeletons, with its order fixed: not one slot may move.

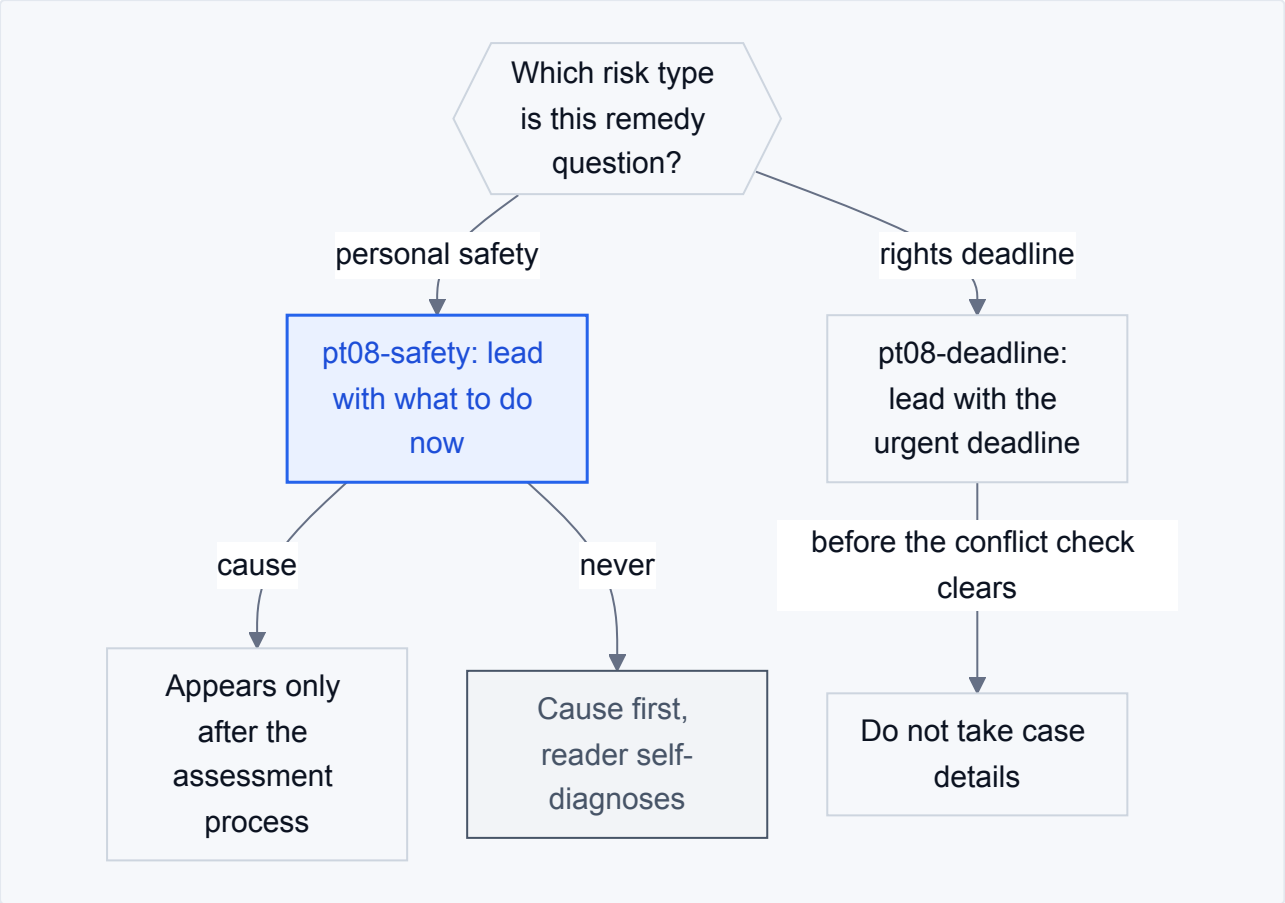


Figure: sort a remedy question by risk type first; each skeleton has a fixed opening that must not move, and the two must never be mixed

Intent: what to do when something went wrong or failed; who to complain to when there is a problem; wanting a second opinion. Which leg: both legs cite it, but not the same kind of page (evidence: see → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types). The subject of the H1 is the buyer's situation, not the organisation; you may state neutrally that earlier treatment took place elsewhere, but these three are banned without exception: assigning blame, disparaging, and any unproven judgement that something failed.

pt08-safety: personal-safety type

This skeleton's order is fixed; not one slot may move [Safety]. The reason: when AI itself answers this kind of question, it also leads with safety, so a page in the same order matches the structure of AI's answer. The sample is only 2 pages; evidence: see → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types.

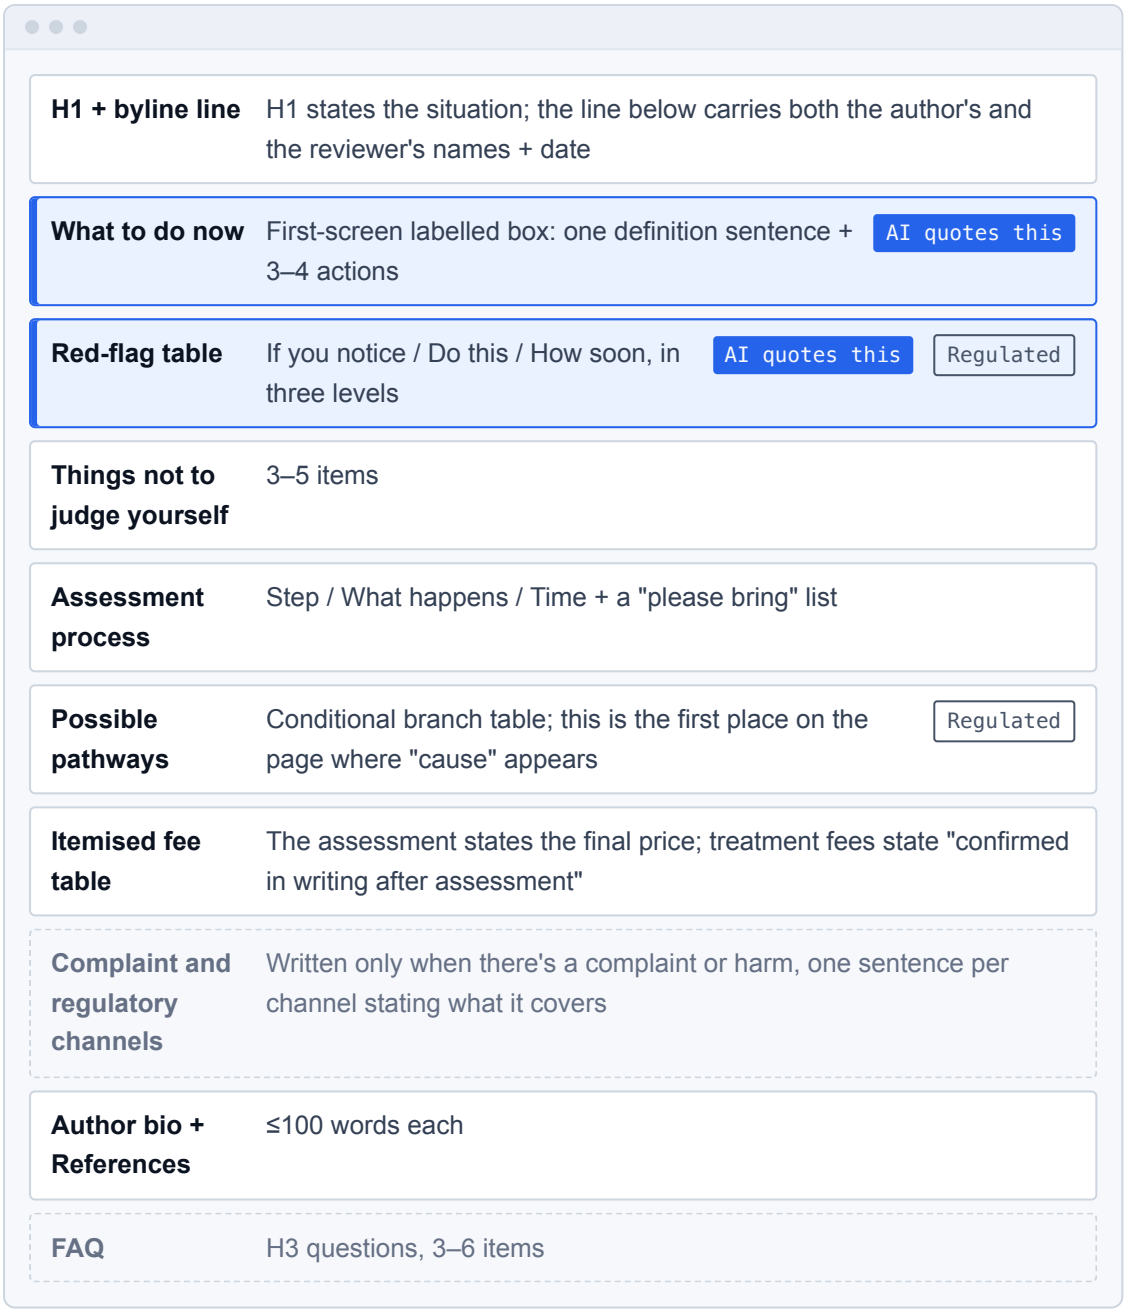


Figure: personal-safety remedy page; not one slot in the order may move; AI copies the first-screen definition sentence and the red-flag table

The red-flag table's three levels are: emergency care immediately, same day, this week.

pt08-deadline: rights-deadline type (not measured)

This whole blueprint is not measured: there is no cited page to compare it with, so no block in the figure is marked cite.

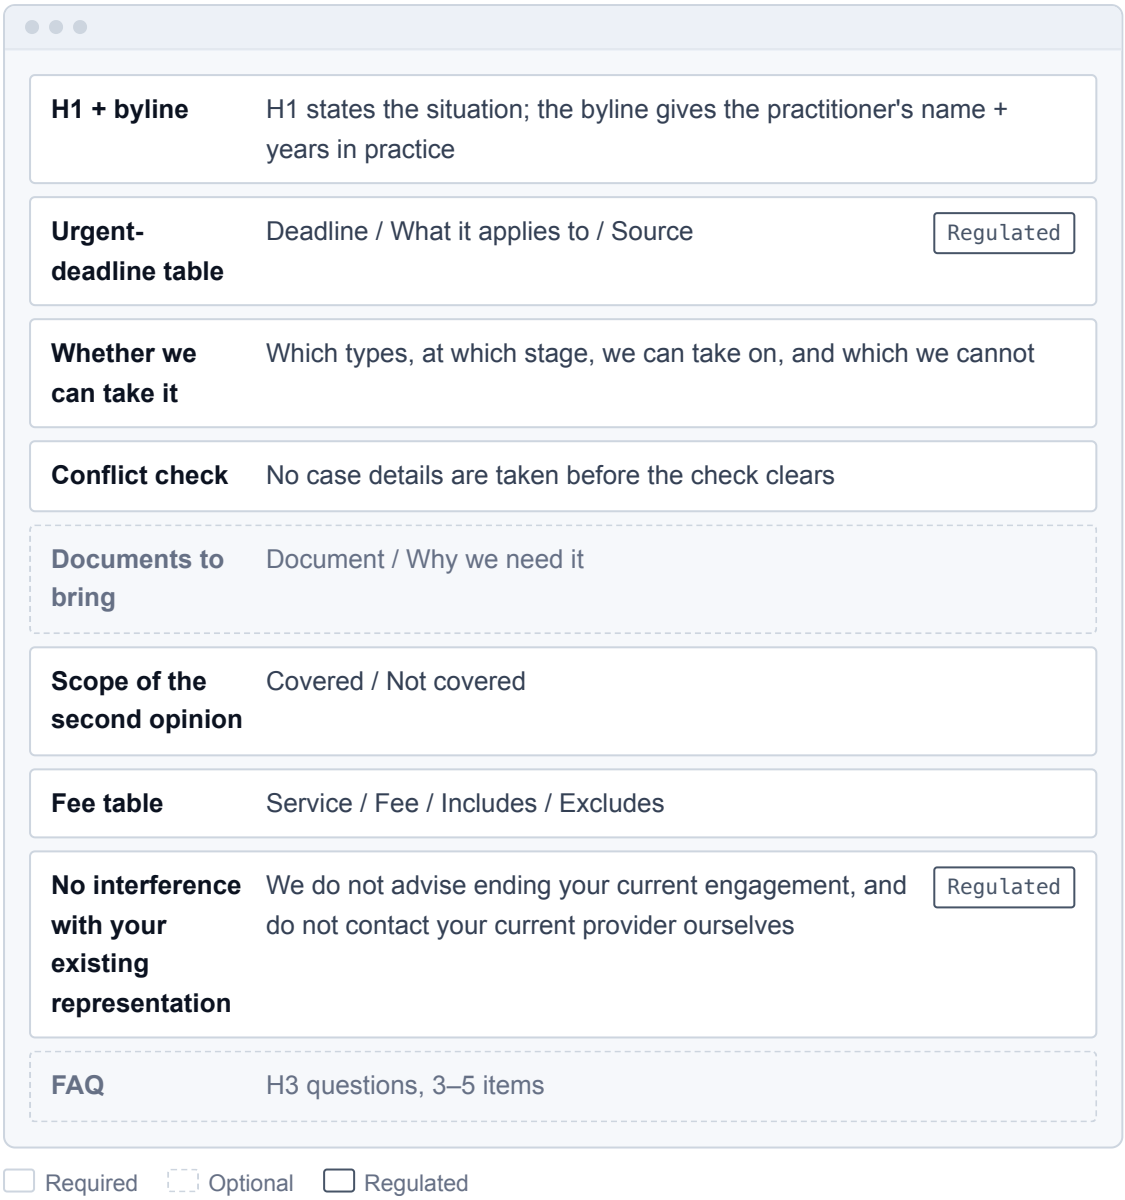


Figure: rights-deadline remedy page, deadlines first, then pathways (not measured as a whole, so no block is marked cite)

Every row of the urgent-deadline table cites the original text of the regulation or an official page, with the date retrieved.

Standard column headers: red-flag table `If you notice / Do this / How soon`; pathway table `If the assessment finds / Possible pathway / Who does it`; deadline table `Deadline / What it applies to / Source`.

Sample sentences: own sentence: "At <Own business>, a review with <assessment type> is S$X (incl. GST); any further fee is confirmed in writing after that examination." Market sentence: "<Situation> is usually managed by <Pathway A>, <Pathway B>, or <Pathway C>, depending on <Variable A>, <Variable B>, and <Variable C>."

Negative conditions in the H2, dual bylines and attached References combine into one general rule: write "when this isn't a fit" into the H2 heading itself, not buried in a paragraph; give both the author's

and the reviewer's names; and attach References at the end of the page.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type [Safety]: no promotional CTA such as a free assessment follows the red-flag section; a professional reviewer's written sign-off is required before launch. That sign-off covers the accuracy of the professional content, not compliance clearance: a written sign-off never turns something banned as Statute text into something publishable; see → General Edition 0.4 The three labels for compliance sentences, and the stop-and-escalate rule.

Most common mistake [Safety]: leading with the cause and letting the reader self-diagnose. That is a safety problem, not a matter of layout preference. A layout mistake: multiple H1s on one page, tagging the summary box, the FAQ and the CTA all as H1.

5.20 Entity anchor pages: three subtypes, and the product detail page (pt09)

What you'll do in this section: sort out which intent each of the three entity anchor subtypes serves. The blueprints for the organisation facts page and the person profile page were already built in chapter 3, so here we only add the product detail page: the first-screen buy box, the feed, structured data and the checkout price must all agree, and specs are written as an HTML table. When you're done, you'll have one more skeleton, for the PDP, plus a rule for deciding which of the three subtypes a question goes to.

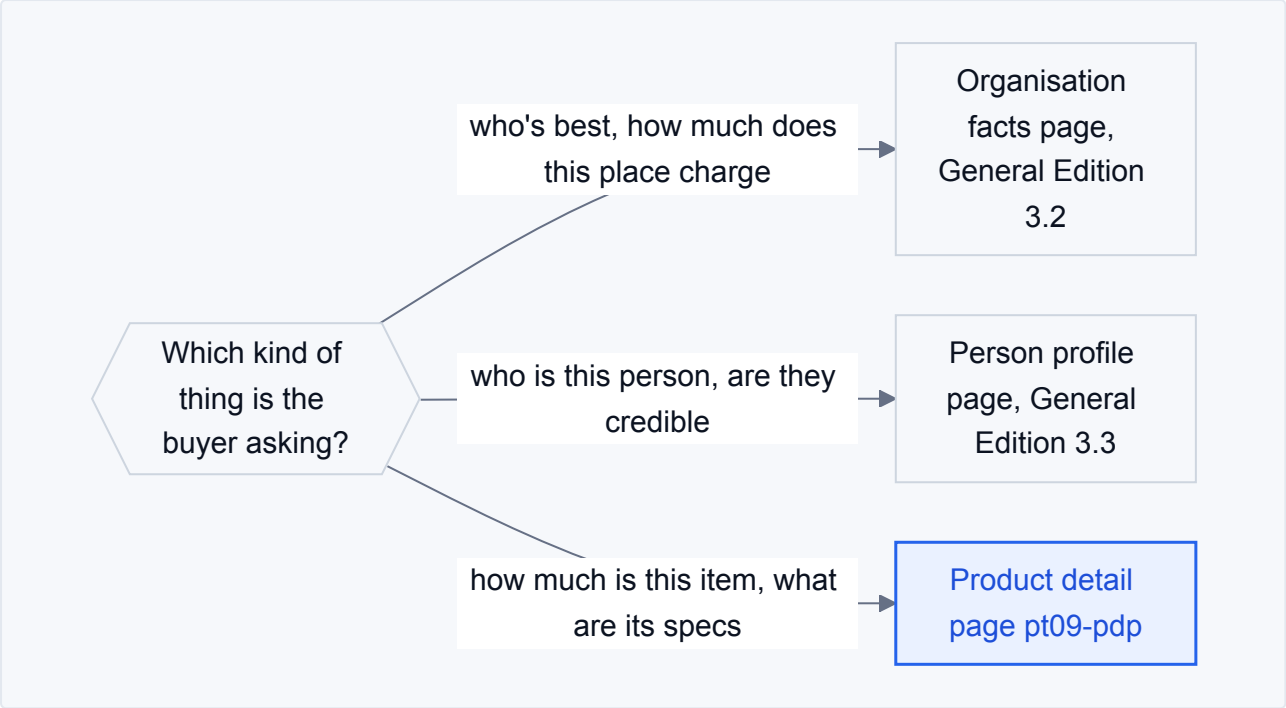


Figure: each subtype serves one kind of intent; don't use the wrong skeleton

For the blueprint of the organisation facts page (pt09), see → General Edition 3.2 The fact table and the organisation facts page; for the person profile page (pt09-person), see → General Edition 3.3 Person pages and six-trace alignment. Neither is redrawn here. On pages related to personal safety, the first line under the H1 is always the named reviewer; for how to write this line of titles and credentials, see → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.

Product detail page (pt09-pdp)

Intent: "best X under $Y", "X 多少钱" (how much is X), "where to buy". Which leg: **ChatGPT only**; AI Mode relies on the shopping card generated from the Merchant Center feed, not on the body text (evidence: see → General Edition A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types).

H1: the product's full name + key specs, on one line.

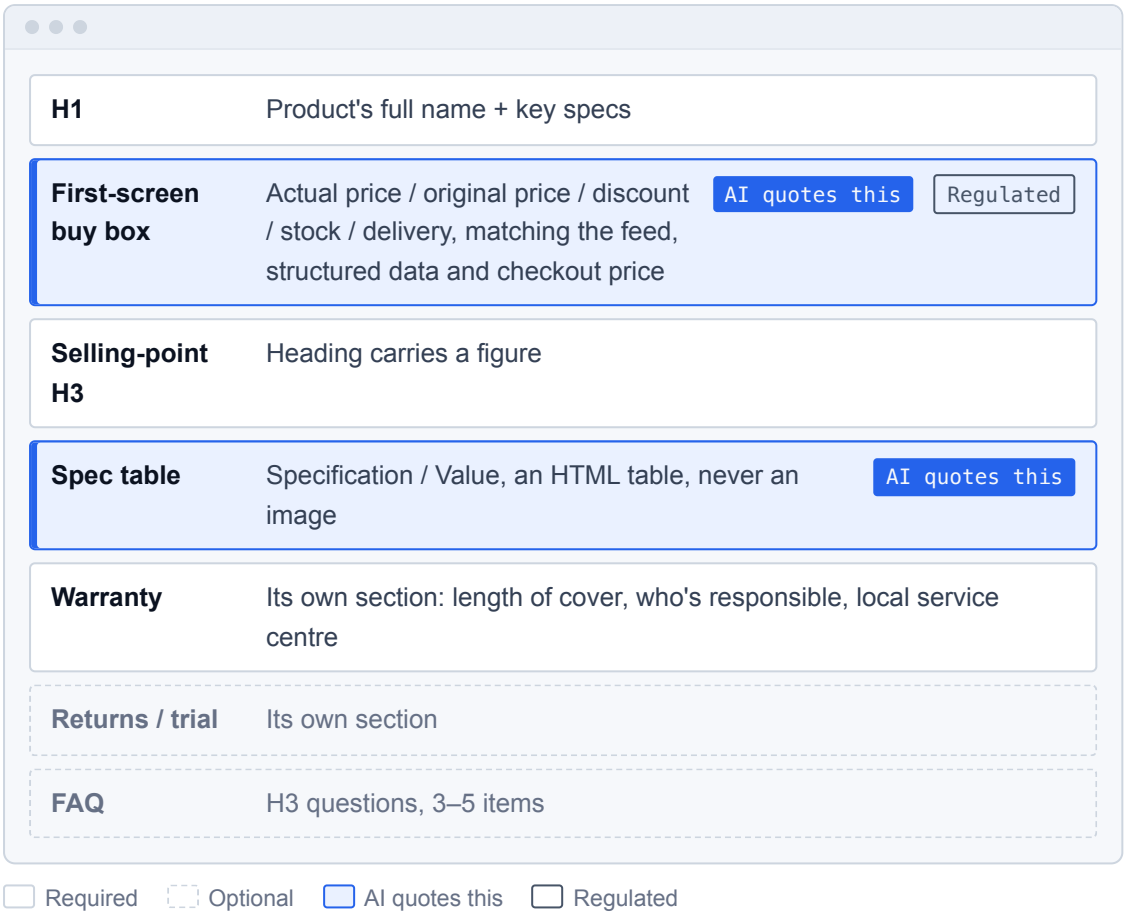


Figure: product detail page; AI copies the first-screen buy box and the HTML spec table

The Product structured data fills in only the price and stock visible on the page; the original price and the discount must be real.

Standard column headers: `Specification / Value`, with rows set by category (such as capacity, battery life, weight, warranty length, local service point).

Sample sentences: own sentence: "The <Model> costs S$X (incl. GST) with <key spec 1>, <key spec 2> and a <N>-year local warranty." Market sentence: "<Category> under S$X in this market generally offer <spec range A> and <spec range B>, depending on <variable>." Fill in the market sentence from current peer data, with the price-check date. The unregulated side writes it as it is; the light side (switch E (peer comparison banned) is not on) writes it only in a separate fact section with its source (Conservative line (not statute text); see 5.2); the strict side, and anyone with switch E on, do not write it.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.

- Specific to this page type: the price must be the same in all four places (the page, the feed, the structured data and the checkout). Once a price is out of date, take it down; never leave an old price up.

Most common mistake: the H1 is a slogan, the first screen carries only promotions, and the hard facts (price, specs) are hidden in JS or images, unfindable in the static HTML. A big brand can still get cited on the strength of its authority; a weak domain written the same way gets skipped.

Chapter 5

Writing pages (3): quick cards for 33 more page types

Subsidies, reviews, FAQs, step-by-step guides, branch pages, legislation and more

5.21 Tier-B price types: subsidy limits, parameter basis, category lists, calculators, time-limited promotions

5.22 Tier-B selection and reputation types: third-party reviews, sentiment and forums, review aggregates, self-built reputation, verdict-first

5.23 Tier-B questions and rules types (1): regulatory obligations, definitions

5.24 Tier-B questions and rules types (2): procedure steps, preparation lists

5.25 Tier-B questions and rules types (3): schedules, collected FAQ, policy hubs, change notices, misconceptions

5.26 Tier-B entity and product facts types (1): branches, verification pages, directory listings

5.27 Tier-B entity and product facts types (2): compliance proof, terms, how-to docs and integrations

5.28 Tier-B primary sources (1): legislation, regulator guidance, official replies, device documents, case law

5.29 Tier-B primary sources (2): statistics sources, registers

5.21 Tier-B price types: subsidy limits, parameter basis, category lists, calculators, time-limited promotions

What you'll do in this section: give each of these question shapes its own page: "how much do I pay after the subsidy" / "how much would I pay in my situation" / "which brands are there and what price range" / "how is it calculated". When you're done, you'll have five blueprints, each with the numbers AI cannot work out for itself (the price you actually charge, the unit price, the effective period, whether tax is included) fixed in the first screen.

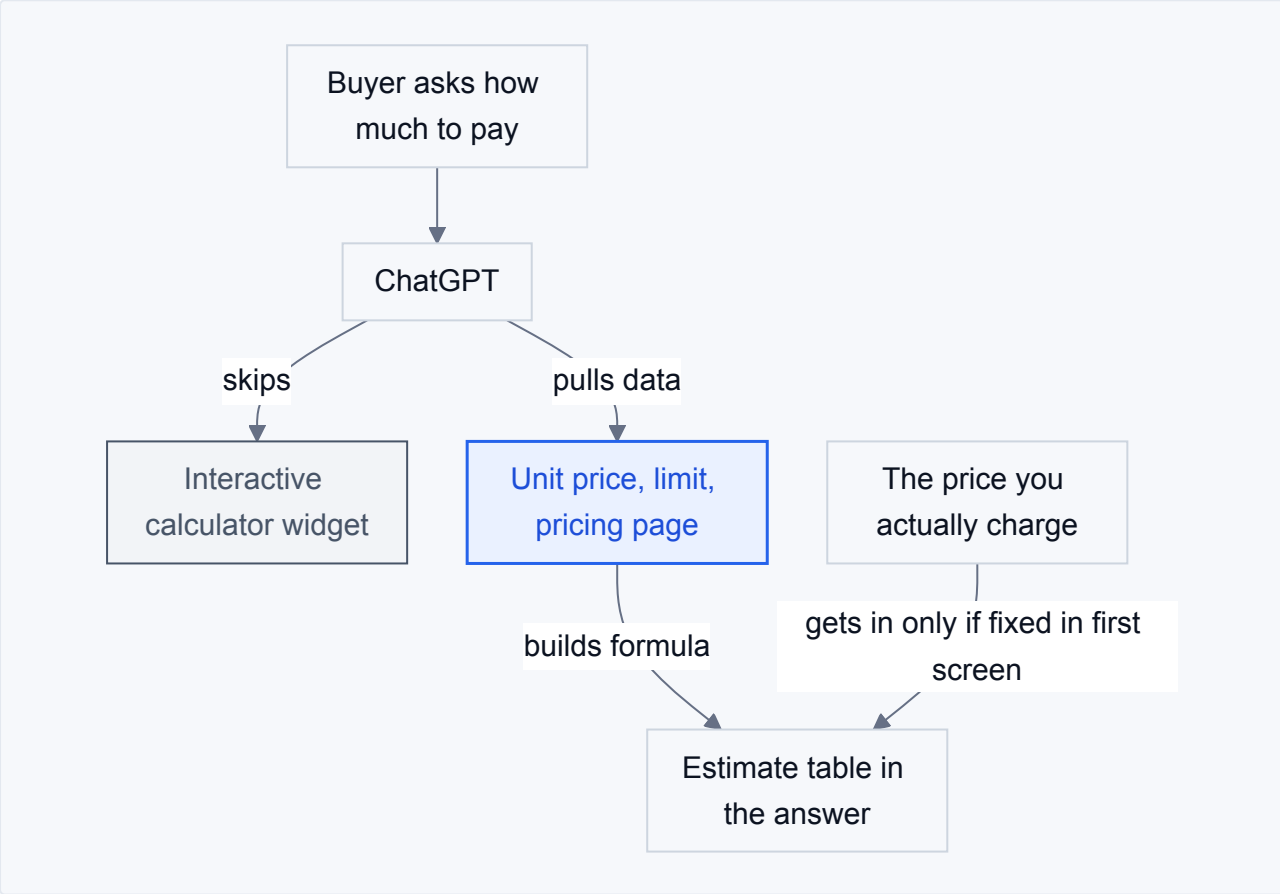


Figure: AI skips the calculator, does the maths itself from the unit price and the limit, and only lacks your own number

This pattern recurs across three industries (evidence in Appendix A.4). Each of the five types below supplies one cell of this estimate table; the cards do not repeat this point.

5.21.1 Subsidy and limit rules page (pt17)

Intent: "How do I apply, how do I claim" and "how much do I pay after the subsidy" share this type. **Leg**: ChatGPT (AI Mode not measured). **title and H1**: H1 = the scheme's full name with its abbreviation; a short summary on the line under H1, with `Last updated <day month year>` in the first screen. This type wins on number density, not on images.

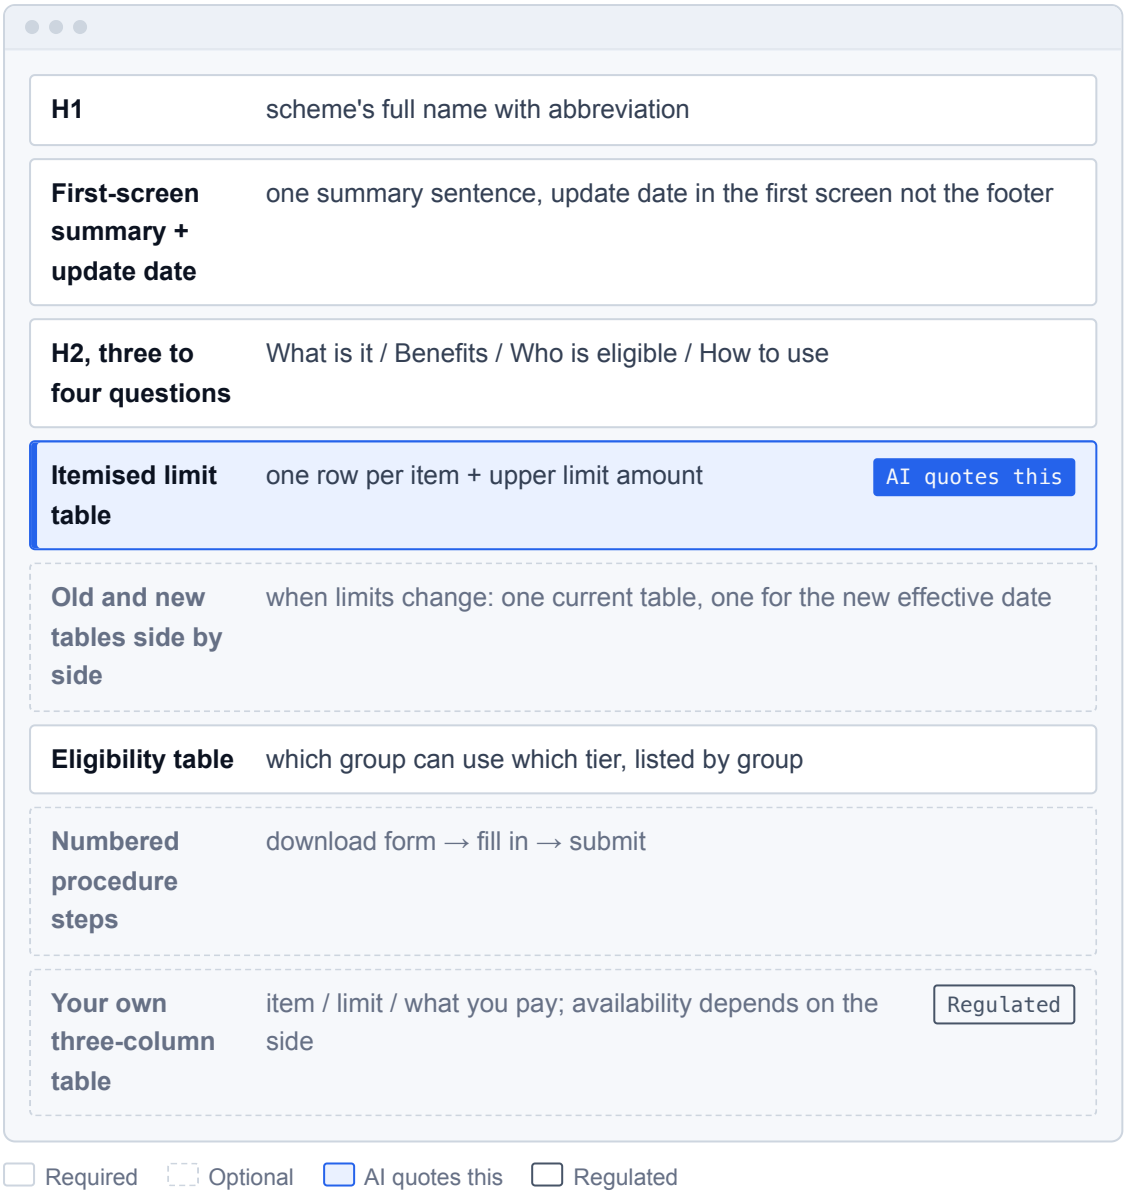


Figure: Subsidy and limit rules page; AI copies cells from the itemised limit table

Standard column headers: official table `S/N / <service category> / Claim limits / Current Subsidy amount (Up to $)`; your own version `Item / Limit / What you pay`. **Sample sentences**: own sentence (holds only on the lightly regulated and unregulated sides) — "<item name>: we charge S$<final tax-inclusive price>; with us, this item's <subsidy scheme> code is <code>, which offsets S$<amount>; you pay S$<amount>." Market sentence — copy the official limit as is, label the table number and update date, and give it its own section.

> **Example** (healthcare subsidies) The sentence giving the limits (up to $830 per day, $240–$5,290 for surgical items; paraphrased) was lifted whole from the official subsidy page; the source is a cell in the itemised limit table. AI then combines it with the fee benchmark and the "can I claim" rule to do the maths on the spot. Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: your own three-column table and the "you pay" sentence hold only on the lightly regulated and unregulated sides; on the strictly regulated side your own version may be built after rewriting: the subsidy stays in its own factual section, in a different table from your own price, not adjacent to it, and there is no case-by-case confirmation tier. See the price row and the dedicated price-family row in 5.2 for the wording.

Most common mistake: building a subsidy calculator. What you need to give is the number AI cannot calculate: the price you actually charge and your own coding result.

5.21.2 Parameter and rate basis page (pt27)

Intent: "how much would I pay in my situation", "X calculator", "how much a month". It is used as a calculation input, not for comparing prices — that is what sets it apart from → General Edition 5.14 Official pricing page (pt03). **Leg**: ChatGPT (AI Mode not measured). **title and H1**: both forms get cited — the original-question form (`What is the <parameter name>?` , one question one answer), or the category + location form (H1 without a year, without "calculator").

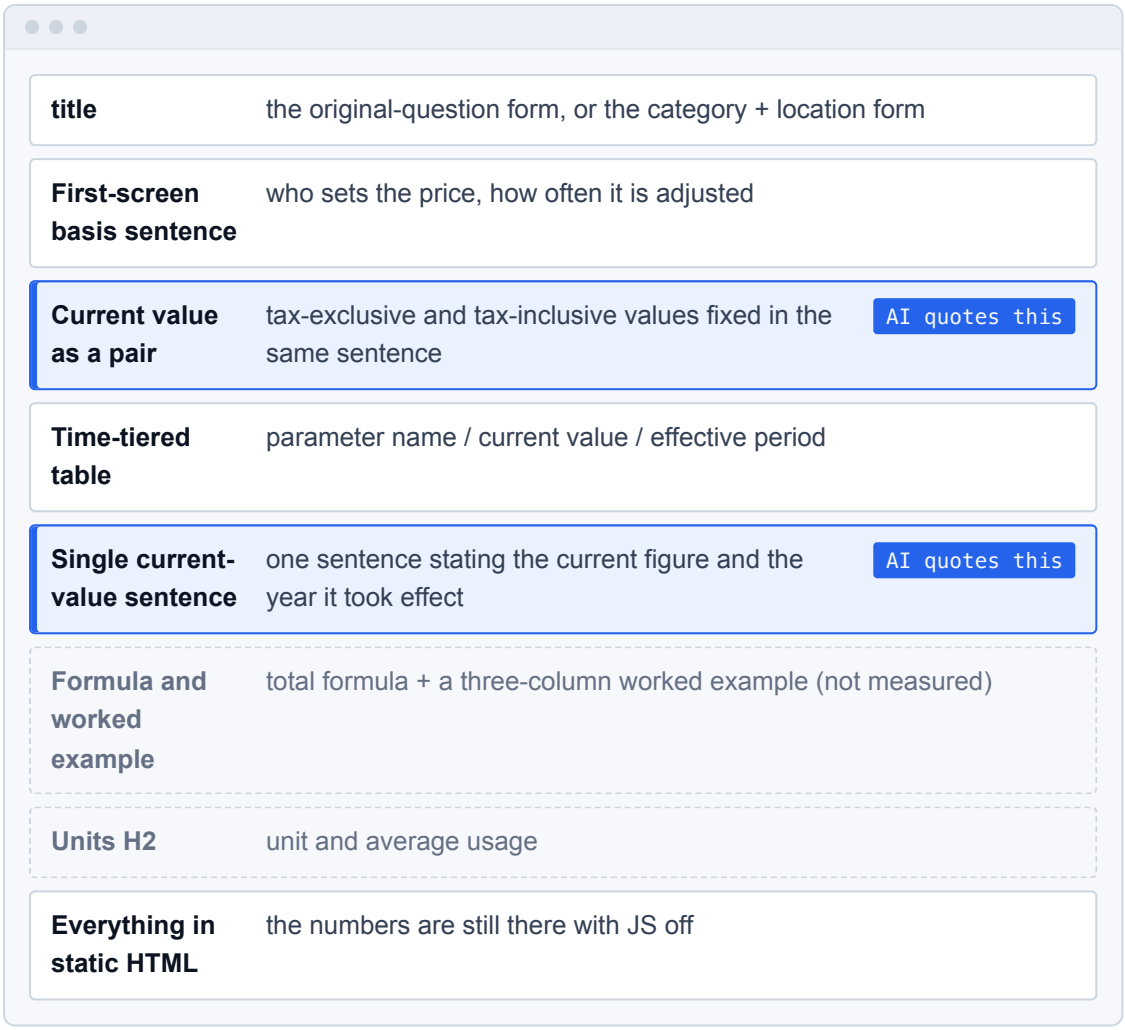


Figure: Parameter and rate basis page; AI copies the current value given as a tax-exclusive and tax-inclusive pair

Standard column headers: `<parameter name> / <annual cap or unit price> / effective period`, one row per tier. **Sample sentences**: own sentence — "<parameter name> has been <value> (<tax basis>) since <effective date>, adjusted every <period>." Market sentence — write out the formula: "<total> = <parameter A> × <parameter B> × <parameter C>".

> **Example** (payroll contributions) Two things were lifted from the official wage-ceiling page: the sentence "It was gradually raised and has reached $8,000 in 2026." and the last row of the tiered table. What AI wants is an authoritative source for each parameter in the formula. Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify

the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.

- Specific to this page type: none. For how to turn a parameter into a final tax-inclusive price, see the dedicated price-family row in 5.2.

Most common mistake: spending effort on an interactive calculator page. The value, the basis and the effective date are all required — fix all three in static HTML in the first screen.

5.21.3 Category list page (pt37)

Intent: the "which brands are there, what's the price range" part of a full-guide question. **Leg**: ChatGPT (AI Mode not measured). **title and H1**: title = `Shop <synonym A> & <synonym B> for <audience>` + site name. H1 can be a selling-point sentence; it is not the part that gets cited, so don't spend effort on it.

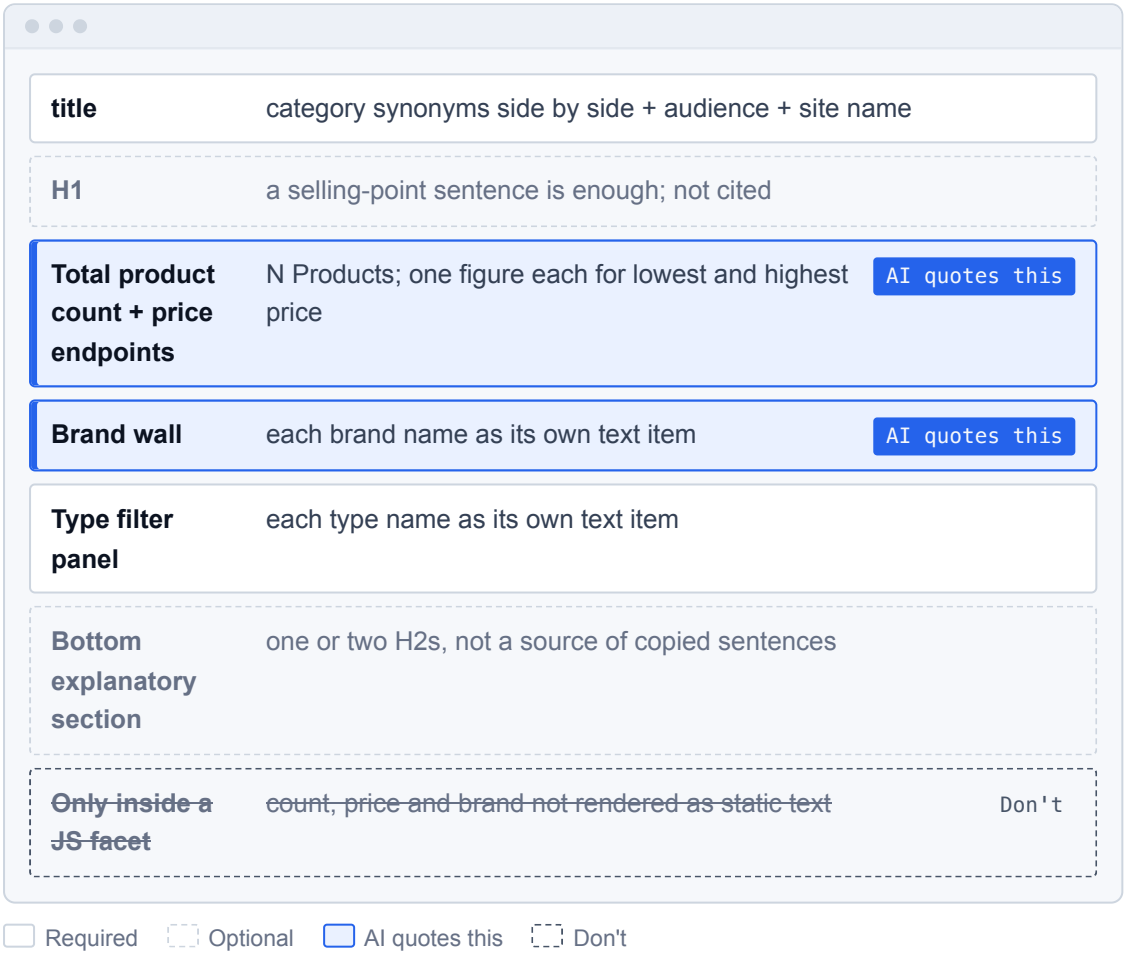


Figure: Category list page; AI copies the total product count, the price endpoints and the brand wall

Standard column headers: this type relies on plain text in the filter panel, not a table; if you add one, use `Brand / Price range / Number of models on sale` (not measured). **Sample sentences**: "<category> on <this site> has <N> models in total, priced from S$<lowest> to S$<highest>, covering <brand 1, brand 2…>." This sentence is also the market sentence — AI takes it as the size and price range of the category, one of the few page types where the two sentences merge into one.

> **Example** (mother-and-baby e-commerce) A stroller category page was copied whole into a market overview sentence: 169 models, about S$112 to S$3,045, nine brands. A comparable page on another site rendered the whole thing in JS, with 0 words of static HTML body text; it was listed as a citation but not one sentence was actually copied. Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: on the strictly regulated side and where switch E (peer comparison banned) is on, do not build this (no listing of peers, no side-by-side line-up across the category); build the organisation's own service catalogue instead. On the lightly regulated side it may be built after rewriting, as the organisation's own catalogue. Where switch B (legally required fields) is on, every item's name in the catalogue must match the registered name word for word.

Most common mistake: the product count, price endpoints and brand list exist only in the facet's JS. If they don't render as static text, being cited counts for nothing.

5.21.4 Calculator page (pt38)

Intent: "How do I calculate my X", "which tier does <score> fall into", "X calculator". **Leg**: ChatGPT (AI Mode not measured). **title and H1**: title = H1 = `<tool name> <year> – <what it calculates>` + site name.

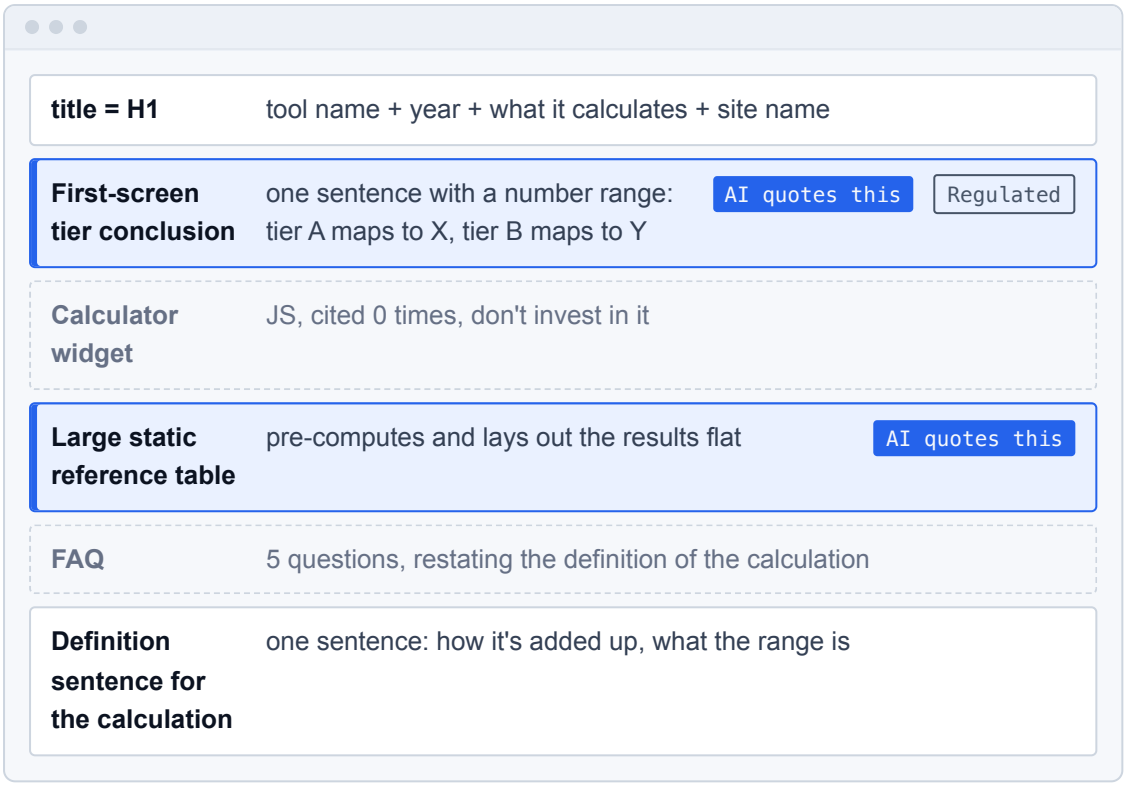


Figure: Calculator page; what gets copied is the first-screen tier conclusion and the static reference table, not the widget

Standard column headers: `<input tier> / <result> / <conditions it suits>` (cross-industry analogy, not measured). **Sample sentences**: own sentence — "<input> under <basis> maps to <result>; <result range A> applies to <scenario A>, <result range B> applies to <scenario B>." It must carry a number range; never write "it varies from person to person". Market sentence — write the public distribution under that basis as one sentence, labelled with the data date.

> **Example** (school admissions) Three things were copied from a score-calculator page: the first-screen tier conclusion, the body rows of the reference table, and the sentence defining the calculation (4 subjects added up, range 4–32); the calculator widget itself was not copied. Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: on the strictly regulated side the calculator's output must never be a price commitment — it may only estimate process (number of visits, duration); the static table should only

carry fixed-price items. Where switch E (peer comparison banned) is on, write "subject to formal engagement", and do not compare prices with peers.

Most common mistake: building a beautiful JS calculator and nothing else on the page. Handover test: turn off JS — the page should still have a tier conclusion sentence and a static table left.

5.21.5 Time-limited promotion page (pt42)

Intent: when AI answers "how is this business", it goes looking for the business's actual current price. **Leg**: ChatGPT (AI Mode not measured). **title and H1**: the page body itself could not be retrieved and can only be worked out backwards from the URL, so we give no formula for title and H1.

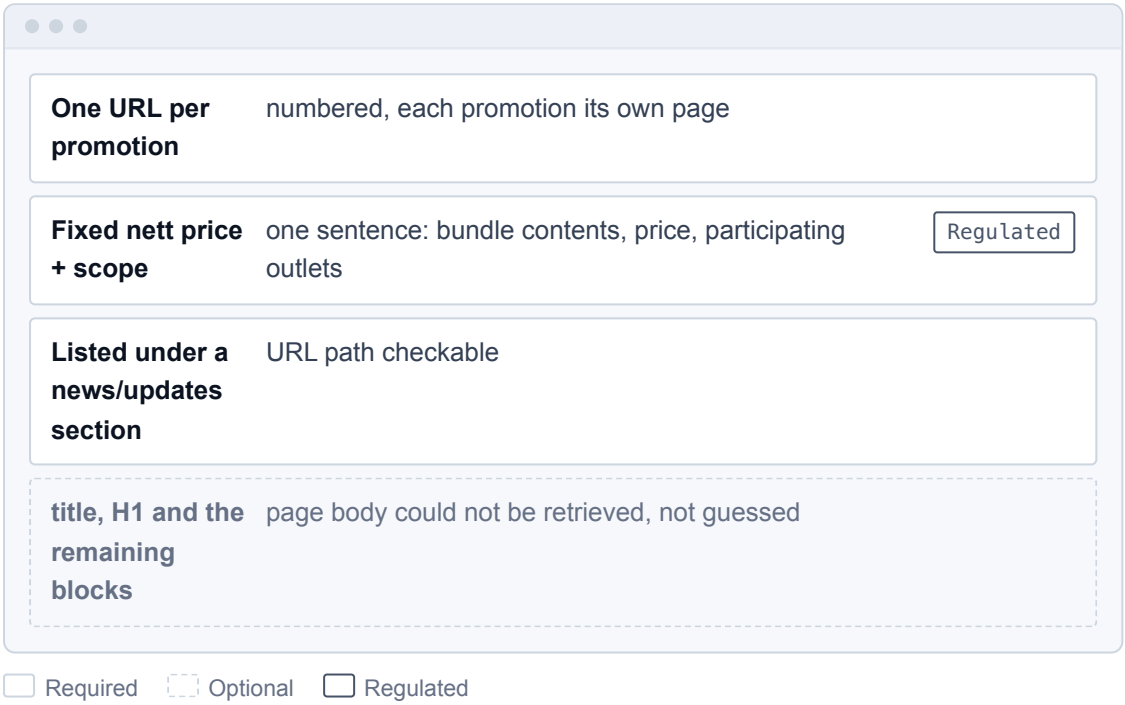


Figure: Time-limited promotion page; only three things were verified, nothing else was guessed

Use all three as clues, not as conclusions. In the one measured question, this type of page was the only citation from the business itself, with a price, that AI wrote into the answer as a positive fact; the homepage only contributed a neutral fact — the number of branches (evidence in Appendix A.4).

Standard column headers: none (layout not retrieved). **Sample sentences**: own sentence (the shape of the cited sentence) — "<item A + item B + item C> S$<price> nett, valid at <participating outlets>." No market sentence applies.

The cited promotion wording	The strictly regulated side's replacement wording
One URL per promotion, listed under news/updates	One URL per bundle, standing
<bundle> S$<price> nett, time-limited	S$<final tax-inclusive price>, including <mandatory items>
Valid at participating outlets	State which outlets it applies to
Carries a validity period and promotional wording	Drop the time-limited and promotional wording

Figure: the cited promotion wording, and how the strictly regulated side replaces it with a standing fixed price in the same spot

This kind of page as cited cannot be built as-is on the strictly regulated side — what wins rankings does not decide compliance; the rules and labels for the promotion cell are in the promotion row of 5.2.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: where switch B (legally required fields) is on, write the discount in full: validity period, conditions, currency, whether tax is included, payment method.

Most common mistake: cramming every promotion into one "latest offers" list page. With ten promotions listed on one page, AI cannot tell which price maps to which item.

5.22 Tier-B selection and reputation types: third-party reviews, sentiment and forums, review aggregates, self-built reputation, verdict-first

What you'll do in this section: work out, inside the answer to "is this business good", which page type each of the four blocks — neutral facts, pros, cons, and the scored rating — comes from. When you're done, you'll have five blueprints, and you'll know which ones you can build yourself and which ones you can only nudge into position by supplying facts.



Figure: the answer to "is this business good" splits into four blocks; the cons block comes only from pages you don't control

For pt25, marked "cannot build" in the figure, on every side you can only supply checkable facts, make sure you are listed on the platform, and keep the platform information matching your own site. pt28, the review aggregate page, has different availability by side: on the strictly regulated side and where switch E (peer comparison banned) is on, you may not self-build it (a rating is a testimonial, Statute text); on the lightly regulated side it may be built after rewriting (testimonials carry a name, relationship and year); on the unregulated side it may be built. Only the unregulated side can self-build the review page: you can write a review of a competitor and cast yourself as the alternative, to take that seat. The verdict-first page as a whole type is only built on the unregulated side; on the lightly regulated side, after rewriting, only the scenario comparison table (5.22.5) may be built. The self-built reputation page only

supplies the positive half, and is not built on the strictly regulated side. For whether each side may write testimonials and how to write a rating, see → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions; not repeated here.

5.22.1 Third-party single-business review (pt24)

Intent: "is <a business> good", "<a model> review", "is it worth it". In the same question, ChatGPT gets specs and price from the official PDP; subjective pros and cons can only come from this type. **Leg**: ChatGPT (AI Mode not measured). **title and H1**: the media version is a first-person verdict sentence, a long headline, naming competitors and stating reservations; the institutional version is written as "<institution> vs <alternative>". The line under H1 has the byline + publish date — this type has the highest byline rate of any type in the book.

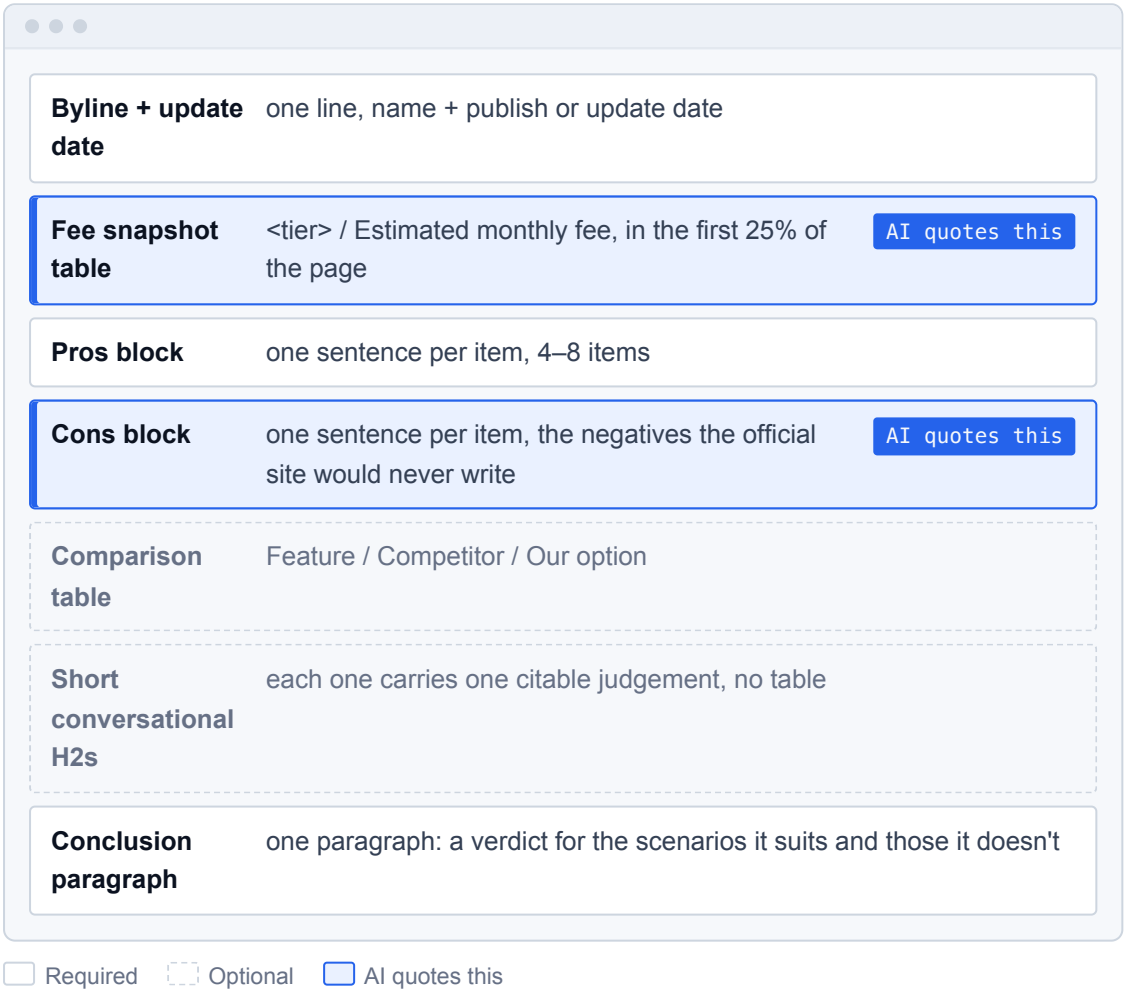


Figure: Third-party single-business review; AI copies the fee snapshot and the Cons

Standard column headers: `<tier> / Estimated monthly fee (<year>)` (fee snapshot table); `Feature / <competitor> / <own>` (comparison table, unregulated side only). **Sample sentences**: this type is written by a third party, so there is no own sentence; the market sentence is the Cons item itself, shaped like "trigger operation can be annoying — you must hold the trigger down" or "motor-heavy design feels awkward in handheld mode". ChatGPT attributes the Pros to the official site and the Cons to this type.

> **Example** (education and training) The fee snapshot table of a third-party review page (Sec 3–4 S$360–420/month, JC H2 S$400–480/month) and its Pros/Cons wording were copied, from a position in the first 25% of the page. Evidence in Appendix A.4.

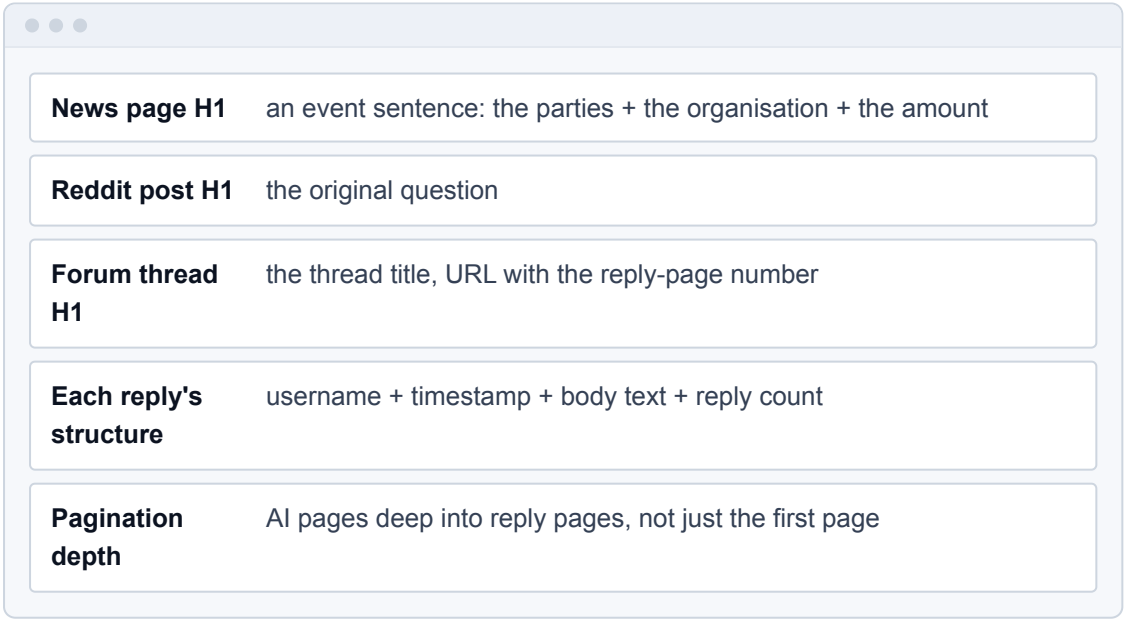
Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: not built on either the strictly or lightly regulated side (Conservative line (not statute text)); not built where switch E (peer comparison banned) is on. It may only be self-built on the unregulated side, where you can cast yourself as the alternative in a review of a competitor, capturing the question that names the competitor's brand.

Most common mistake: putting the rating and conclusion entirely in JS. With JS off, the body text is nothing but navigation — listed as a citation, but with not one sentence to copy.

5.22.2 Sentiment, forum and news pages (pt25)

Intent: the negative side of "is this business good". You cannot build this type, but it sets the ceiling for your own site's pages. **Leg**: ChatGPT (AI Mode not measured). **Carriers**: news page H1 = an event sentence, amounts down to the cent; Reddit post H1 = the original question, with multiple replies; forum thread URL paginated by reply page, H1 = the thread title, each reply has username + timestamp + body text, no rating system.



News page H1	an event sentence: the parties + the organisation + the amount
Reddit post H1	the original question
Forum thread H1	the thread title, URL with the reply-page number
Each reply's structure	username + timestamp + body text + reply count
Pagination depth	AI pages deep into reply pages, not just the first page

Figure: the three sentiment, forum and news carriers (observed; cannot self-build)

All three carriers are observed (some pages could not be retrieved; read only from the cited sentence) — **you cannot build this type**. For the forum part, AI summarises; it does not copy a single sentence.

> **Example** (dental) Of the 5 citations in an answer to "is this business good", 2 were sentiment sources: a news article was copied into an account of a bill a woman was given after a scaling appointment (the amount down to the cent), and a Reddit post was copied into "some users say [it] was okay; others complain about rushed cleaning, high bills, or upselling." AI added its own line, "treat those as anecdotes, not proof", but still wrote it into the body of the answer. Evidence in Appendix A.4.

Defence (the only three things you can do):

1. Lay out checkable facts (price list, nett price, credentials, branch opening hours) so AI finds them on the first search, crowding out its motive to go digging through forums. Brand-review questions are not won by adding more pages.
2. Respond to common complaints directly on your own page, so that when AI restates a complaint it may cite your response. Note This is inferred, not measured.
3. Recognise that your own site only supplies neutral facts such as branch count — the half-sentence in an answer that calls the reputation mixed has only this one source; neither your own site nor review pages can supply the negative, and the two sentences that set the tone come from pages you don't control.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: on every side this type cannot be self-built; you can only supply facts and make sure you are listed on the platform with information matching your own site. On the strictly regulated side, you may not respond with comparisons or mobilise testimonials to counter it.

Most common mistake: paying for brand-polishing content. In the measured case, not one of these made it into the answer; what made it in was checkable numbers and the forum's own words.

5.22.3 Review aggregate page (pt28)

Intent: "is <an organisation> good", "how's <brand>", "reviews from parents". This type explains where the half of citations you don't control comes from in "who's best" questions — it can be generated programmatically, one URL per branch the organisation has. **Leg**: ChatGPT (AI Mode not measured). **title and H1**: title puts the review count in the title itself (`<organisation (branch)> – N Reviews – <category> – <platform>`). H1 = the organisation name with a branch suffix.

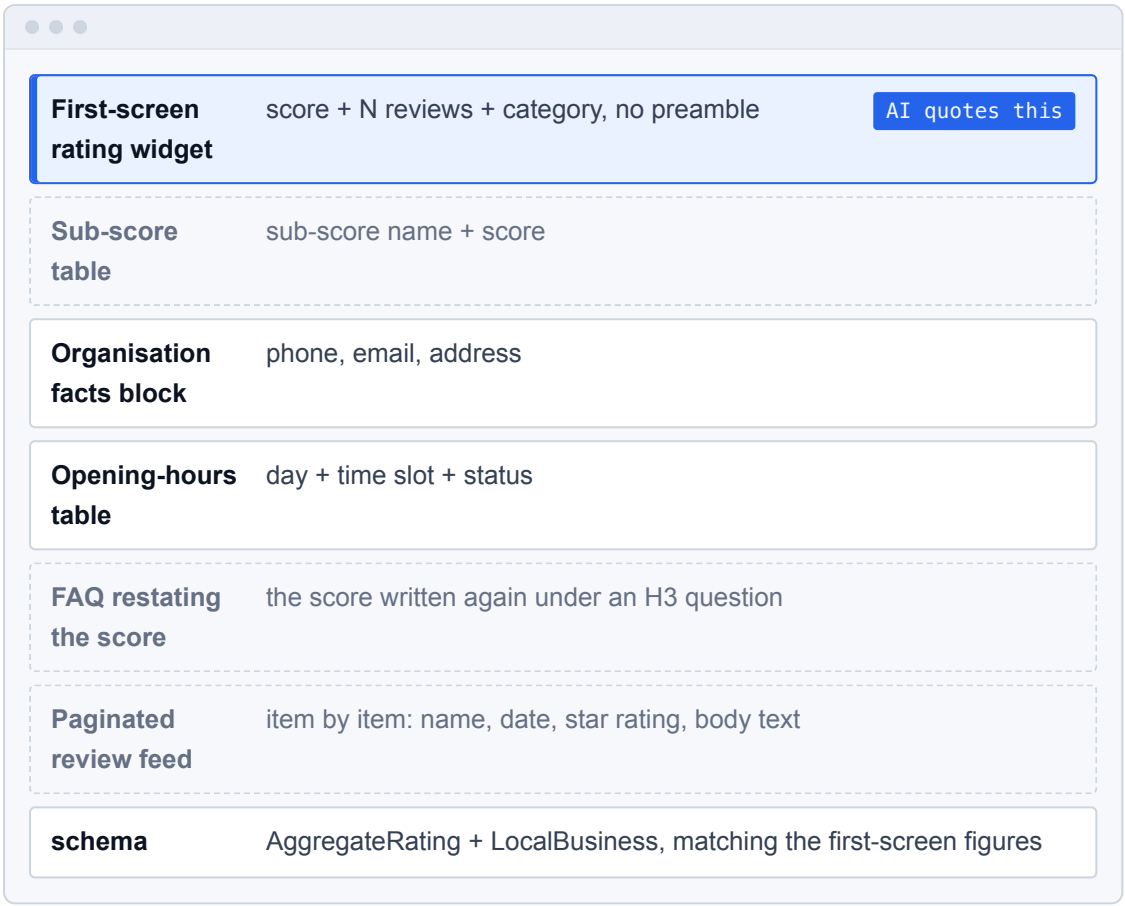


Figure: Review aggregate page; AI copies the aggregate value with its denominator, not the review text itself

Standard column headers: `Weekday / Schedule / Status` (opening hours). **Sample sentences**: does not apply on the strictly regulated side or where switch E is on. Market sentence (the shape AI wants) — "<organisation> (<branch>) <rating>/5, <N> reviews", the score plus the sample size; the star rating alone is not used.

> **Example** (dental) The first-screen "4.2/5 with 113 reviews" was copied whole, and copied side by side across three branches within the same answer. Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: not self-built on the strictly regulated side or where switch E (peer comparison banned) is on (a rating is a testimonial, Statute text); on the lightly regulated side it may be built after rewriting, with testimonials carrying a name, relationship and year; what can be done is

limited to making sure you're listed on local directories and review platforms, that reviews carry dates, that each branch has its own URL, and that the address and phone number on the platform match your own site.

Most common mistake: thinking that putting a positive review on your own site offsets this. AI wants a third-party aggregate value with a denominator, which your own testimonial page cannot supply.

5.22.4 Self-built reputation and credentials page (pt29)

Intent: "is <a business> good". This is the only type you can self-build for review questions. **Leg**: ChatGPT (AI Mode not measured). **title and H1**: the title carries a ranking claim, or is written as "Award Winning <business> | <brand>". Numbers must go into an H2 or a stand-alone sentence, not scattered through a paragraph.

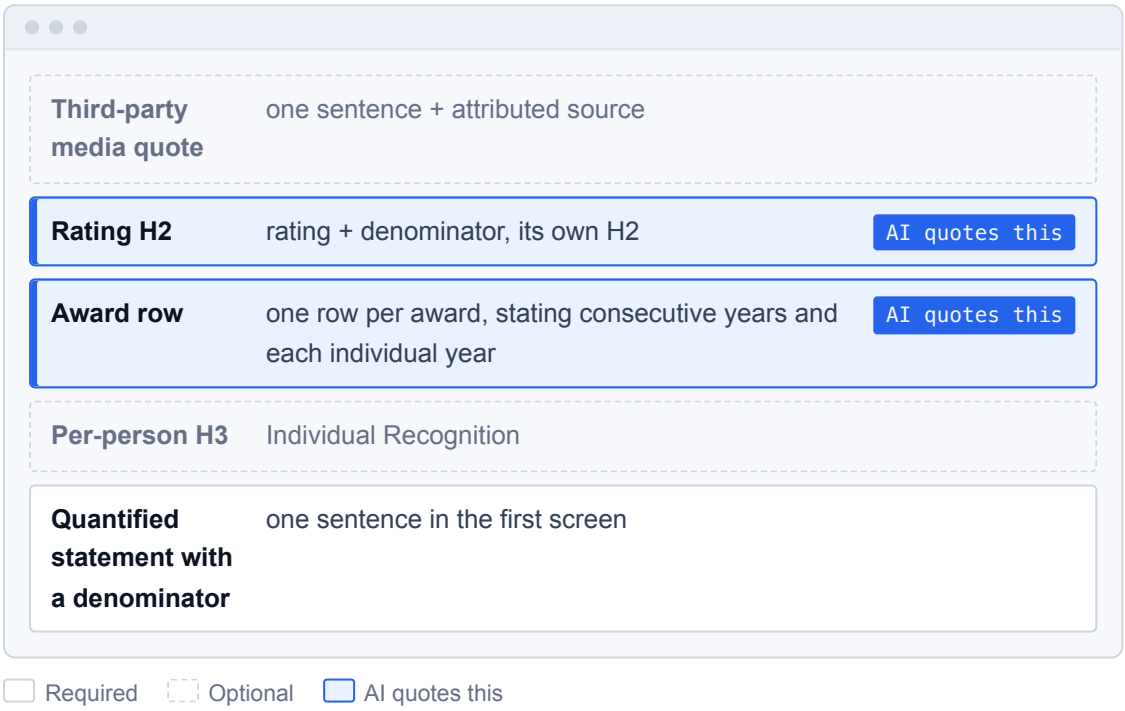


Figure: Self-built reputation and credentials page; AI copies the rating sentence with its denominator and the string of award years

Standard column headers: no table; awards and ratings are carried by an H2 and a single row. **Sample sentences**: own sentence — "<N> reviews on <platform name>, <rating>/5 (<check date>)." "<award name> — <judging body>, <N> consecutive years: <string of individual years>." Market sentence: none — this type does not do side-by-side comparison.

As long as the rating on your own page carries a denominator and the source platform, AI accepts it directly and notes the source as your own page. But it only takes the positive half from this type; it gets the negative from review pages and forums — a self-built reputation page's role is real but limited.

> **Example** (legal) The H2 on a law firm's self-built reputation page was copied for the sentence "4,000+ Google Reviews | 4.9★ Rating", and the award row was copied for the string "Ranked for 13 consecutive years: 2015 to 2026". Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: not built on the strictly regulated side (testimonial star ratings, Statute text; do not write success rates, do not write Leading / Expert / best, do not include testimonials or star ratings); on the lightly regulated side it may be built after rewriting, with testimonials carrying a name, relationship and year. Where switch E (peer comparison banned) is on, only provable expertise is allowed (qualifications, years in practice, share of business, a string of third-party ranking years); do not write success rates, and do not compare fees or service quality with peers.

Most common mistake: piling up 150 testimonials. AI only copies one or two quantified statements with a denominator; without a "Y out of X + year + basis" sentence, the whole page counts for nothing.

5.22.5 Verdict-first page (pt41)

Intent: a mixed question type combining a single-product review with price — "<model> review, is it worth it in <location>", supplying both "is it worth it" and "how much, where's it cheapest" at once. **Leg**: ChatGPT (AI Mode not measured). **title and H1**: title = H1, with a year and a location, packing three intents into one title. The line under H1 has byline + role + date, then a line of dek with the key number fixed directly in it — this is where the copied sentence lives, at the 1–2% mark of the body.

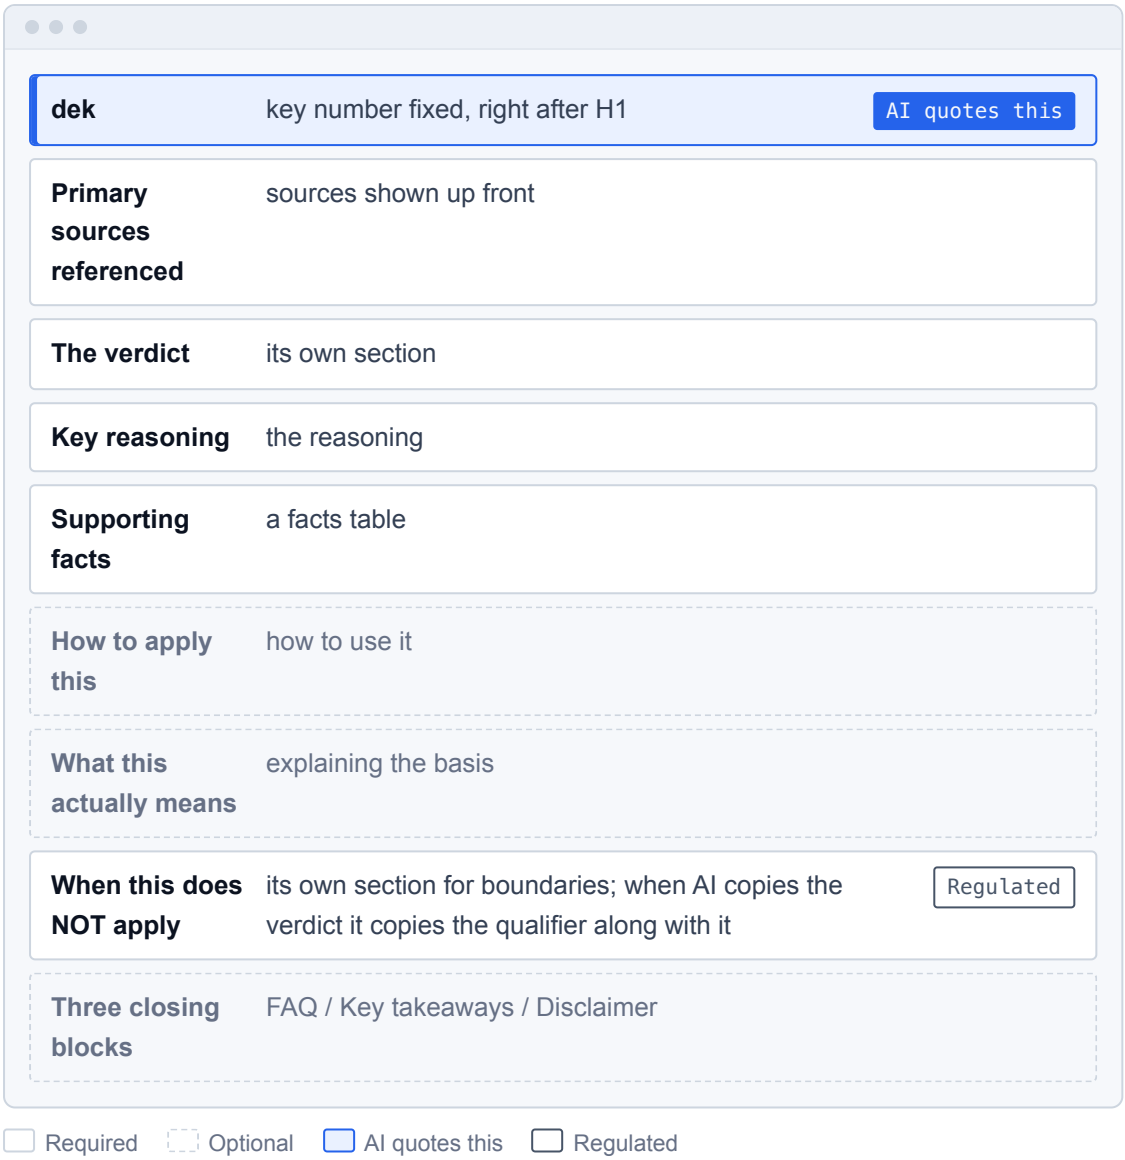


Figure: Verdict-first page; the verdict, the reasoning and the boundary each get their own section

Standard column headers: `Variant / Retail <currency> / Ideal for / Key feature` (variant comparison); `Channel / Cashback / Warranty / Instalment` (channel comparison); `Scenario / Recommendation / Reason` (scenario comparison). **Sample sentences**: own sentence — "<model> <variant A> reference retail price S$<X>, <variant B> S$<Y> (<price check date>); <channel> also has <condition>." Market sentence (this type's unique value) — "the price gap for the same model between <official site / promotion / grey market> is S$<X>–S$<Y>, driven by <warranty / variant / cashback>."

> **Example** (appliance e-commerce) The static HTML of a model's official PDP had no retail price, only specs; a third-party content site wrote the reference retail prices for two variants into its body text, and AI copied it as the price source, attributing the source to this content site rather than the official one. Evidence in Appendix A.4; see also → General Edition 5.20 Entity anchor pages: three subtypes, and the product detail page (pt09).

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: the strictly regulated side cannot build this type as a whole (at the page-type level, Conservative line (not statute text); the verdict is itself a comparison and recommendation, and comparing prices across channels counts as price promotion), and it is not built where switch E (peer comparison banned) is on; on the lightly regulated side it may be built after rewriting, as a scenario comparison table only, without writing "best" or "best value". The only thing the strictly regulated side can borrow is the skeleton idea — giving "boundary conditions" their own section, which actually helps compliance. See → General Edition 0.3 Which page types are limited for you, and what to build instead.

Most common mistake: burying the verdict in a "summary" at the end. This type's whole design is to split the verdict, the reasoning and the boundary into separate H2s, so AI copies all three in one pass.

5.23 Tier-B questions and rules types (1): regulatory obligations, definitions

What you'll do in this section: for rule questions, put the obligation sentence right under H1, shaped as "subject + deadline + penalty"; for definition questions, write each H3 as the buyer's original question, with the first sentence under it a conditional sentence — "if … then …". What AI is looking for is "under what condition is this required", not the definition itself. When you're done, you'll have two blueprints, and you'll know that under rule questions you are usually not the source, and how to schedule around that.

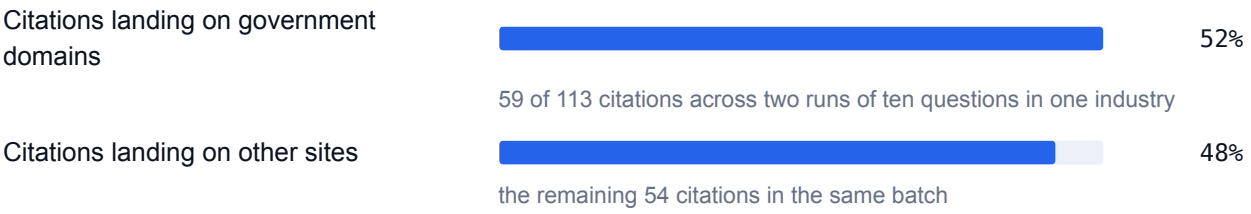


Figure: under rule questions, government pages dominate the cited surface

The figure is from an industry billed as lightly regulated: once the question shape shifts from "how much / who's best" to "how do I do it / is it regulated / what do I need to prepare", it looks like this, and vendor pages still get in only for pricing, review and integration questions. So for the two types below, you are usually not the source, and can only build a trust asset or a judgement page. Evidence in Appendix A.4.

5.23.1 Regulatory obligations and penalties page (pt10)

Intent: "is X legal", "what are the rules for X", "do I need to register", "what happens if I don't". **Leg**: ChatGPT (AI Mode not measured); this type has the largest sample in tier B. **title and H1**: title = a noun

phrase naming the obligation, no year; H1 same as title, a statement. Clarification-piece variant: the H1 is itself a judgement sentence.

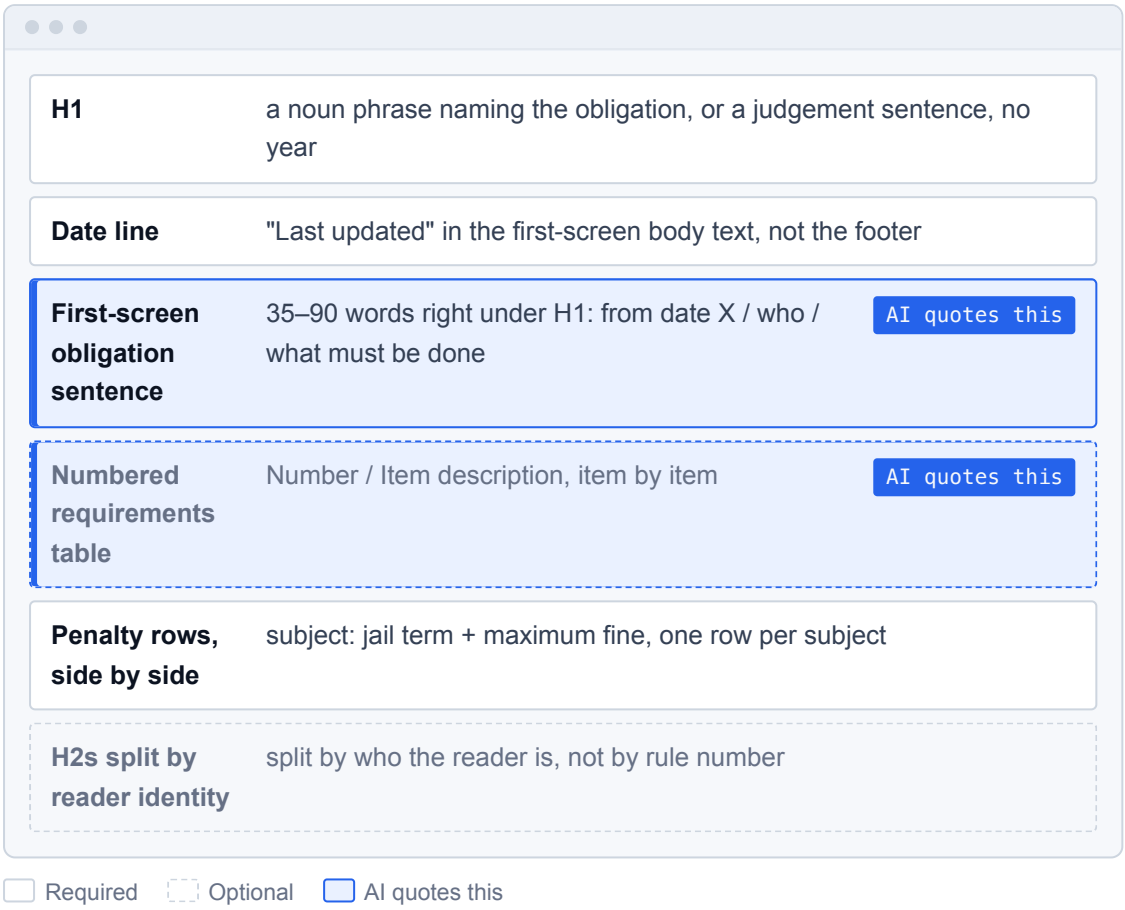


Figure: Regulatory obligations and penalties page; AI copies the obligation sentence right under H1

Standard column headers: `Number / Item description` (requirements table); penalties do not use a table — use side-by-side rows of "subject: fine + jail term". **Sample sentences**: own sentence (downstream citation page wording, not measured) — "Under <full name of the legislation + year>, <subject> must <action> by <date>; failure to do so may lead to a fine of up to S$<amount>. <Our organisation>'s approach is: <one sentence describing our own process>." Market sentence: does not apply — this type has no market range, only what the statute says.

> **Example** (corporate compliance) The first sentence right under H1 on an official obligations page was copied whole: "AIS employers are required to submit … by 1 Mar each year. Late submission may lead to a fine of up to $5,000." The subject, the deadline and the fine are compressed into one sentence — AI cannot work this out on its own. Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: you are usually not the source for this type. On the strictly regulated side, only build a "regulation summary page" (a verbatim quote + a link back + a check date), scheduled as a trust asset, not a traffic page. On the lightly regulated side you may self-build a judgement page on "does this rule apply to you", with threshold numbers matching the original word for word.

Most common mistake: writing the obligation sentence as an explanatory lead-in paragraph. In the measured cases, mid-page explanatory paragraphs were never copied once; AI only took the sentence right under H1 or a row from the table.

5.23.2 Definition page (pt12)

Intent: "what is X", "what is X, why do it", "how is X calculated", and it is also reused as a factual source for checklist and full-guide questions. **Leg**: ChatGPT (AI Mode not measured). Note The definition of a generic consumer term is part of the model's internal knowledge and needs no external retrieval — without a local variable (a local regulation name, a standard number, a local price), don't put budget into a definition page. **title and H1**: title = the term itself + up to three sub-questions, no place name, no price; the government/institutional version goes further, using the bare term alone. H1 is never written as a question — this is the hardest line dividing it from "one question, one page".

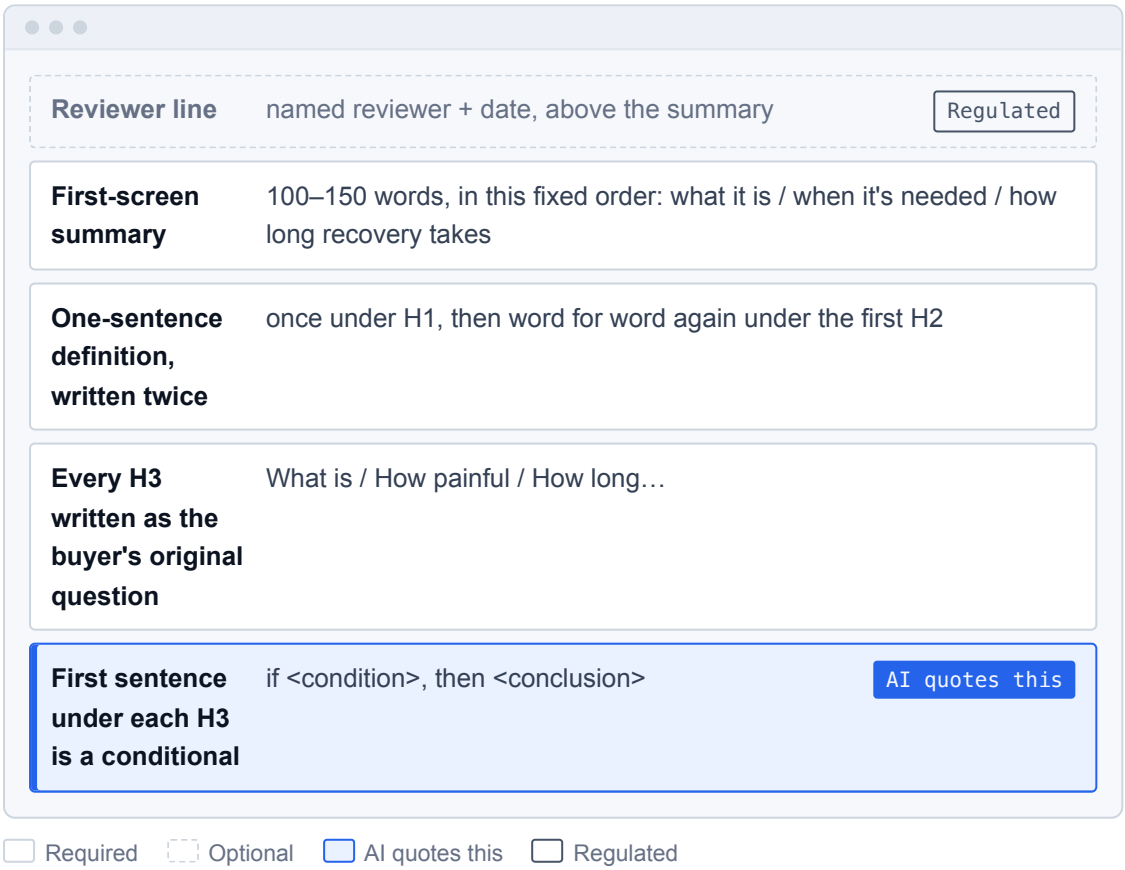


Figure: Definition page; AI copies the conditional sentence under each H3 question

Standard column headers: this type mostly has no table. When it does, use `Option / What it is / When it applies` (not measured). **Sample sentences**: own sentence (this type's make-or-break line) — "if <condition>, <conclusion>", the first sentence under every H3 question must be a decidable conditional sentence, never a lead-in. Market sentence (the vendor version's only way to win) — put facts carrying a local variable into Key takeaways; drafts that explain the term in generalities get 0 citations.

> **Example** (dental) On a medical-term definition page, all 10 copied sentences fell on the first sentence under an H3 question, and all were conditional sentences; not once was the first-screen summary paragraph copied. Evidence in Appendix A.4.

Wording copied 0 times	Wording copied 10 times
H3 = a noun phrase, e.g. "Risks" or "Recovery period"	H3 = the buyer's original question
First sentence under H3 is a lead-in	First sentence under H3 is a conditional
No copied sentence in the summary paragraph	Copied sentence falls on the first sentence under an H3 question

Figure: on the same page, the wording that got copied 0 times versus 10 times

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: a definition touching personal safety needs a named reviewer + registration number + date, with risk wording matching the official device documentation. On the strictly regulated side, the first-screen "how long recovery takes" item and the recovery-period block may only carry content quoted verbatim from official device documentation with the source noted; if it cannot be quoted, delete it (Conservative line (not statute text)).

Most common mistake: writing H3 as a noun phrase, or writing the first sentence under H3 as a lead-in ("everyone's situation is different…"). AI's retrieval hook is the H3; its answer hook is the first sentence under it. Get either one wrong and nothing gets copied.

5.24 Tier-B questions and rules types (2): procedure steps, preparation lists

What you'll do in this section: for procedure questions, write the steps as a Step / Result table, fix the time limits, and state the official fees; for preparation questions, group by identity or scenario, and write out "what we already provide" — that half is the part no one else can copy. When you're done, you'll have two blueprints.

5.24.1 Step-by-step procedure page (pt13)

Intent: the tutorial type asking "how do I do it, step by step"; the checklist type asking "which documents do I need to submit" also lands here. **Leg**: ChatGPT (AI Mode not measured). **title and H1**: title = opens with a verb + a bracket saying which track; H1 same as title. One first-screen sentence states clearly what this page delivers.

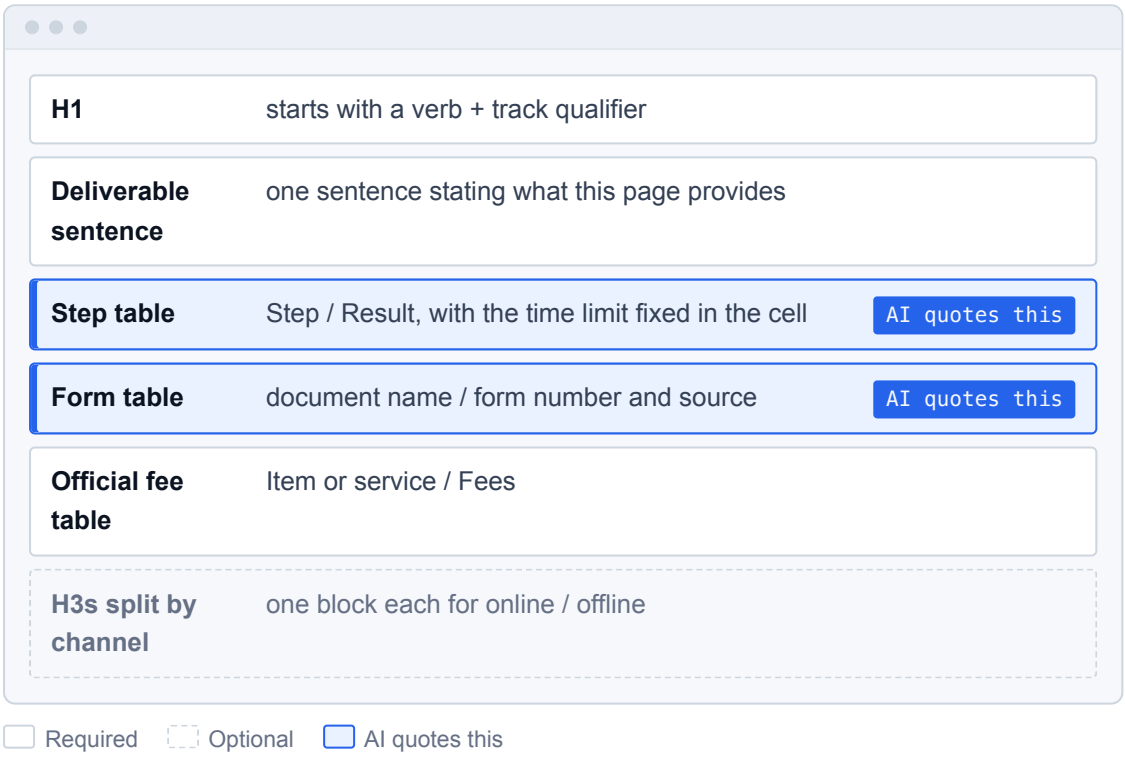


Figure: Step-by-step procedure page; what AI copies are cells from the step table and the form table

Eligibility and process page (pt07)	Step-by-step procedure page (pt13)
Answers whether you qualify, how many visits, how long it takes	Answers what to do at each step, which form to submit, what the official fee is
What gets copied is the conditional sentence	What gets copied is a table cell

Figure: pt13 and pt07 can both exist on the same site, and what gets copied differs between them

In an industry where the rules just changed or the process is complex, this is the type you should fight hardest for: the government page gives the general track, and you give "how this step works with us, what our handling fee is" — never comparing fees with peers.

Standard column headers: `Step / Result`; `<document name> / <form number>, <source> (format, size)`; `Item or service / Fees`. **Sample sentences**: own sentence — "For <matter> at <our organisation>, step <N> is <action>, the result is <a verifiable state>, and the time limit is <N> days; you need to submit <full form name> (<number>); our handling fee is S$<price> nett, the official fee is S$<price>." Market sentence — the official fee is the only market figure; copy it directly from the regulator's table and label it with the check date.

> **Example** (legal) The "Step | Result" table on an official procedure page was reordered by AI into a numbered list, with the days carried over as-is; the fee table let it write out the exact official fee figure in passing. Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: the time limit is only how long the process takes; it must never be written as the time for a result to materialise (see the eligibility and process page row in 5.3 for the wording).

Most common mistake: writing the steps as a narrative paragraph. In the measured cases, 100% of copied sentences came from table cells; the narrative paragraph was never copied once.

5.24.2 Preparation and bring-list page (pt14)

Intent: "what do I need to prepare before X", "what should I bring", "X checklist". **Leg**: ChatGPT (AI Mode not measured). **Note** Unstable: running the same question twice, the number of citations can range from zero to several (evidence in Appendix A.4); when scheduling, treat it as "worth doing, but don't expect it to be cited every time". **title and H1**: title / H1 = "What to Do Before <matter>" / "What to Bring" / "<matter> Checklist". Neither a short link nor a PDF stops it from being cited.

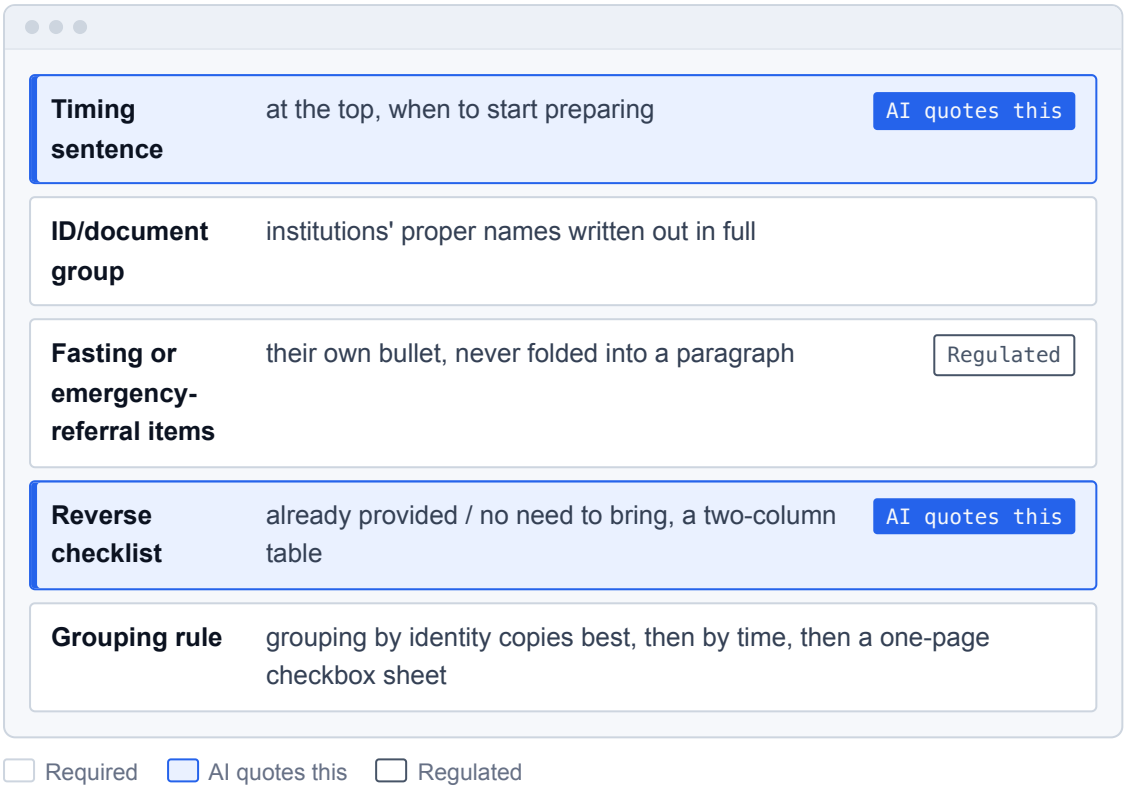


Figure: Preparation and bring-list page; AI copies the timing sentence and "what we already provide"

Standard column headers: as measured, the reverse checklist is two columns split by person (`Baby / Mother`), with the cells listing what the organisation already provides, so you don't need to bring it; the generic wording "what we already provide / what you need to bring" is not measured. **Sample sentences**: own sentence — "Start preparing <N> weeks before <matter>. Must bring: <item A>, <item B>, <full document name>. We've already prepared <item C>, <item D> for you — no need to bring your own." Market sentence: does not apply — this type gives a trade-off, not a range.

> **Example** (dental) 3 items were copied from a pre-treatment preparation checklist page; AI laid the answer out directly as checkboxes, one checkbox per page item, copying the imperative sentence itself. Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: fasting or emergency-referral items are written as personal-safety items and need review; the rest of the page should only carry factual items, never touching price or efficacy.

Most common mistake: writing only "what to bring" and not "what we already provide". The latter is the half no one else can copy.

5.25 Tier-B questions and rules types (3): schedules, collected FAQ, policy hubs, change notices, misconceptions

What you'll do in this section: follow one card for each of these five question types — dates, a bundle of multiple questions, a full guide, rules that just changed, and a popular false premise. Put dates in a table, not in H1; make the first sentence of every 40–80-word answer the conclusion; build the full guide as a single hub page; publish a notice page every time a price changes; write misconceptions in the buyer's own first-person words. When you're done, you'll have five blueprints.

5.25.1 Schedule and deadline page (pt16)

Intent: "when does registration open", "which days is the window", "can I still make it now". **Leg**: ChatGPT (AI Mode not measured). AI cross-checks process dates, so every business publishing its own version of a page with a date table has a place — it does not count as duplicate content. **title and H1**: title = "<process name> | <organisation name>"; H1 = the process name, no year — so each year you update the table without changing the URL, and the year goes in the dates in the body.

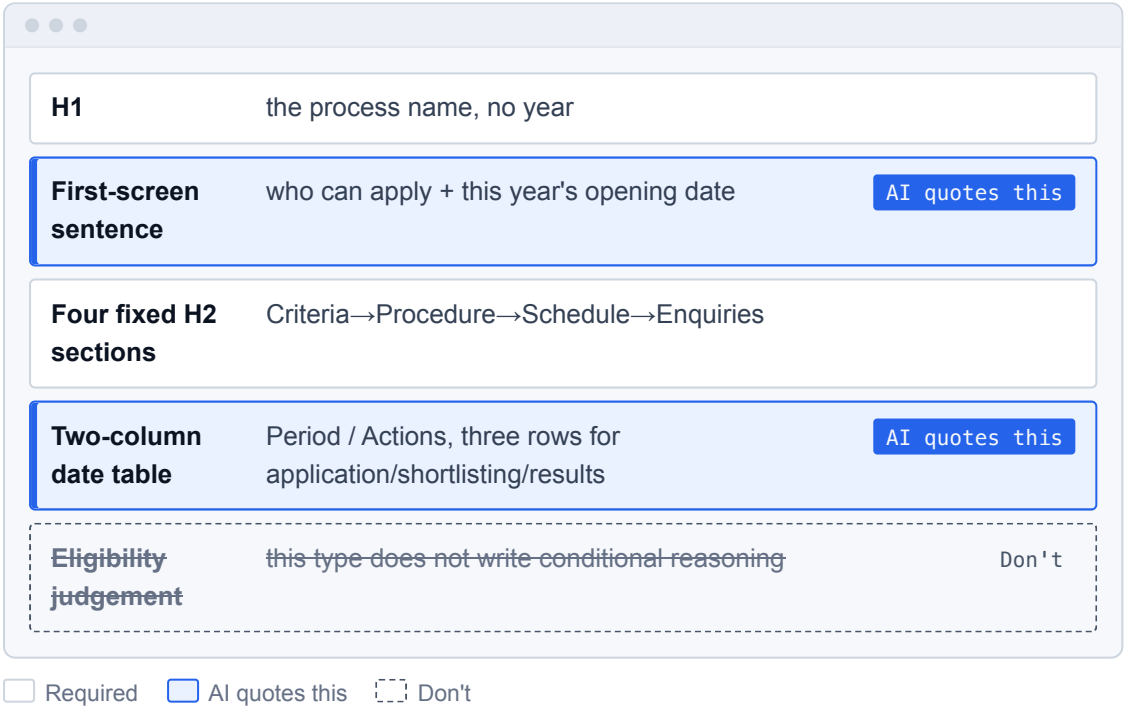


Figure: Schedule and deadline page; AI copies the first-screen date range and the date table

Standard column headers: `Period / Actions`. **Sample sentences**: own sentence — "The application period for <process> is <date A> to <date B>." A date range with the year, written in one sentence. Market sentence — also write in the next window; AI uses the date table for time-based reasoning.

> **Example** (school admissions) The first-screen sentence giving the date range with the year on an admissions process page was copied, and rows from the two-column date table were reordered into a timeline in the answer; in the same question, AI cited nine different organisations' pages of this same type in one go, to cross-check them. Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: where switch B (legally required fields) is on, the course or service name attached to each class or service date must match the registered details. On the strictly regulated side, booking windows and intervals may only state how long the process takes, never the time for a result.

Most common mistake: writing the date into H1, so the URL changes every year and historical weight resets to zero. Put the date in the table and keep H1 unchanged.

5.25.2 Collected FAQ page (pt22)

Intent: definitions and risks asked together; long-tail follow-up questions from tutorials and rules. **Leg**: ChatGPT (AI Mode not measured). The difference from "one question, one page": that type is one question per URL, with a yes/no first sentence; this type is many questions on one page, each answer 40–80 words, winning by supplying multiple answer fragments from the same page. **title and H1**: H1 = "Frequently asked questions" or "<item name> FAQs". The question itself is the hook, not H1. Three forms all get cited: an H3-question version, a numbered plain-text version, and a bold-question-lead-in version.

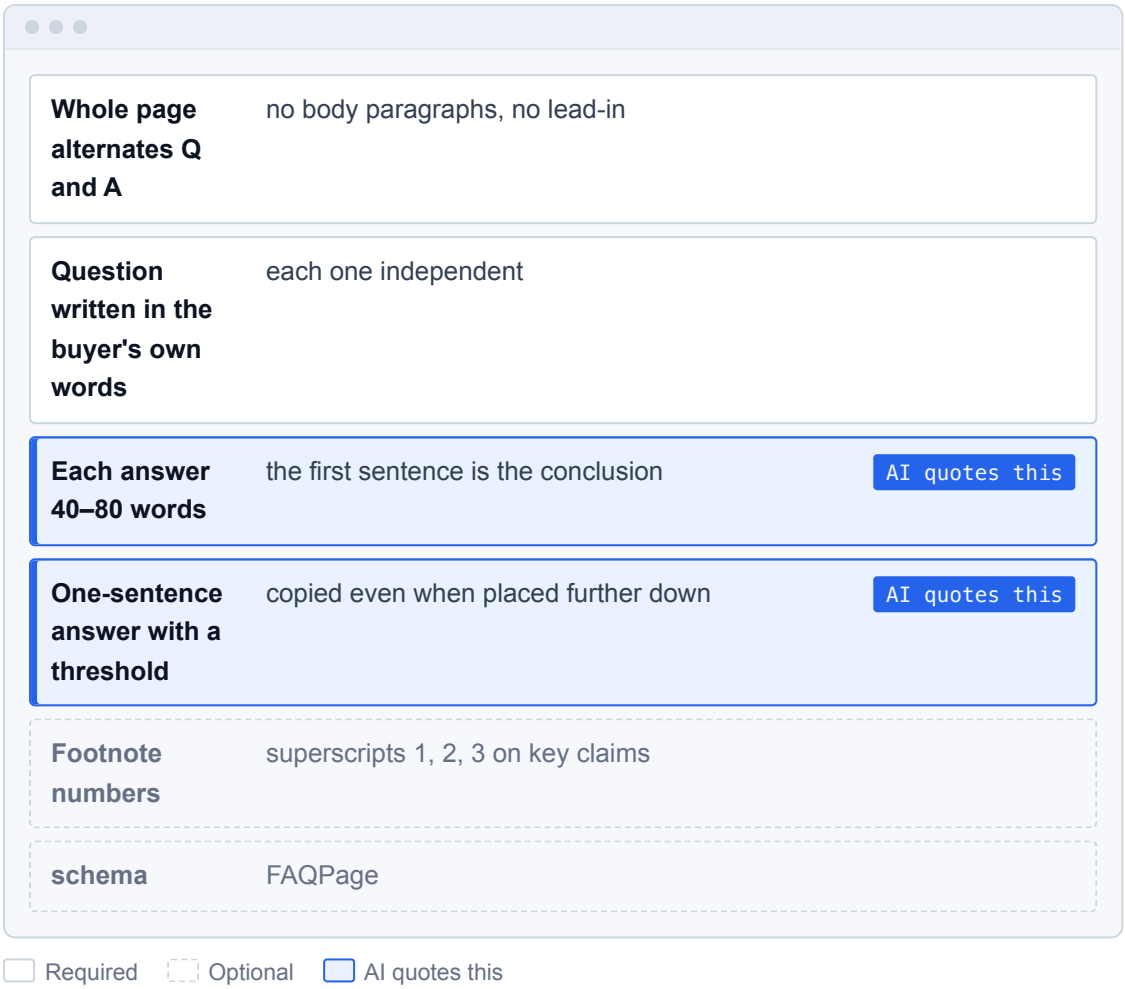


Figure: Collected FAQ page; several Q&A blocks on the same page get copied separately

Standard column headers: no table, the whole page is question and answer. **Sample sentences**: own sentence — "<buyer's original question>? — <one-sentence conclusion>. <one sentence with a condition or number>[footnote N]" Market sentence — write the industry's common basis into the second sentence of the same answer, and label the source.

> **Example** (aesthetics) On one FAQ page, 3 of the 7 citations in a single answer came from the same page, copied separately from three different Q&A blocks: what it is came from the first block, slight redness from the pain block, and swelling, tingling and tenderness from the recovery block — one page supplying multiple fragments. Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.

- Specific to this page type: can be built on every side; it is the most practical type for the regulated sides. Any claim involving scale must first pass the six checks for numbers (see → General Edition 5.10 Shared parts for all page types: six checks for numbers, no public price, billing units, dates and schema); where switch B (legally required fields) is on, answers about fees must be written out in full against the legally required terms.

Most common mistake: writing the answer as a three-part structure (lead-in + explanation + conclusion). Each answer should be 40–80 words with the first sentence as the conclusion, or chunking will miss the conclusion.

5.25.3 Policy hub page (pt26)

Intent: the full-guide type, "complete guide to X", "all the rules for X". **Leg**: ChatGPT (AI Mode not measured). Under a full-guide question, AI's answer structure is the official hub's own anchor structure — it writes its sections almost exactly in the order of the page's own in-page navigation. **title and H1**: H1 = the policy name, no year, no "guide".

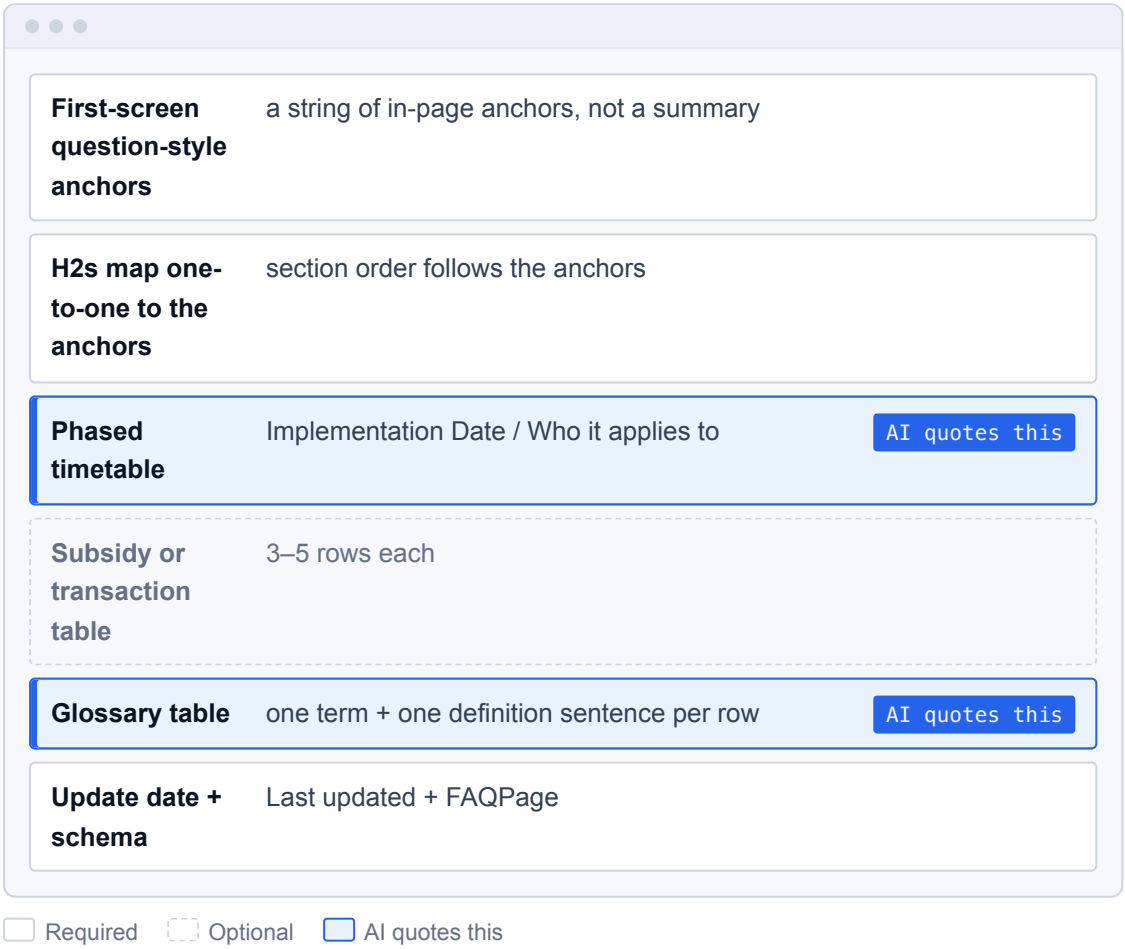


Figure: Policy hub page; AI copies rows from the phased timetable and definition rows from the glossary table

Standard column headers: `Implementation Date / Who it applies to`. **Sample sentences**: own sentence — "From <date> — applies to <which class of subject>", one row per tier. Market sentence — the glossary table, one term plus one definition sentence per row.

> **Example** (corporate compliance) A policy hub page had rows copied from its phased timetable ("from date X — applies to class Y") and definition rows copied from its glossary table; the string of question-style anchors in the first screen only had its structure copied, which does not count as a copied sentence. Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: can be built on all three sides; on the strictly regulated side the hub may only cover process and the official position, never a side-by-side review of peers.

Most common mistake: splitting the same topic into ten blog posts. In the measured cases, splitting it up actually kept it out; build it as a single hub page instead — opening with question-style anchors + a phased timetable + a glossary table.

5.25.4 Change notice / old-vs-new page (pt34)

Intent: once AI has a price or a rule, it goes looking for "does this still hold now"; industries where the rules just changed depend heavily on this type. **Leg**: ChatGPT (AI Mode not measured). This page's job is not to move rankings, but to stop AI citing an old price or old rule long after it changed — once pricing changes, AI keeps copying the old number until a page spells out the relationship between old and new. **title and H1**: the commercial version starts with "Important:", with the effective date written into the title itself. The official/regulatory version's H1 states only the reference number, with a one-sentence scope statement in the first screen. Different from the News and policy explainer (pt45): this type only covers changes to yourself or your own jurisdiction; commentary on outside events belongs to pt45.

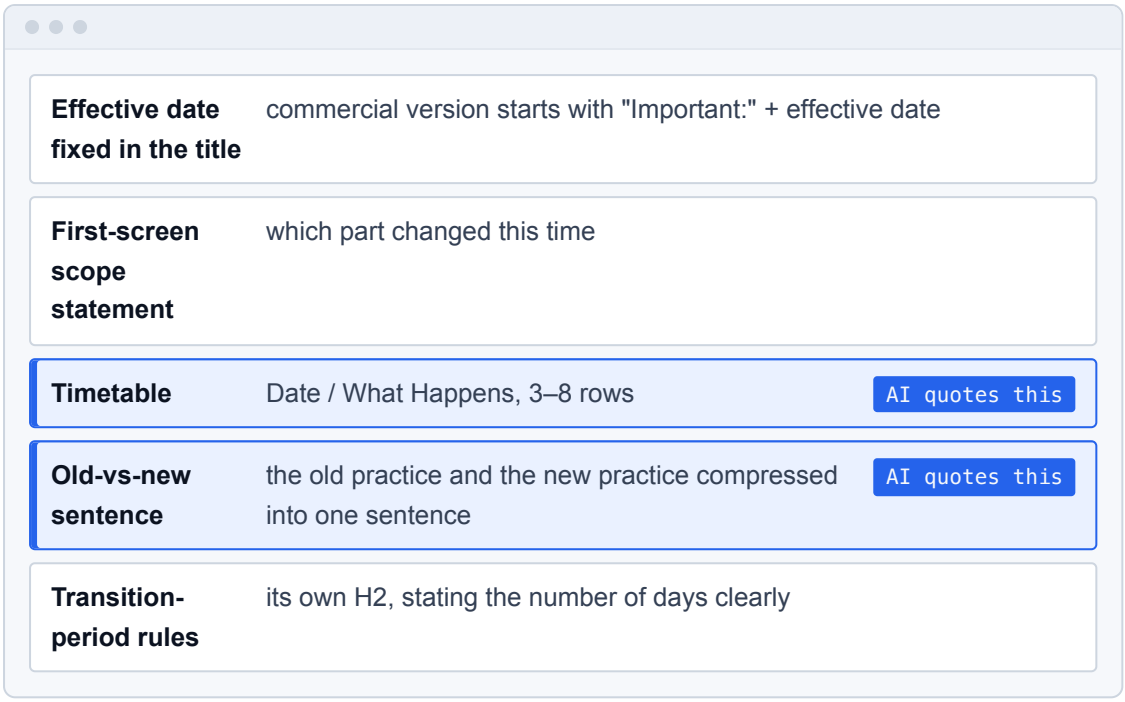


Figure: Change notice page; AI copies "from date X, the new practice replaces the old"

Standard column headers: `Date / What Happens`. **Sample sentences**: own sentence (old-vs-new sentence) — "From <effective date>, <new practice> replaces <old practice>; <a class of people> must <do what> within <N days>." Market sentence: does not apply — this type only covers changes to yourself or your own jurisdiction, and gives no market figures.

> **Example** (SaaS) What was copied from a pricing change notice was the transition date in the timetable, and the original transition-period rule sentence "After your 30-day SUITE trial ends, your account enters 'View-only' mode; you'll have another 30 days to upgrade or export your data"; AI used it to timestamp the figures on the official pricing page, marking the current price's effective date in the answer. Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: on the strictly regulated side, write only the rule and the date, never the rise or fall itself; where switch B (legally required fields) is on, update the registration record before publishing the public notice for any fee or course content change.

Most common mistake: quietly changing the number on the old page without publishing a notice page. Publish a notice page every time a price changes, and add a line at the top of the old page linking back to the notice.

5.25.5 Misconception page (pt40)

Intent: within the rules type, the question shape of "everyone thinks it's allowed, but it isn't" — triggered when a popular false premise is hidden inside the question. **Leg**: ChatGPT (AI Mode not measured). Note This type does not follow the rule that "copied sentences fall in the first 30%" — every entry carries its own full context, and AI can take any one from the middle. **title and H1**: H1 = an opposing-pair phrase ("<topic> Myths vs Facts"). Only one H2 is needed, with no further sections under it.

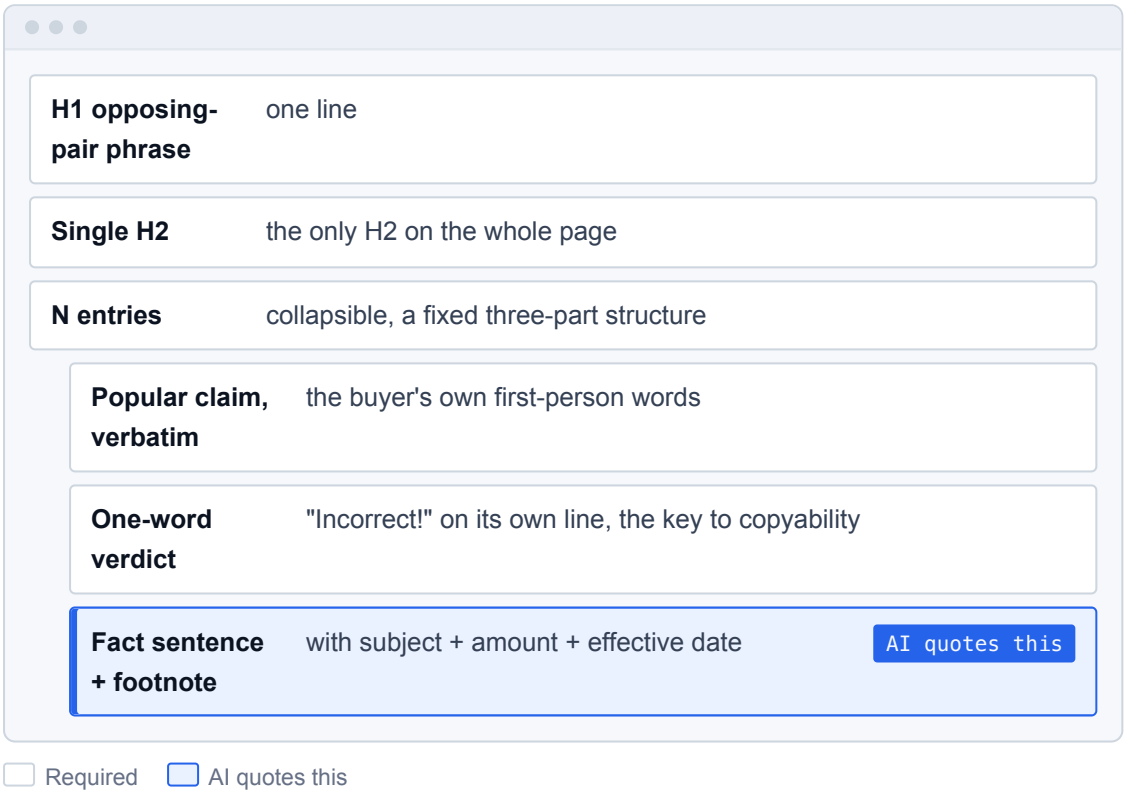


Figure: Misconception page; each entry carries its own context and does not follow the first-30% rule

Standard column headers: no table; if you build one, use "What people think / What the rule actually says / The rule it's based on / The penalty" (not measured). **Sample sentences**: own sentence — "'<buyer's original words, first person>' — Incorrect. <fact sentence, with subject + amount + effective date> (<source number>)." Market sentence: does not apply — this type gives a verdict, not a range.

> **Example** (public health) The two things copied from a misconception page were both in the middle of the article and both fact sentences: one said that placing and paying for an order counts as a purchase and is equally illegal, the other that from a given date, possessing, using or purchasing carries a maximum fine of $X. The legislation page on the same topic said only that purchase is prohibited — only this page matched the buyer's own words against the rule. Evidence in Appendix A.4.

Compliance

- Availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.
- Specific to this page type: this is one of the few types on the regulated sides that may state a "conclusion" (the conclusion comes from the regulation, not from self-description); every fact sentence must match the regulatory text word for word and carry a source number; a "misconception" must never be written as a veiled dig at peers.

Most common mistake: writing "what people think" in formal written language. Use the buyer's own words, first person — the more colloquial, the better.

5.26 Tier-B entity and product facts types (1): branches, verification pages, directory listings

What you'll do in this section: give each of two question shapes, "where is it, is it open on Sunday" and "how do I confirm they're qualified", its own plain-text facts page, then check your own entries on every directory site. When you're done, you'll have three things: one branch page per location, a list of registration numbers on your own page with links to the official register, and a check table that matches your directory entries word for word against your own site.

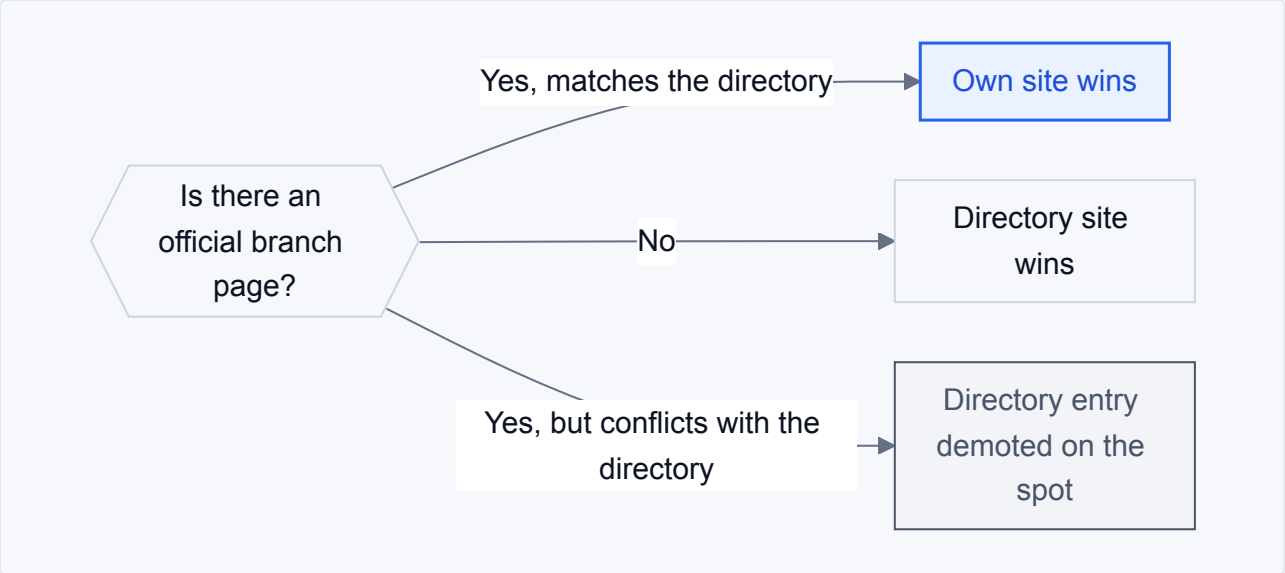


Figure: Who wins a location question depends on whether you have an official branch page that matches the directory

This figure rests on a single observation; treat it as directional only (evidence in Appendix A.4). Do this section's three cards in this order: build the branch page first; when the staff lines on the branch page go live, the verification-page wording goes live with them; check the directories last. The branch page is the safest type on regulated sides and the one to do first: the whole page is address, hours, phone and practitioner registration fields. Zero adjectives, no price, no results.

5.26.1 Store / branch page (pt11)

- **Intent**: location questions, such as "X near <MRT station / area>", "is X open on Sunday", "where in <place> can I try before I buy". Whenever a question names a place or mentions going in person, this type takes almost the whole answer's citations. **One location, one URL**: a single hub page listing ten branches loses to one page per branch on this kind of question.
- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: write title as `<service word> <place> | <brand>`, with the service word and place first and the brand last; write H1 as `<service word> <place>` or `Welcome to <branch name>`. Both the single-branch page and the top-level locator page get cited. On the hub page, use H2 for the area name and H3 for each branch, with three lines under every H3: address, phone, opening hours.

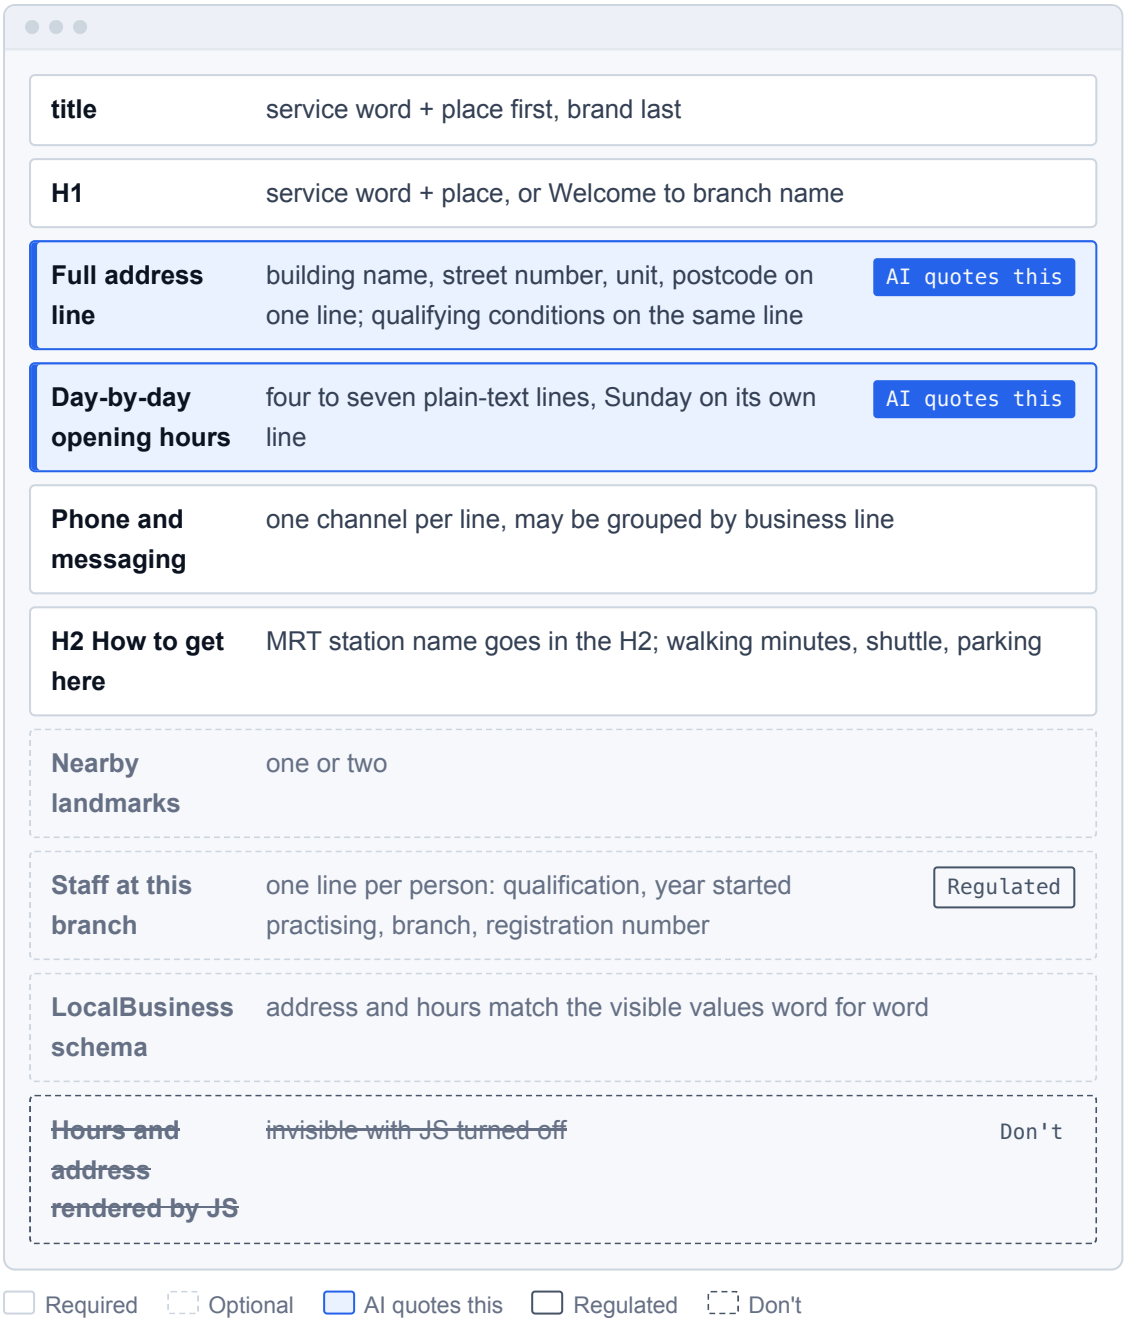


Figure: Block order on the branch page; what AI copies is the whole address line and the day-by-day hours — it doesn't read the body text

- **Whether to build a page**: a location needs at least 5 facts that belong only to it and can be checked — such as its hours table, the staff on duty there, that location's certificate number, real photos of the premises and facilities, and the actual walking time to the nearest MRT station. If you can't reach 5, don't build the page; set up only the business profile. Building ten pages for ten locations by copying the same paragraph gets none of the ten into any answer. For "near me" questions, the main entry point is the business profile (→ General Edition 3.4 Third-party credential tiers, Wikidata and the five business profiles). The branch page's job is to give the profile somewhere to link back to, and to supply the local facts that AI extracts sentences from. Never fill a profile with the URL of a branch page that isn't live yet: either publish a minimal version first (address, transit, service lines, certificate

number), or have the profile link to `/facts` first and swap the link over the day the branch page goes live.

- **Qualifying conditions**: "this location only handles a certain kind of business" must be written into the address line itself, not into an asterisked footnote. AI drops a limitation that sits in a footnote and still recommends you for the kind of question you don't handle (evidence in Appendix A.4). Page length doesn't affect citation; don't pad it with adjectives to hit a word count.
- **Standard column headers**: the locator hub page uses `STORES | LOCATION | ACCESSIBILITY | OPENING HRS`; for the hours table, use the day of the week as the column header directly (`Monday | 11:00 AM – 8:00 PM`). ChatGPT's own answer table has four cells: location name, Sunday hours, address, phone. Make sure your page can fill all four.
- **Sample sentences**: write the own sentence as "<branch name>, <building name> <street number> #<unit> <city> <postcode>, about <N> minutes' walk from <MRT station>; Monday to Friday <hours>, Saturday <hours>, Sunday <hours>, closed on public holidays; phone <number>." This type doesn't need a market sentence — the whole page can go without a single adjective.

> **Example** (dental): a chain clinic's branch page was pulled into an AI answer table using three lines from the Opening Hours / Where to Find Us block at the bottom of the page: `Sun: 9.30am – 1pm`, the street-address line and the phone line, none of them on the first screen. Three things got it into the table: one URL per branch; Sunday on its own line, with an actual time given (`9.30am`, not "by appointment"); and an MRT station name in the address.

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Specific to this page type: if switch B (legally required fields) is on, add the registered name and registration number after each location's address line.
- **Most common mistake**: hours and address rendered by JS. On the static page only a row of column headers is left, so even when AI cites you, it gets almost no information. Before handing over, run `curl` once; the page passes only if `grep` finds the street number and "Sun".

5.26.2 Verification and lookup page (pt30)

- **Intent**: the closing paragraph of a full-guide answer, such as "how do I confirm this <practitioner> is qualified", "where do I find a certified X", "how do I check this myself". This type gets cited not because it says a lot, but because it's the only step in the whole answer that lets the reader go and actually do something. You don't build the official register page or the locator; your job is to move that step onto your own page.
- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: for the register version, write the H1 as the body's name, immediately followed by a sentence on its statutory authority; for the locator version, write the H1 as `Find a <category>`

`<role>`, with the first H2 as the search box itself.

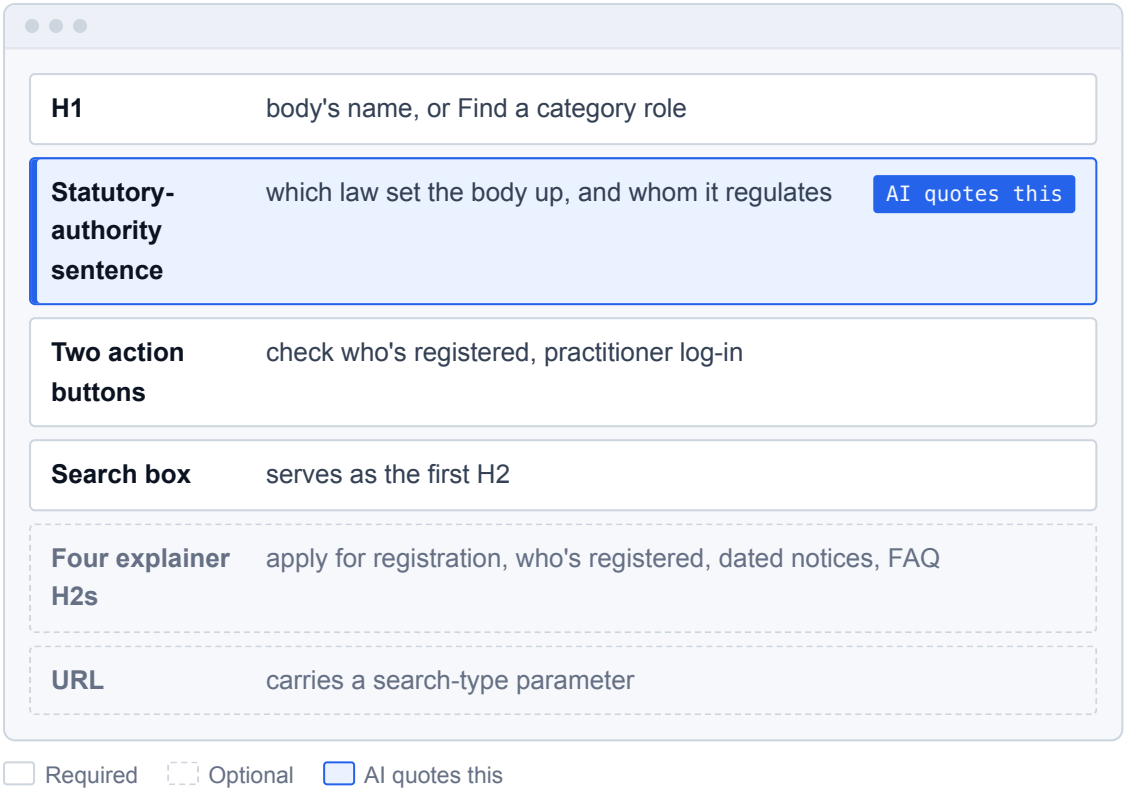


Figure: Verification and lookup page; its value is the citation slot at the end of the answer — for the register version, what's copied is the sentence on its statutory authority

- **Standard column headers**: no table.
- **Sample sentences**: put the own sentence on the team page, the `/facts` page and the branch page's staff line: "<Register> registration number of every <practitioner> at this <organisation>: <name — registration number>. You can check them yourself at <register's official URL>." This type has no market sentence. All three sides can build it; on regulated sides it's the safest trust asset of all: it states only registration numbers and the official lookup link, and makes no claim.

> **Example** (legal): a lawyer and law-firm register page didn't have a single word copied into the answer — the model just attached it at the very end, as the "go check for yourself" exit. This type gets a citation slot, not a content slot, so it has limited value for ranking on its own. But it does show that the end of an answer really does hold a spot reserved for an official lookup entry point. If your page has the registration number and the lookup link, it can take that spot. This is one observation, directional only (evidence in Appendix A.4).

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and

captions. Specific to this page type: on the strict side, the verification entry only points to a government register, never to a manufacturer locator (outbound-link row: see → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions).

- **Most common mistake**: copying the manufacturer-locator approach and linking back to the brand's own site. On the strict side, this runs straight into the ethical code's own wording.

5.26.3 Third-party directory listing (pt33)

- **Intent**: the fallback for location questions. When AI can't find a business's own official branch page, it falls back to citing the directory site's entry instead. This is a page you can't build yourself; you can do two things: publish an official branch page (pt11) for every location, and check your entries across the major directories so the opening hours match your own site word for word.
- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: title in three parts: English name + local-language name + category and place; H1 as the organisation's name with a branch suffix.

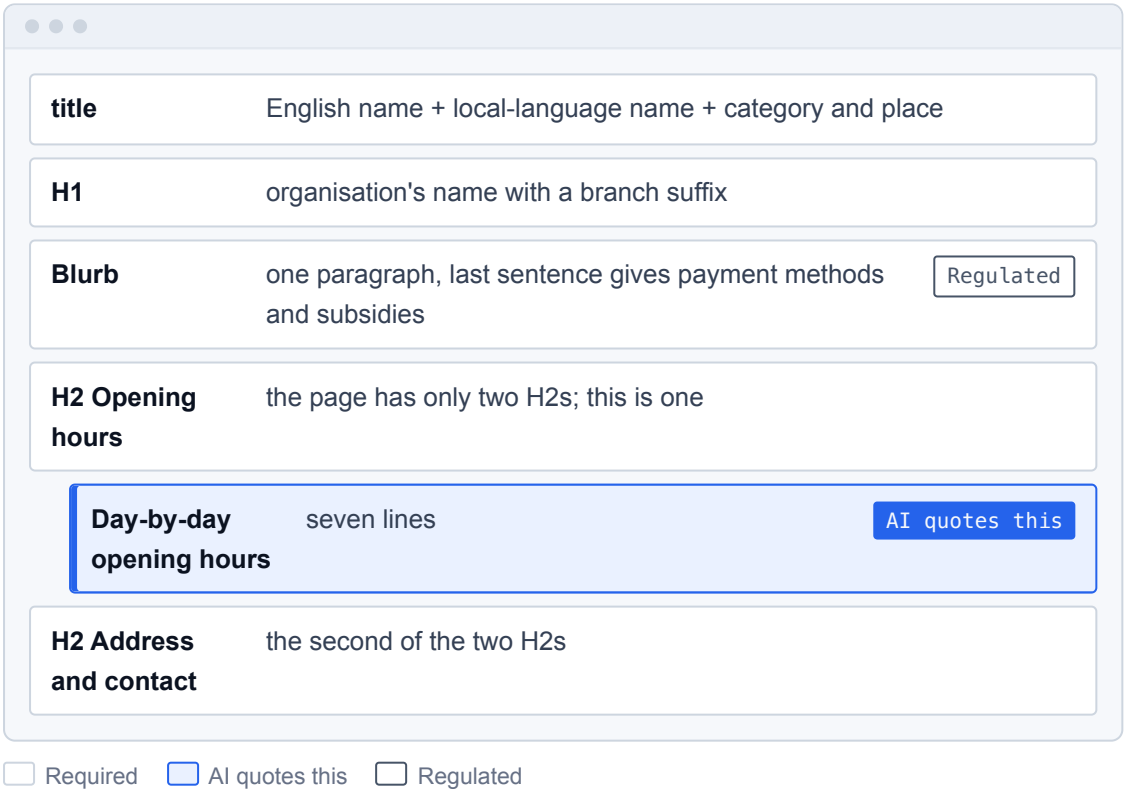


Figure: Third-party directory entry; what AI copies is the day-by-day opening hours, and that's also the line challenged when it conflicts with the official site

- **Standard column headers**: no table.
- **Sample sentences**: the last sentence of the directory blurb states the billing basis and checkable facts: "<organisation> accepts <payment methods>; <a certain service> can be paid for with <subsidy or insurance scheme>."

> **Example** (dental): a chain clinic's entry on a directory site was copied in round one as Sunday hours of 9am–9pm; in round two, ChatGPT instead flagged it as the unreliable one, in these words: "this directory lists it as open Sunday, but another current listing marks the branch closed…I wouldn't rely on that one without calling first".

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Specific to this page type: the self-supplied blurb on the directory is bound by the same price and laudatory-terms rules as your own site, and must not carry From, discounts or best (5.2 price row).
- **Most common mistake**: the opening hours on the directory were filled in three years ago and don't match your own site. The moment there's a conflict, that directory entry is demoted on the spot.

5.27 Tier-B entity and product facts types (2): compliance proof, terms, how-to docs and integrations

What you'll do in this section: give each of four question shapes (compliance due diligence, terms, how-to instructions, and whether two products can integrate) its own plain-text facts page. When you're done, you'll have a proof wall showing each issuer and validity period, terms that name the countries, a how-to document that spells out the objects covered and a negative list, and an integration page with every field filled in.

Archetype (software / data-service provider)	Regulated-side equivalent
Compliance proof wall (pt39)	Licence, credential and insurance wall
Vendor legal terms (pt21)	Personal data and confidentiality policy page
Help centre how-to page (pt31)	Service process how-to page
Integration listing (pt35)	Page listing the institutions, insurers and payment methods you work with

Figure: The archetypes for these four types all come from companies selling software or data services; on regulated sides, swap in the material in the right column — the way you write it doesn't change

Why the way you write it doesn't change: for compliance, terms, how-to and integration questions, AI wants plain-text facts it can check against, not persuasive copy — that holds whichever material you

swap in.

One more thing this family shares, and even the unregulated side can't dodge it: a certificate name or compliance-standard name is a legal claim, not an adjective. Every time one appears on the page, put next to it the form of the credential, who issued it, its validity period and the date it was last checked; writing "compliant" when the reality is only "aligned" can get you sued. Full criteria: see → General Edition 8.2 Five steps to decide the side, and three criteria.

5.27.1 Trust centre / compliance proof page (pt39)

- **Intent**: B2B procurement due-diligence questions, such as "is X compliant" or "how safe is X".
- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: a separate subdomain is optional (`trust.<brand>.com`), the main domain works too; write H1 as `<brand> | Trust Center` .

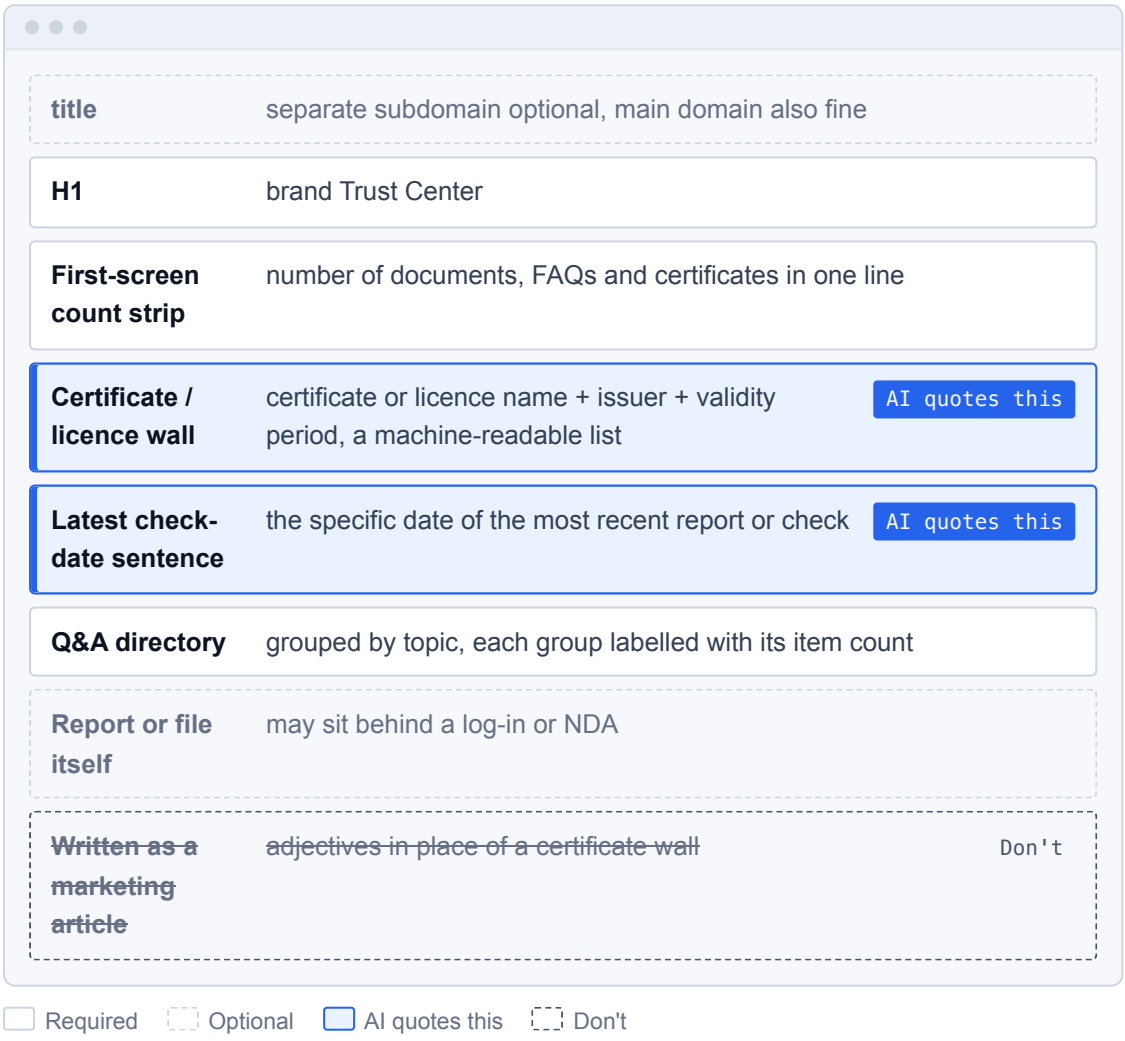


Figure: Trust centre / compliance proof page; AI copies the certificate wall and the latest-report-date sentence

- **Standard column headers**: no table; the certificate wall and the grouped Q&A directory take the place of one.

- **Sample sentences**: "<brand> holds <certificate or licence name>, issued by <issuing body>, valid until <date>, last checked <date>; the full Q&A library has <N> entries, and the report itself is provided <on request or under NDA>."

> **Example** (software services): on a dated certificate-wall page, one sentence was copied whole into the answer: the list of certificates held, plus the specific date of the latest report update. A marketing-style security page from the same vendor, whose H2s were all product navigation and whose certificate names appeared only inside a sales paragraph, was cited once in the same round but had not a single sentence copied. Evidence in Appendix A.4.

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Specific to this page type: on regulated sides, build this as a licence and credential proof wall; each licence or certificate must be written with its issuer and validity period next to it, never just the certificate name.
- **Most common mistake**: writing compliance as an article instead of a certificate wall. A compliance question wants a checkable list, not a persuasive paragraph.

5.27.2 Vendor legal terms page (pt21)

- **Intent**: questions such as "is it okay if data sits overseas" or "are you compliant" that need a country named in black and white.
- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: H1 gives the terms' full name with no marketing language; this type doesn't need an H2 table of contents, an FAQ, an update date or a named author.

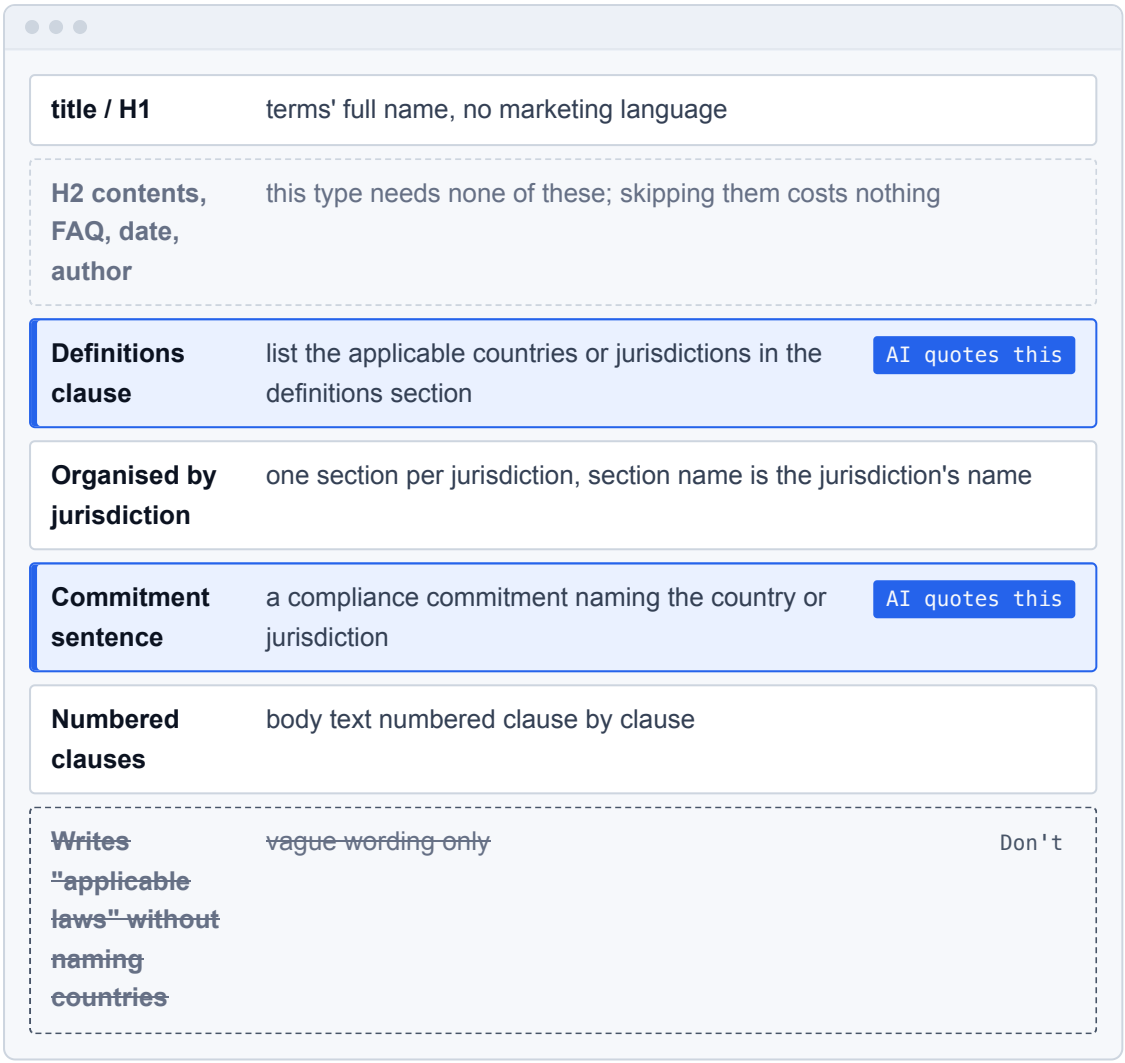


Figure: Vendor legal terms page; AI copies the commitment sentence that names the country or jurisdiction

- **Standard column headers**: no table.
- **Sample sentences**: "<brand> complies with <regulation's full name> when providing services in <country>; clause <N> of this agreement makes a commitment on <obligation name> for <country>."

> **Example** (software services): a terms page listed its applicable countries in the definitions section, and a jurisdiction-specific clause elsewhere held a compliance commitment sentence naming a country. Both were copied verbatim into the answer. The examples on the next two cards come from the same vendor. Evidence in Appendix A.4.

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and

captions. Specific to this page type: the regulated-side equivalent is the personal data and confidentiality policy page, which likewise must state the applicable law and its scope; cite only rule numbers you have verified, with the rule text in the industry edition's appendix.

- **Most common mistake**: writing "applicable laws" in the clause without naming a country. Terms with no country named are effectively invisible to local compliance questions.

5.27.3 Help centre / how-to page (pt31)

- **Intent**: tutorial and integration questions such as "how do I set up B inside A" or "how do A and B connect". A help-centre page beats a product page on the same site because it writes facts as short sentences rather than marketing language; having lots of screenshots on the page doesn't hurt citation.
- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: H1 uses a verb-first task name; one line above the H1 gives the update date and a scope condition.

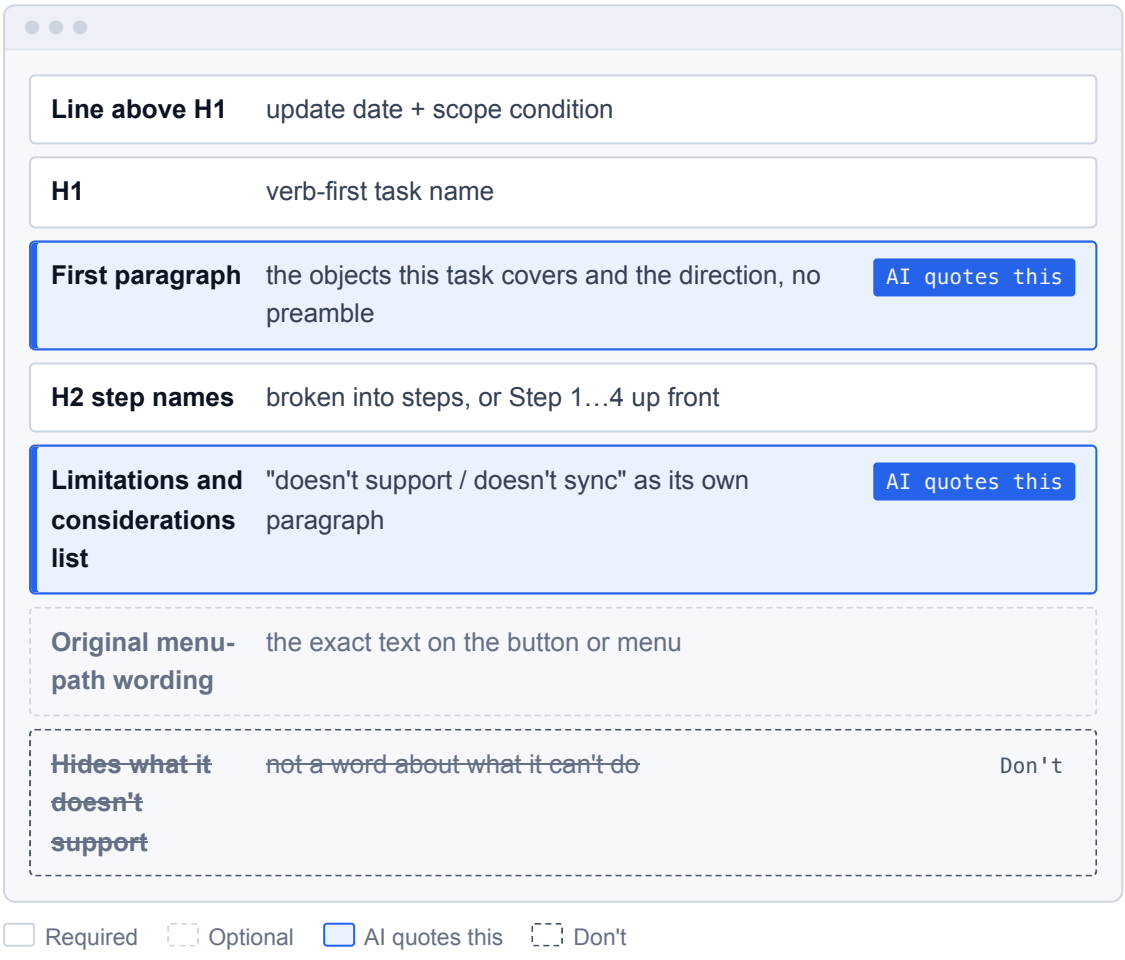


Figure: Help and how-to page; AI copies the object list and the negative list in the first paragraph

- **Standard column headers**: no table.

- **Sample sentences**: "<A> syncs <object list> with <B>, <one-way / two-way>; it does not sync <negative list>. Do this from <menu path>."

> **Example** (software services): this vendor's help-centre page opens without preamble, stating in the first paragraph exactly what it syncs and in which direction, and that sentence was copied verbatim into the answer. The negative list of limitations and considerations further down was also copied whole, as a caution. Evidence in Appendix A.4.

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Specific to this page type: the regulated-side equivalent is the service process how-to page (how to book, how to reschedule, how to get a report), covering process only, never results.
- **Most common mistake**: hiding what's not supported. AI specifically goes looking for a negative list to use as a caution — if you don't write one, it will find one on a third-party site instead.

5.27.4 Integration / marketplace listing (pt35)

- **Intent**: B2B-only questions such as "can A connect with B" or "does A have an integration".
- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: title and H1 match, one integration per URL; don't build a logo-wall page claiming "N+ integrations".

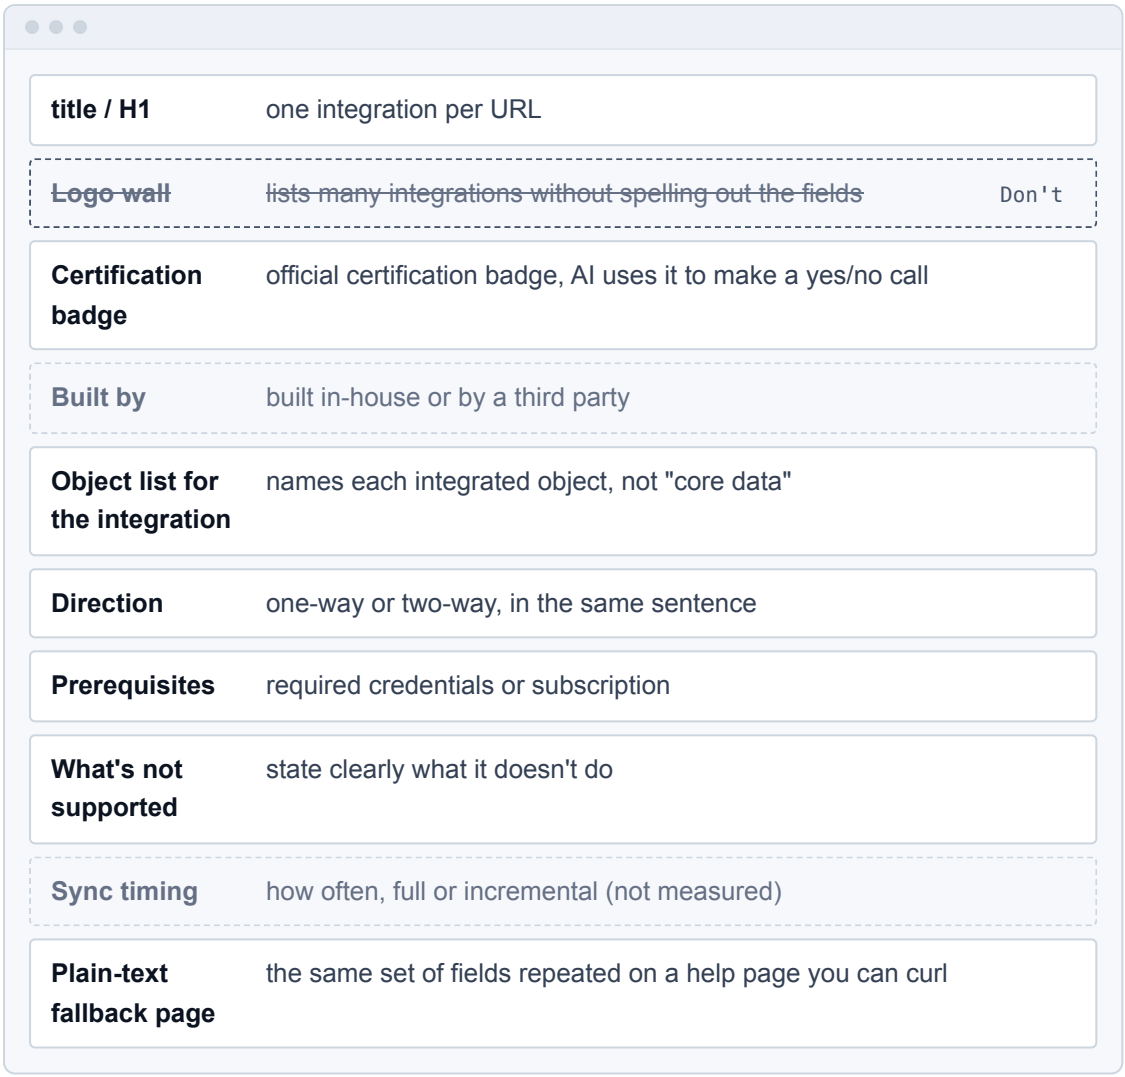


Figure: Integration listing page; it gives the yes/no call, while the object list that gets copied sits on a plain-text page on the same site

- **Standard column headers**: suggested `Integrated object / Direction / Trigger timing / Not supported` (not measured).
- **Sample sentences**: "<A>'s integration with <B> is <one-way / two-way>, covering <object 1, object 2, object 3>; it does not integrate <exclusions>. Requires <prerequisite>."

> **Example** (software services): this vendor's integration listing page is rendered by JS, so no body text could be retrieved from it; all AI got from this page was the yes/no call on whether the two connect, plus the certification badge. The object list and unsupported items that were actually copied word for word sit not on this page but in the first paragraph of the same site's help-centre page (5.27.3), which can be fetched with curl. Evidence in Appendix A.4.

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by

side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Specific to this page type: the regulated-side equivalent isn't software integration but which institutions, insurers and payment methods you work with, written with the same four fields — object list, direction, prerequisites, not-supported items; on the strict side, don't link out to commercial companies — name them only (outbound-link row: see → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions).

- **Most common mistake**: putting integration information only inside JS-rendered cards, where plain text can't reach it. Repeat the same set of fields on a page you can curl.

5.28 Tier-B primary sources (1): legislation, regulator guidance, official replies, device documents, case law

What you'll do in this section: for three types (legislation, official replies and device regulatory documents) you can't write your own version to stand in for the official one. What you can do is write your downstream citation page as a verbatim quote + section-number anchor + check date. For regulator guidance, you can publish your own compliance long-form PDF with a question-style contents page; for case law, you can build a summary version, but the subject of the ruling sentence must be the court, not you. When you're done, you'll have a downstream citation sentence template, a self-buildable PDF contents-page format, and a case-summary format.

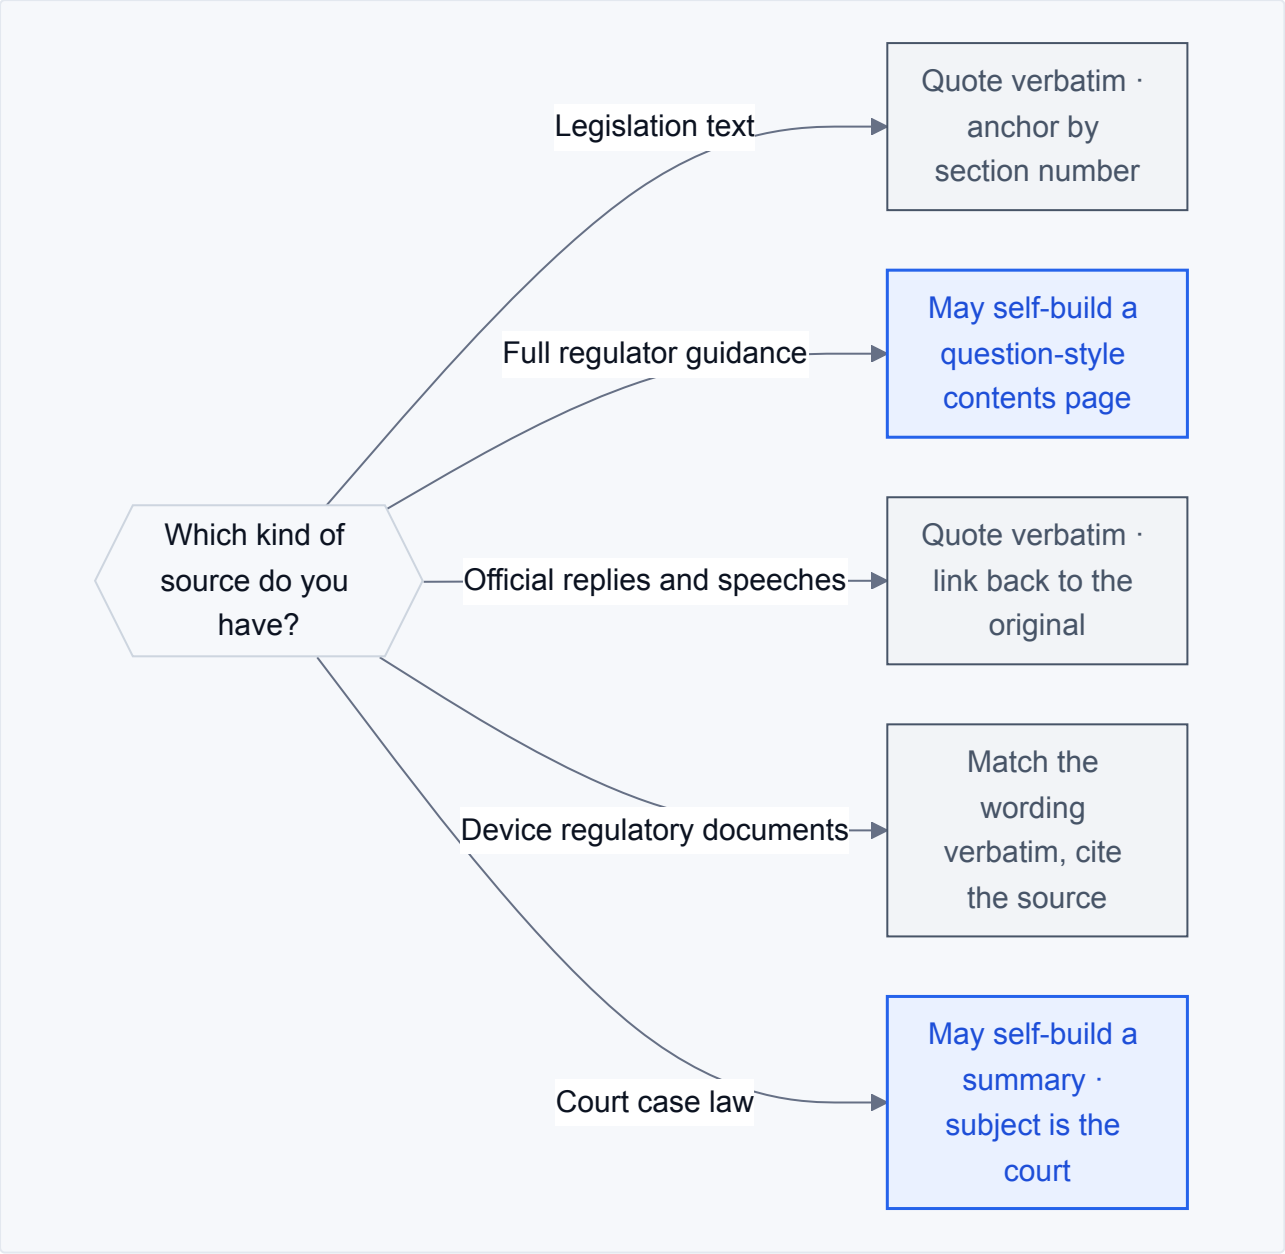


Figure: Which kind of source you have decides whether you can only cite it or can build your own page

Why you can only cite it: the original is official. If you write your own version to stand in for it, AI won't cite you; it goes straight to the original page instead. Measured: compliance explainer pages the business wrote itself were not cited once, so treat such a page as a trust asset, not a traffic page. The standard sentence for a downstream citation page is: Under <regulation's full name (year)>, section <N>: '<verbatim original text>' (link with a section-number anchor, <check date>).

All five types also share one rule: compliance constraints follow the content, not the carrier. Prices and claims that appear inside a PDF or a register are bound by exactly the same rules as on a web page.

5.28.1 Legislation text page (pt15)

- **Intent**: the "what gives it authority" step a rules-type question needs — AI wants a URL that points to a specific section number.

- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: write title as `Part <part number> – <part name> | <parent document's full name>`; the legislation page itself often has no ordinary H1/H2 structure — it opens with publication details and a list of commencement dates, then runs entirely as numbered clauses. What matters isn't the title but the URL: every clause has an address that anchors straight to its section number.

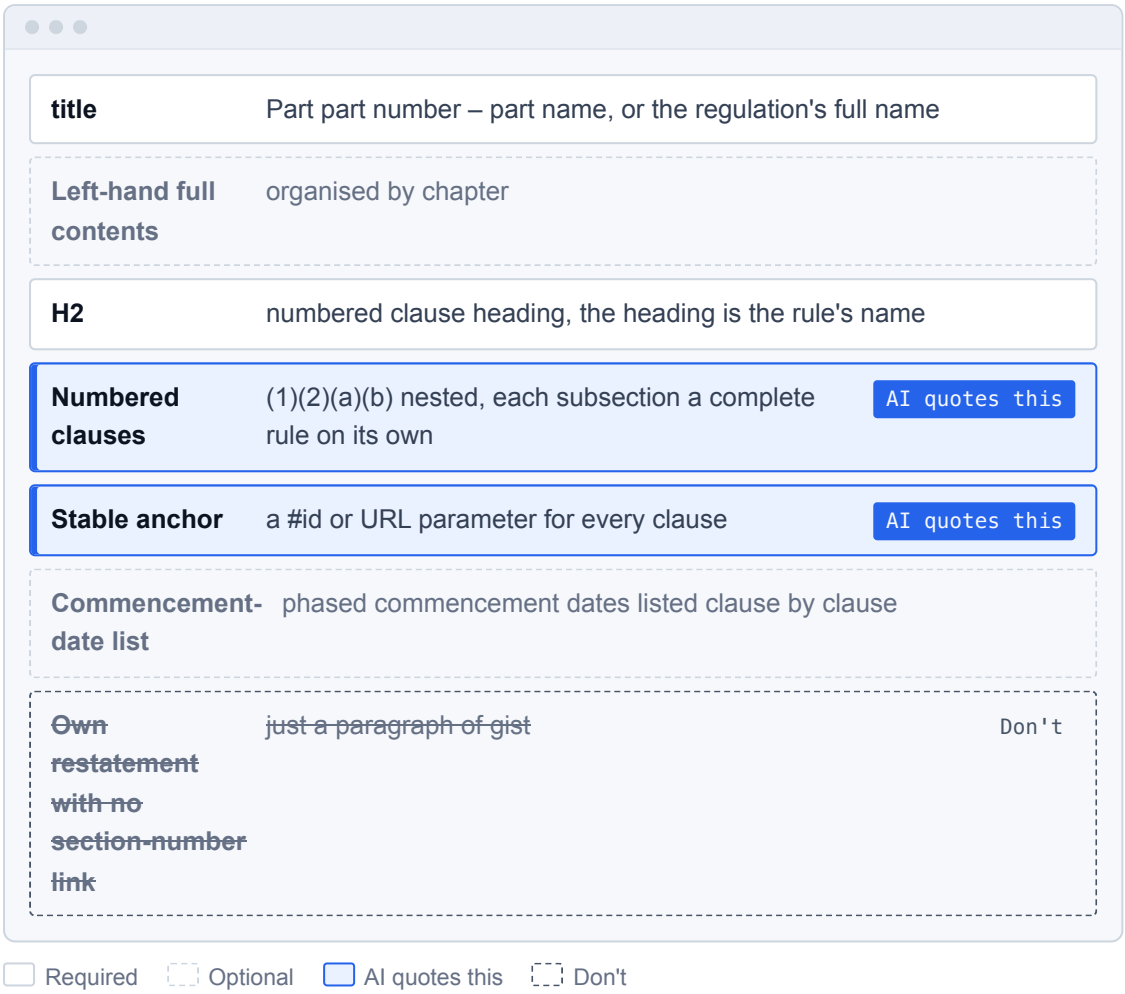


Figure: Legislation text page; what AI wants is one clause, not a whole piece of legislation

- **Standard column headers**: no table.
- **Sample sentences**: for the downstream citation, write "Under <regulation's full name (year)>, section <N>: '<verbatim original text>' (<link with a section-number anchor>, <check date>)."

> **Example** (legal): an official legislation database was cited 8 times, and every single citation carried a URL parameter that pointed to the exact clause; because the buyer asked "what gives it authority", the model wanted one specific clause, not a whole piece of legislation. Evidence in Appendix A.4.

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Specific to this page type: no side may self-build this; the downstream citation page only quotes verbatim + anchors by section number + gives a check date, and never implies an official endorsement of your own business.
- **Most common mistake**: restating the legislation without a section number or link. AI will go straight to the original page instead, and your restatement page might as well not exist.

5.28.2 Regulator guidance PDF (pt18)

- **Intent**: rules-type and full-guide questions, triggered in both English and Chinese; for Chinese-language regulatory questions, this type dominates. A PDF is not demoted: in rules-type questions it's actually the preferred source, provided it has a question-style contents page.
- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: the cover gives the guidance's full name + version + revision date; what actually gets the citation is the contents page — every line of it should be written as a question in the buyer's own words plus a page number, not a noun phrase.

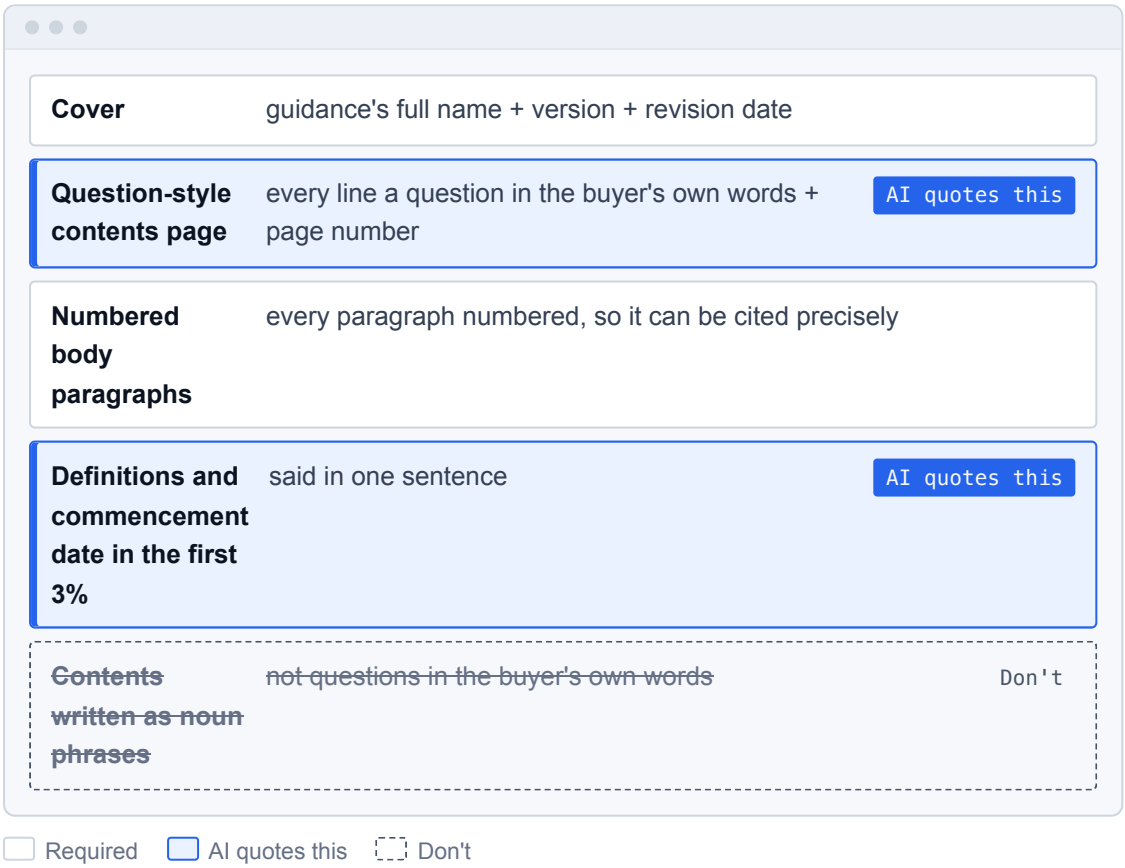


Figure: Regulator guidance PDF; a Chinese-language question hits the matching question on the contents page

- **Standard column headers**: the contents page itself is `<question> … <page number>`.

- **Sample sentences**: write the contents page of your own compliance long-form document as questions in the buyer's own words, and open each section with one sentence giving a definite conclusion + a number.

> **Example** (software services): a buyer asked in Chinese "公司把员工资料存在海外 SaaS 服务器上合法吗" (is it legal for a company to store employee data on an overseas SaaS server). It hit a line on the contents page of a regulator guidance PDF phrased the same way; what AI copied was the numbered body section that the contents line pointed to. Evidence in Appendix A.4.

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Specific to this page type: all three sides may self-build a long-form document with a question-style contents page, but any price that appears inside the PDF is still bound by the same price rule: on the strict side, give only a fixed final price, and where switch B (legally required fields) is on, write charged items in full as the law requires, with the rule text in the industry edition's appendix.
- **Most common mistake**: writing the contents page as noun phrases. Switch it to questions in the buyer's own words, and the hit rate changes immediately.

5.28.3 Official replies and speech records (pt23)

- **Intent**: in statistics-type and rules-type questions, when a number or a rule has no official document behind it, AI falls back to "what an official actually said on the record".
- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: title gives the official's title and name + occasion + year. There is no keyword optimisation at all, and it still gets cited; the first screen is just one line of date and category.

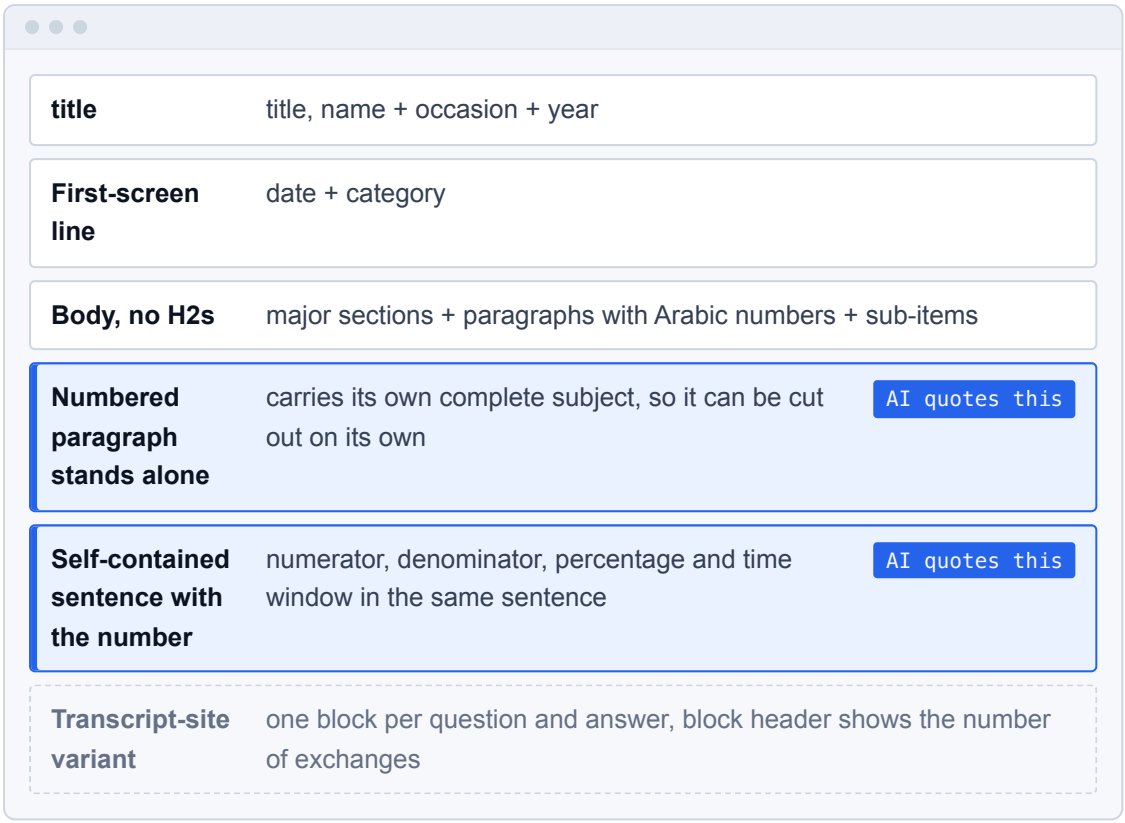


Figure: Official replies and speech records page; each numbered passage carries its own subject

- **Standard column headers**: no table.
- **Sample sentences**: "On <date>, <title, name>, replying at <occasion> on <topic>, said: '<verbatim original text>' (<link to the original>). The effect on <your audience> is <one sentence>." This is currently a source slot no one has claimed.

> **Example** (legal): what got copied from a record of an official reply was a numbered sub-item in the middle of the body, which put the numerator, denominator, percentage and time window all in the same sentence — the model could copy it whole without doing any cross-sentence arithmetic. Evidence in Appendix A.4.

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Specific to this page type: all three sides can build an official reply restatement page: quote verbatim + link back to the original + label the date; on the strict side, never use a reply to imply an official endorsement of your own organisation.
- **Most common mistake**: putting an industry figure only into a chart. Write it as a self-contained sentence — "M out of N, that is X%, over <time window>" — that's what gets copied.

5.28.4 Device and product regulatory documents (pt32)

- **Intent**: definition-type questions asking about side effects, indications, or whether something can be used on a certain body part. For claims about efficacy and side effects, AI only trusts the regulator's own wording, not a seller's self-description.
- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: any of three carriers may be cited: an IFU relay page (very short body + PDF link), a regulator explainer page (single H2 grouping), or an official approval document (indications listed clause by clause).

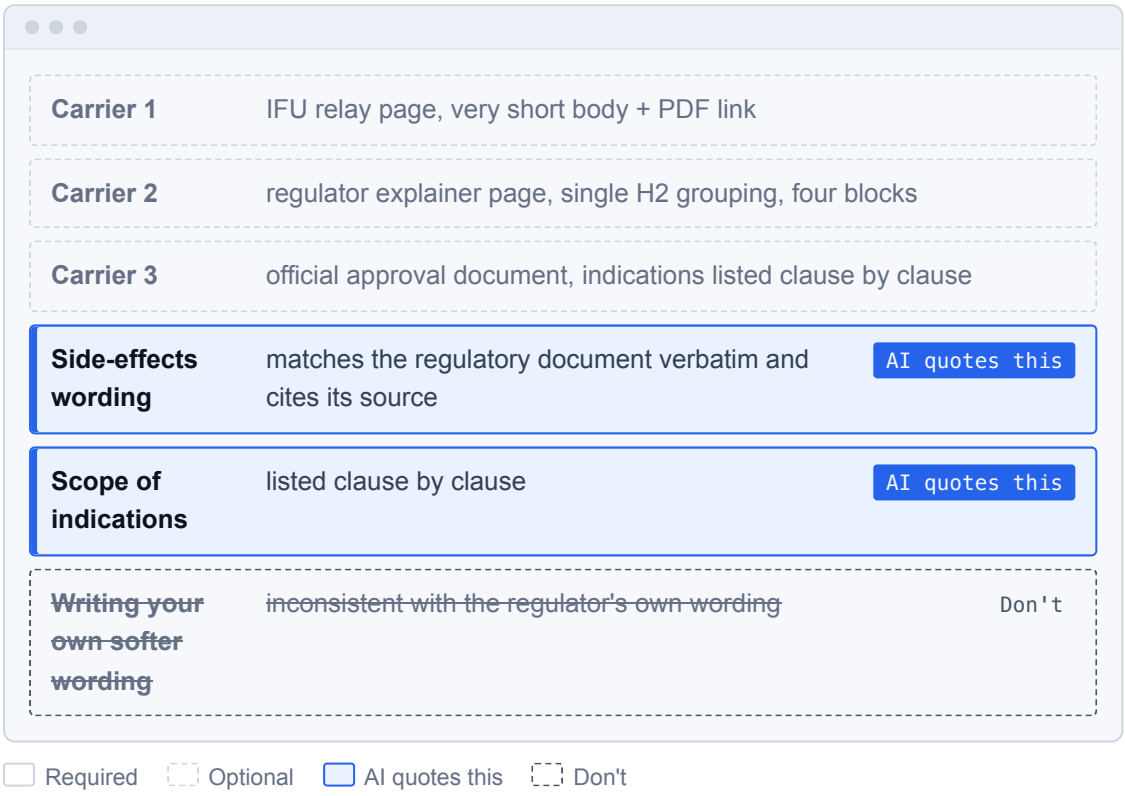


Figure: Device and product regulatory documents; the side-effects sentence goes into the answer almost unchanged

- **Standard column headers**: no table.
- **Sample sentences**: "Per <regulatory document name> (<version / date>): '<verbatim original text>'." Matching the regulator's own wording clause for clause and citing its source is the only way to get your own passage into the same frame as the regulatory document.

> **Example** (aesthetics): wording on side effects from a regulator's explainer page was carried almost word for word into an AI answer, and the scope of indications from a separate official approval document was also copied; a softer version of the same wording that the business wrote itself was never cited once. Evidence in Appendix A.4.

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Specific to this page type: on the strict side, side-effects wording must match the regulatory document verbatim — it may not be softened or omitted; you may link to the regulator's own document, but not to a manufacturer's IFU page — name only the document and its version (outbound-link row: see → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions).
- **Most common mistake**: writing your own softer version. AI won't cite it, and it's inconsistent with the regulator's own wording — which is itself a compliance risk.

5.28.5 Case-law page (pt36)

- **Intent**: in full-guide and estimate-type questions, when the model needs the step of "what the court actually ruled".
- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: for the summary version, title = H1 = case name + subject + year + court abbreviation + case number; for the full-text version, H1 gives the case number, with a subline giving court / date / case number together.

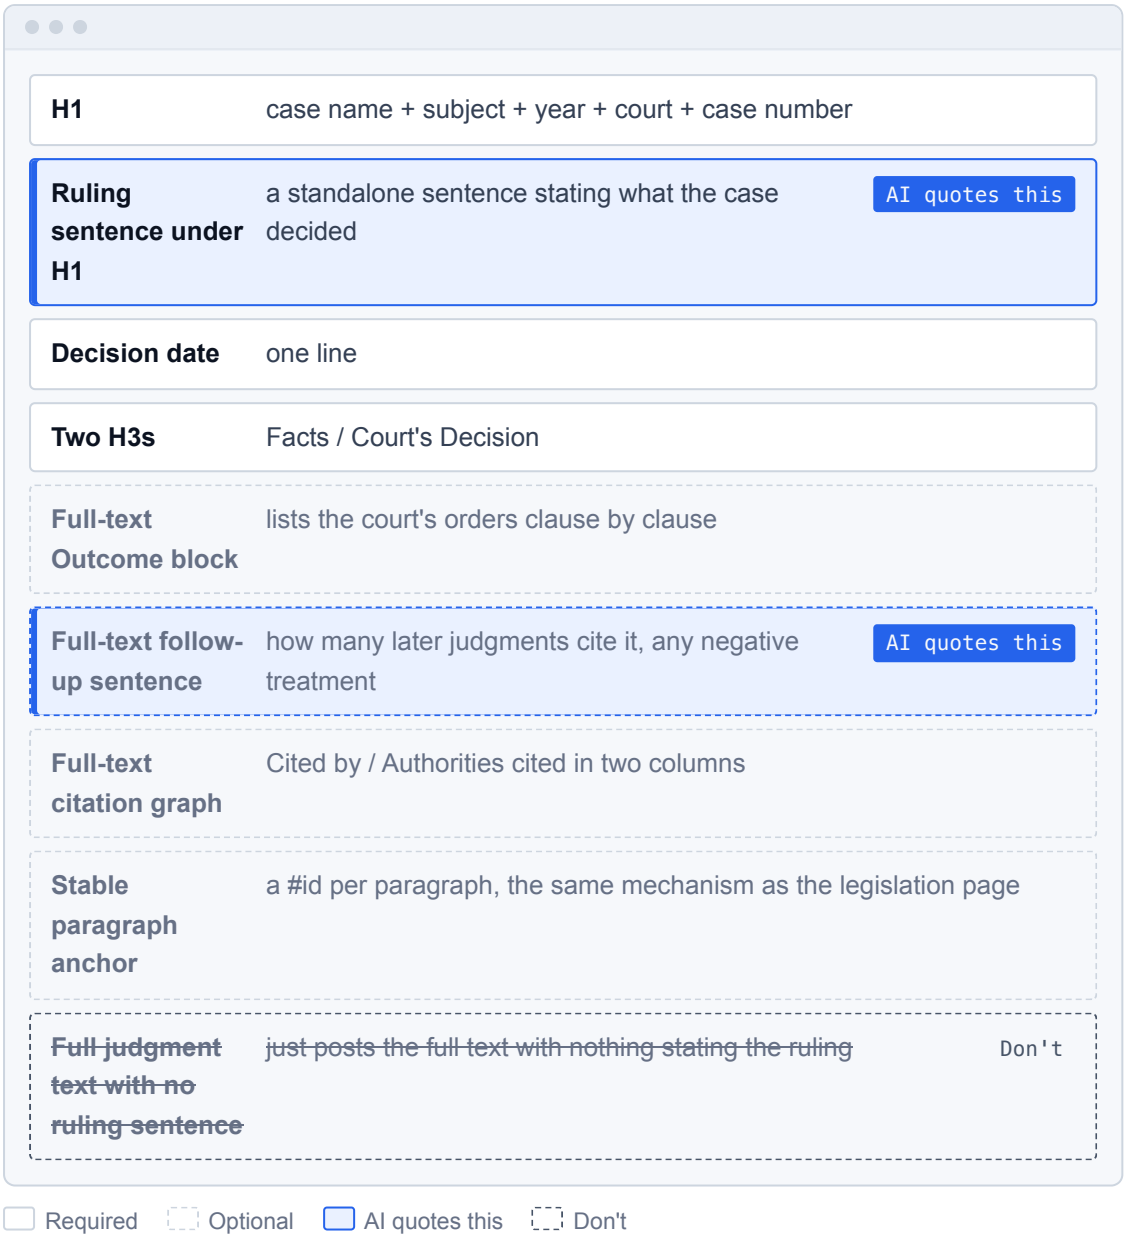


Figure: Case-law page, summary version; AI copies the ruling sentence under the H1; the full-text version separately has its sentence on whether the case is still good law copied

- **Standard column headers**: no table; the citation graph uses a two-column list (Cited by / Authorities cited).
- **Sample sentences**: "In <case name [year] court case number>, the court held that <one-sentence rule>." The subject is the case, not you — never write it as "a case we handled" or a win rate.

> **Example** (legal): a case summary page of only a few hundred words was cited alongside a full judgment running to over ten thousand words. Right under its H1, the summary gave a ruling sentence stating what the case decided, doing the summarising for the model; from the full-text version, the sentence copied was the one saying the case is still cited by later decisions with no negative treatment (the page's own wording: "No negative treatment detected"). Evidence in Appendix A.4.

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Specific to this page type: you may build a case summary page, but not a track-record page. The subject of the ruling sentence must be the court, never a win rate; don't write "represented by us" unless the client has agreed in writing and it has cleared the six checks (result-numbers row: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers).
- **Most common mistake**: just posting the full judgment and calling it done, with no ruling sentence on the first screen. The model doesn't read long text — it wants one sentence stating what the case decided; if you don't write it, it will go find the site that did.

5.29 Tier-B primary sources (2): statistics sources, registers

What you'll do in this section: get onto every register you can, and fill in every field on your own row. When you have first-party data in hand, write it up the way a statistics source page does (method, denominator, a scope-exclusion sentence), so you become the source other people cite. Having no official number isn't a dead end either: spelling out the gap may itself get cited.

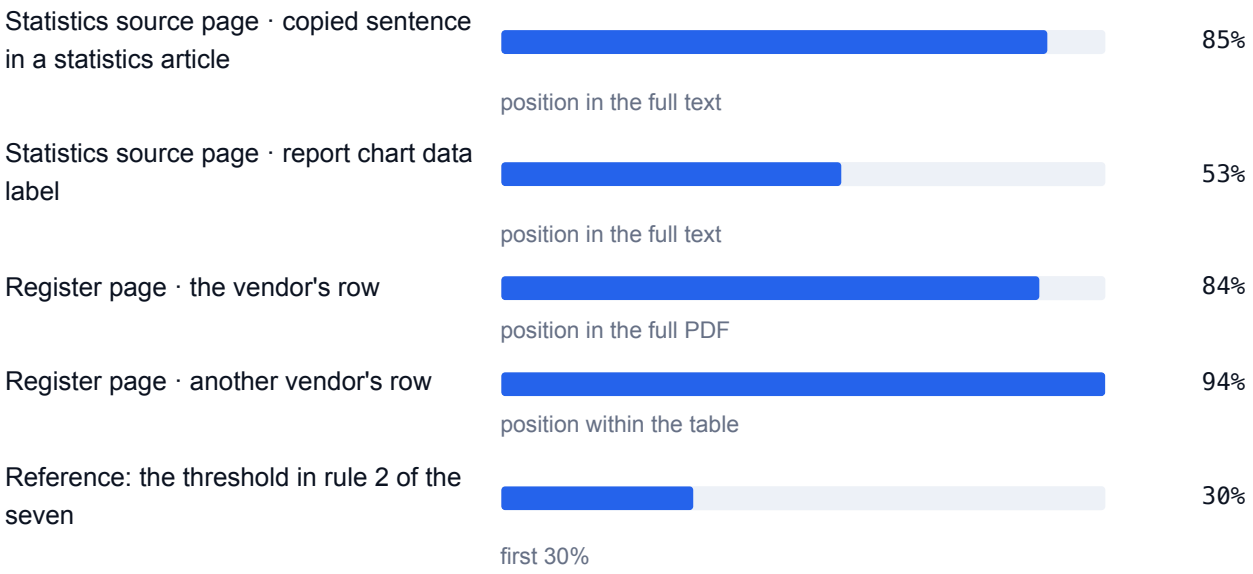


Figure: On statistics-source pages and register pages, the copied sentences all sit past the halfway point, so of the seven rules, "put the core facts in the first 30%" does not hold for these two types

Why it doesn't hold: a statistics question only has one correct answer, and AI will keep digging until it finds the sentence carrying a year, a survey name and a sample size; a register question wants a specific row, and however long the table is, that's how deep it will dig. So don't shave words off these two types on the instinct of "put the conclusion up front" — write out the methodology section and the tier explanations in full. Full text of the seven rules: see → General Edition 5.7 Seven rules that hold for every page type; sample in Appendix A.4.

5.29.1 Statistics source page (pt19)

- **Intent**: statistics-type questions, such as "what percentage of people X" or "average X" or "how many businesses use X".
- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: for the academic-paper version, title gives the journal + year + paper's full name, with the paper's own name carrying a year, region and population; for the report version, every chart's caption gives the full basis; for the research-blog version, H1 can be short, but the numbers must be in the first two sentences of the body.

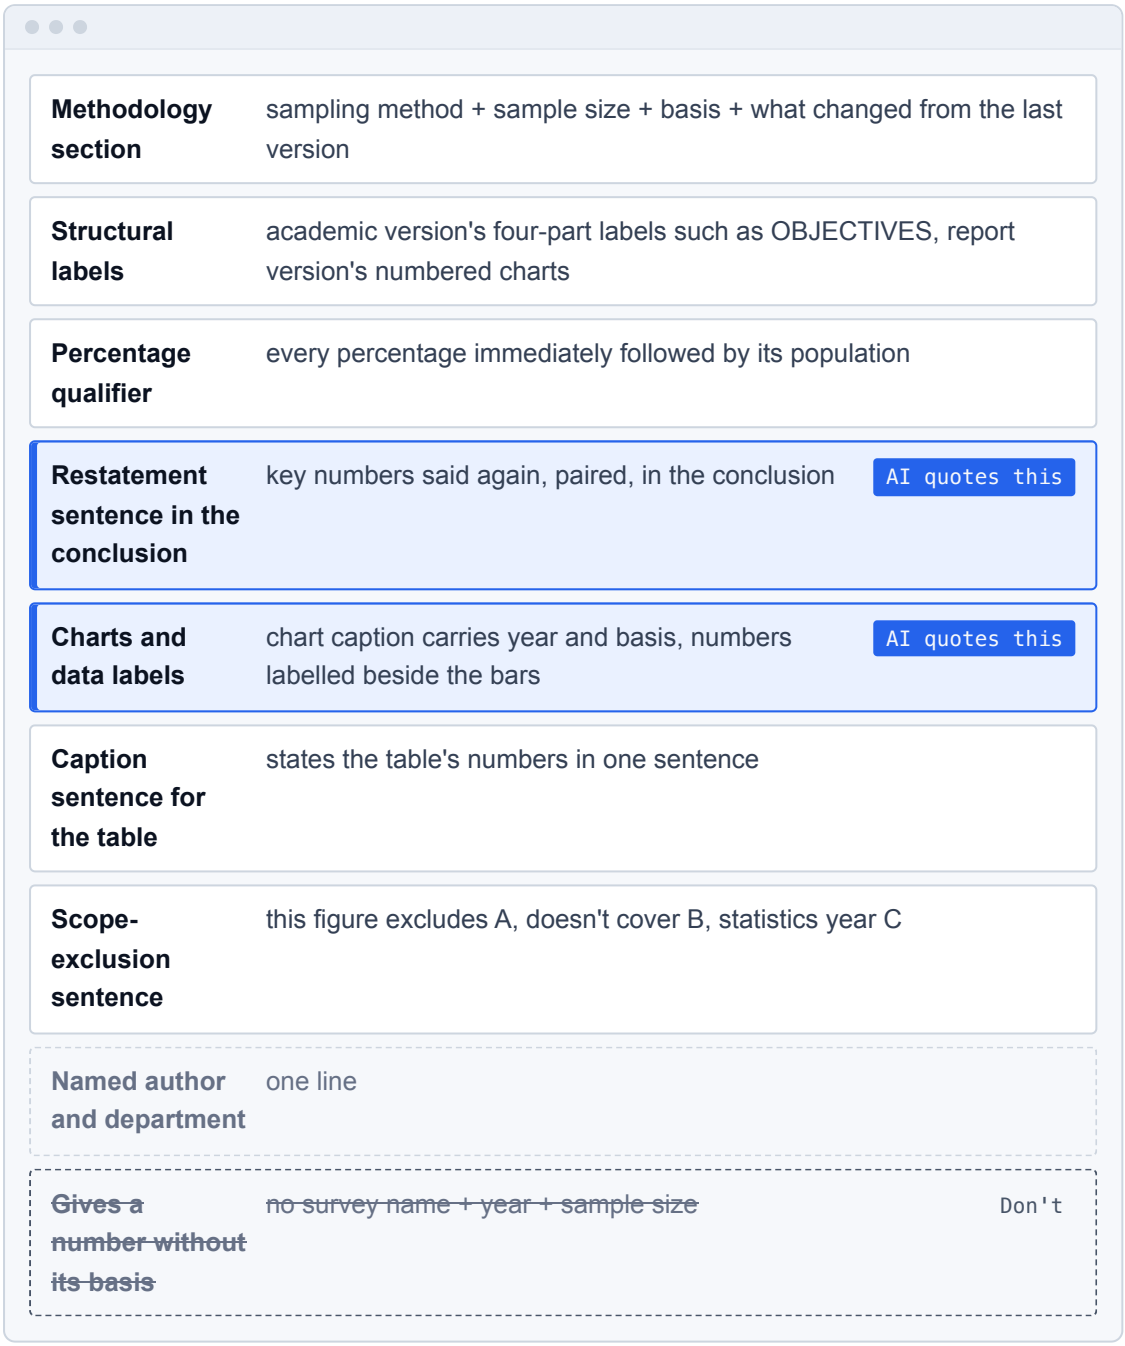


Figure: Statistics source page; AI copies the paired restatement sentence in the conclusion and the chart caption that carries the basis

- **Standard column headers**: follows the source's own basis; commonly `<grouping dimension> / <metric 1> / <metric 2>`, one row per group.
- **Sample sentences**: own — "<year> <survey name> (n = <sample size>, <sampling method>) shows: <subject>'s <metric> is <value>. This figure excludes <A>, doesn't cover <B>." Market — where older figures on the same question use an older basis, say so yourself and explain why they should be discounted. The scope-exclusion sentence itself has been copied whole before, so don't leave it out to save words.

When there's no official number, you can self-build a first-party data report page. It counts as the same type as the statistics source page; here is its self-build spec (an external teardown, not measured):

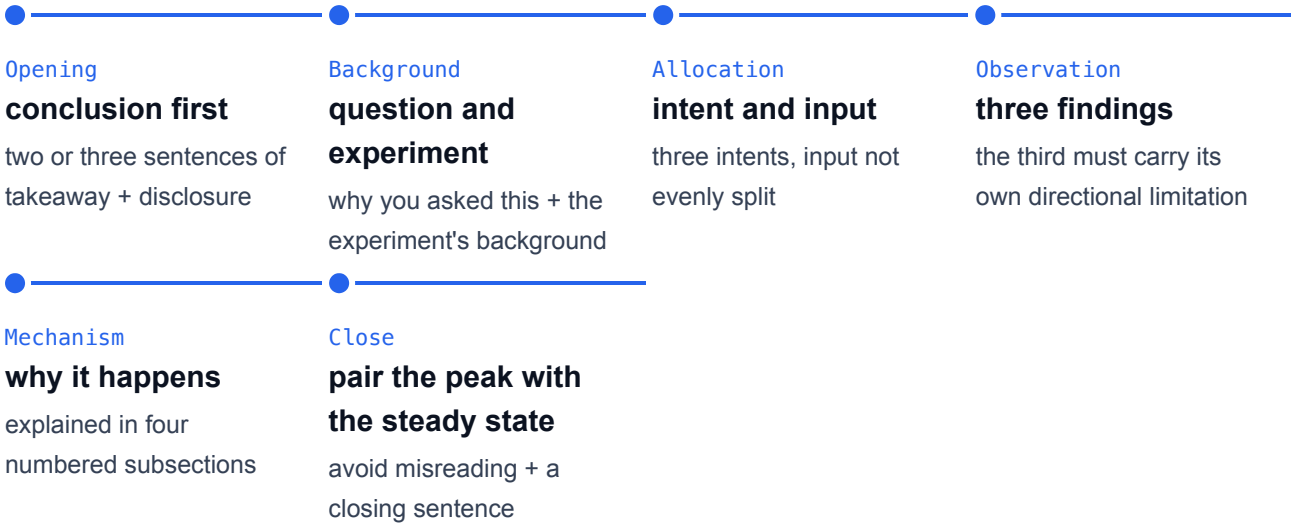


Figure: A self-built first-party data report page gives the conclusion first, then background and three observations, and closes by pairing the peak with the steady state (external teardown, not measured)

What the figure cannot show: 9–10 numbered H2 sections across the whole piece, 3,500–4,600 words, 5–7 small decision tables, 3–5 images, written in the first person in the voice of the operator who ran it; title formatted as `The <name you coin for the phenomenon> Experiment — <brand> Data Report`; lock the opening's conclusion in with a metaphor. This spec has not been measured for citation; treat it only as a starting template.

> **Example** (medical): a paper abstract from a national adult oral-health survey had one sentence copied — the sentence in the conclusion that restates two key numbers side by side — not any of the raw figures in the results section. Evidence in Appendix A.4.

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Specific to this page type: your own results data must clear the six checks first (→ General Edition 5.10 Shared parts for all page types: six checks for numbers, no public price, billing units, dates and schema) and get a written compliance sign-off; on the strict side, never write a success rate. A sign-off cannot turn something banned as Statute text into something publishable — see → General Edition 0.4 The three labels for compliance sentences, and the stop-and-escalate rule.
- **Most common mistake**: giving a number with no basis. A number with no survey name + year + sample size gets demoted; conversely, clearly stating "there's no official number" may itself be copied as the answer.

5.29.2 Register / approved list page (pt20)

- **Intent**: "which organisations are recognised", or a location question that, underneath, needs a whole batch of vendor names at once. No side may self-build this, but all three should do two things: get onto every official and platform register you can reach, and write a good self-supplied blurb wherever the register allows one (see below).
- **Leg**: ChatGPT (AI Mode not measured).
- **title and H1**: `<organisation><list name> Listing / Register`, no year, no "best".

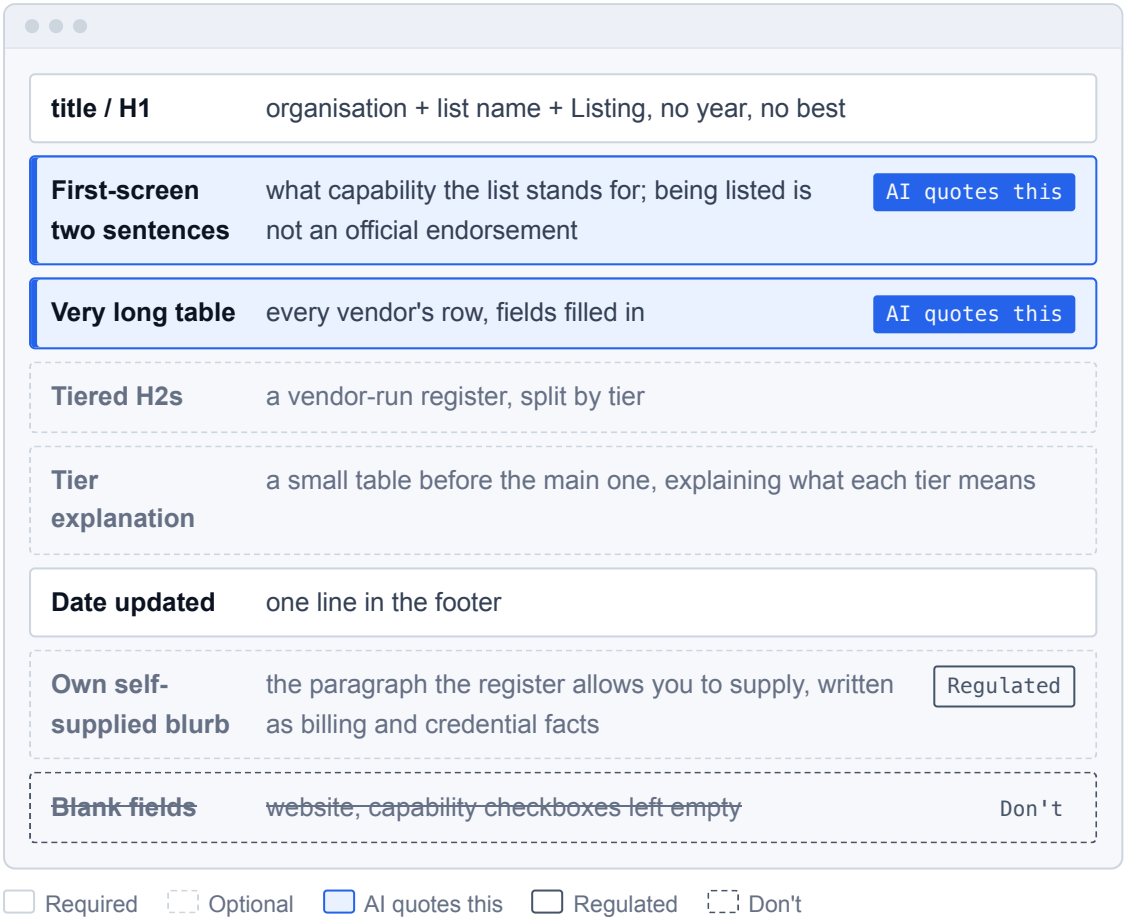


Figure: Register page; AI pairs the row deep in the table with the disclaimer on the first screen into one sentence

- **Standard column headers**: depends on the source; commonly `S/N / <name> / website / contact`, or capabilities ticked off item by item.
- **Sample sentences**: own — not applicable; this type can't be self-built, all you can do is fill in the fields on your own row. Market — on your own site, write "<register> lists this company under <capability>, since <year>, tier <X>", and link back to the register page.
- **Self-supplied blurb**: some registers let you write your own blurb. This is a legitimate channel for putting your own sentences onto an authoritative domain — a self-supplied blurb under an authoritative domain gets cited too. Write this paragraph as billing basis and credential basis — all checkable facts.

> **Example** (payroll software): on an official vendor register, what got copied was a combination of fields from a row deep in the table, immediately paired in the same answer with the disclaimer sentence from the first screen: the first half came from deep in the table, and the qualifier in the second half came from the first screen. Evidence in Appendix A.4.

- **Compliance**: availability by side: see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side. Cell-by-cell wording: see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers, → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Specific to this page type: the self-supplied blurb on a register is bound by the same price rules as your own site; on the strict side, give only a fixed final price, with the rule text in the industry edition's appendix.
- **Most common mistake**: getting onto the register but leaving your row's fields blank. AI copies exactly that row — a blank field is the same as not being on it at all.

Chapter 5

Writing pages (4): 4 unverified page types, and checks before you publish

Five steps before each page goes live, the red line on lying, and common mistakes

5.30 The four tier-C types: unverified page types

5.31 Handing over: five steps per page, the red line on lying, and self-checks

5.32 Master table of common mistakes

5.30 The four tier-C types: unverified page types

What you'll do in this section: recognise that these four types have been seen only in an external teardown and that we have measured none of them ourselves. When you're done, you'll have a blueprint for the Buyer's selection framework (pt43), which is the strictly regulated side's lawful substitute for getting onto the list shelf, plus a one-line scheduling call for each of the other three types: all four types come after tier A and tier B.

> Note **None of the four types has been measured by us**: not one of the citations we measured supports them, and the teardown target's own pages of these types don't even rank organically. The teardown measures organic rankings and we measure AI citations; numbers from the two rulers cannot be read as one. We still write them up below so that the catalogue is complete and the next round can take them straight into testing. Evidence: see → General Edition A.4 Page-type measurements (2): tier B and C samples, and the AI Mode retest list.

pt43 Buyer's selection framework

title	How to Choose a <provider type> for <scope>
Quick answer	2–3 sentences answering the title question directly, followed by a Disclosure statement
H2 What does <thing> measure?	Opens with one punchy sentence, then 5–6 bold labels
H2 What a useful <thing> should show	Numbered, 5–7 items, one sentence each
H2 Research snapshot	Evidence paragraph: what the status quo is missing
H2 Evaluation dimensions	5 "Label:" paragraphs, no bullets
H2 Provider comparison matrix	Column headers: Provider / Dimension 1 / Dimension 2 / Best for `Regulated`
Per-provider card	One card per provider; write each limitation as a buyer task `Regulated`
H2 How to interpret common results	Short sub-sections without bullets, each opening with a label
H2 How to run a defensible baseline	Numbered, 5–10 imperative steps
H2 What a free diagnosis cannot prove	100–160 words, spelling out what the diagnosis cannot prove
H2 Questions to put in the RFP	8 items, each a complete question ending in a question mark

H2 Final recommendation	When to switch from X to Y, 2 paragraphs
FAQ + Method and sources	5–6 questions, at most 2 naming your own brand

Required Optional Regulated

Figure: Block order for the Buyer's selection framework (not measured as a whole; spec taken from an external teardown; no block is marked cite); on the strict side and when switch E is on, the two blocks marked law are removed entirely

When to use: buyers ask "how do I choose a <provider type>?" It looks like → General Edition 5.18 Eligibility and process page (pt07) but the intent differs: pt07 answers "can I do this, what does it take, how do I apply"; this type answers "how do I choose a provider". Which leg cites it has not been measured. Going by the two-leg pattern, we infer it leans toward AI Mode, so only a follow-up measurement on the AI Mode leg can give the first piece of evidence: any citation moves it up from tier C; 0 citations keeps it at n = 0.

Why this is the only tier-C type we draw a blueprint for: the strict side cannot self-build a list (see → General Edition 5.16 List page (pt05)), and a "how to choose X" criteria-based selection guide (it lists criteria, does not rank and names no one) is its lawful substitute for getting onto that shelf. **Not measured, 65% confidence**; we recommend scheduling one per industry as a control.

Compliance points specific to this page type: how comparisons and lists are written on each of the three sides is covered only in → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Applied to this type: on the strict side, and whenever switch E (peer comparison banned) is on, delete the two blocks marked law in the diagram (the comparison matrix and the per-provider cards), and keep only three blocks in the main body: evaluation dimensions, the limits of the diagnosis and the RFP questions. The lightly regulated side may include how fees are charged and a verification checklist, and its result numbers must pass the six checks for numbers. The unregulated side uses the blueprint as it is, per-provider comparison included.

Length and sample sentences: the external teardown gives 2,100–3,700 words, 11–19 H2s and 1–2 tables. Treat these as a reference only: set the length by the page type and stop when the page is complete (see → General Edition 5.7 Seven rules that hold for every page type). Sample sentence (not measured): write the own sentence as "a firm fact about us on <a dimension> + the check date".

Most common mistake (inferred, not measured): writing it as "here's how to choose, and by the way we're the best". Restraint is the precondition for this type: the FAQ names your own brand in at most 2 items, and one item states "always confirm current scope with the vendor directly".

The other three types: one row each

Page type	The measured neighbour to do first	Why it comes later	How to schedule it
pt44 Concept pillar page	pt12 Definition page	pt12 has measured citations; pt44 has 0	Do pt12 first; pt44 is still to be verified. For how to split it from the Comparison page, see the figure below
pt45 News and policy explainer	pt34 Change notice / old-vs-new page (about changes to your own business or your own jurisdiction)	Citations for rule-type questions land almost entirely on the official original page; commercial pages get 0 citations. AI is more likely to cite the original page directly than your explainer	Schedule it for defence and to piggyback on the rule-question leg, not as a traffic page. On the regulated sides, the strongest case is an explainer of a regulation amendment (for wording by side, see the page-type row in → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions)
pt46 Market observation page	None	The weakest evidence of all, and there is counter-evidence: the teardown target's own articles of this type are dead ends in the site's internal linking and do not even appear in the blog index	Without first-party data, don't write it and don't schedule it. A service provider may run it on its own blog for its own use, but there it works as a brand asset, not as a page that gets cited

pt12's card is in → General Edition 5.23 Tier-B questions and rules types (1): regulatory obligations, definitions, pt34's is in → General Edition 5.25 Tier-B questions and rules types (3): schedules, collected FAQ, policy hubs, change notices, misconceptions. The blueprints for these three types are not in the main text; each is labelled in full "not measured, spec taken from an external teardown" and placed in → General Edition B.5 Blueprints for unverified page types (for industry editions to include). When you test the four types, test 43, 44 and 45 first; 46 comes last.

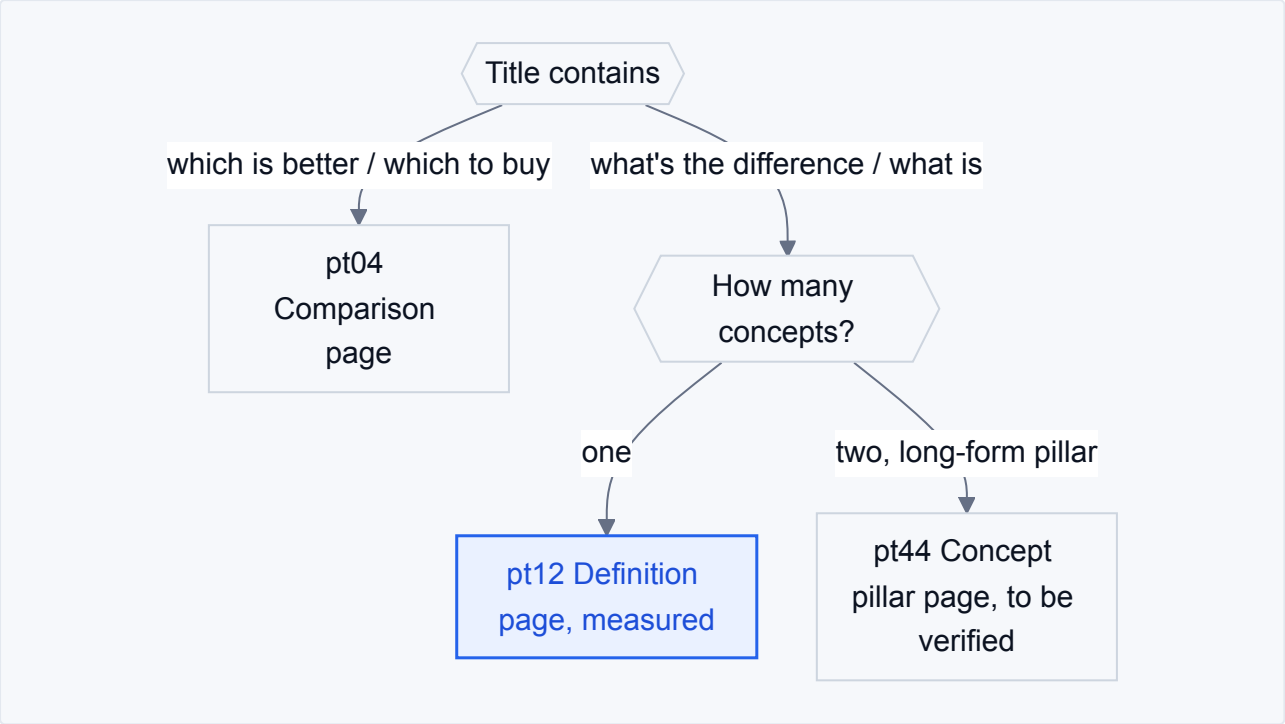


Figure: The title wording decides whether a page is a Comparison page or a Concept pillar page; only one page may take a given term, or the two pages cannibalise each other

The same applies to English titles: "which is better" belongs to pt04, "what's the difference" belongs to pt44. If the two can't be cleanly split, merge them into one page.

> **Example** (medical, not measured): the title "Dental implants vs bridges: differences, indications and cost difference" contains "differences", so it belongs to pt44, not to the Comparison page. The strict side can build it too, but it compares procedures only, not clinics.

For pt45 to piggyback on the rule-question leg, there is only one way to write it: quote the rule word for word, date it, link back to the original page, then add one sentence on "what it specifically means for your line of business". In our measurements, what got copied was always these four things: "who / from which date / what must be done / the penalty for not doing it"; not one sentence of opinion was ever copied (inferred, with indirect support from measurements). Sample sentence (not measured): "As of <effective date>, <the new rule> requires <who> to do <what>; the direct impact on <a category of people in this industry> is <one sentence>."

5.31 Handing over: five steps per page, the red line on lying, and self-checks

What you'll do in this section: take every page through the same five steps before it goes live (brief → draft → non-author fact check → licence-holder compliance sign-off → launch acceptance, about 7 hours for an English page), and stop right there if any step gets stuck. Before launch, grep the whole

site once against the banned-word list; before handing over, go through the self-check list for every page, and any question answered "no" sends the page back.

1 · 90 minutes

Write brief

Fewer than 5 exclusive facts: stop on the spot, change cluster

2 · 3 hours

Write draft

Missing a source: write a placeholder sentence, never make one up

3 · 60–90 minutes

Fact check

Checked by a non-author; if any item can't be found, it doesn't go live

4 · 30–45 minutes

Compliance sign-off

Licence holder checks the draft item by item and signs; unsigned means not live

5 · 45 minutes

Launch acceptance

If a non-cached load doesn't show the new H1, or the page source has no body text: roll back

Figure: Five steps per page; the next step does not start until the one before it is done

What each step delivers, and who accepts it:

Step	Who does it	What it delivers	Who accepts it
1 Brief	The brief writer drafts it; the person running probes produces the gap table	`brief-<intent-cluster>.md` : the cluster's 3–6 intents in the buyers' own words, H1/H2 questions, 5 or more exclusive facts + fact-table row numbers, 3 or more outbound-link candidates, a gap table of cited sources, and a list of banned phrases picked out by side	The owner's signature
2 Draft	The page writer	`draft-<slug>.md` , with frontmatter carrying intent_ids, page type, language and side; wherever a source is missing, write `[Source needed]`	—
3 Fact check	Non-author (a different person)	`factcheck-<slug>.md` : an item-by-item verdict + where the evidence is + the checker's signature and date	The checker signs and is personally responsible
4 Compliance	The owner + the licensed signatory	The compliance memo with every item ticked + the scan result against the Chinese–English banned-word list (hits must be zero) + the written sign-off	The signatory

Step	Who does it	What it delivers	Who accepts it
5 Launch	The person who edits the site	The site change record + the `accept-<slug>.md` acceptance sheet + the page-register entry (intent cluster, buyer type, URL, launch date = split day, version) + one off-site target	The brief writer

A Chinese page counts as about 0.8 of a page of work and is scheduled on a separate day, through the same five steps. Only step 3 differs: it does not look for new evidence; it only checks that the fact-table row numbers and figures match the English page. None of the other four steps is cut: rewriting the question phrasing and writing independent answer sentences are done in step 2, and the three Chinese-page hard gates (the H1 contains the country name; the amount in the first paragraph carries its currency and says whether tax is included; the body names at least one real station or area, see → General Edition 5.9 Content form, FAQ, Chinese pages and other-language pages) are checked in step 5.

At step 5, when you accept the page, also log one off-site target, "which source should cite this page", and hand it to the off-site queue. If you write the page but push nothing off-site, the output never leaves your own site (for the other constraints on assigning tasks, see → General Edition B.4 Headers for work orders, ledgers and monthly reports).

Self-check list

Before handing over, check every page in this order: "positive gate → form → numbers → compliance → launch". Any row answered "no" is handled as the rightmost column says.

Ban gate only (no positive gate)		Positive gate added
The more you delete, the safer	→	Count exclusive facts first; go on only once you have 5
The page you hand over is identical to a competitor's site	→	Fewer than 5: sent back on the spot to change cluster, no last-minute patching
AI has no reason to cite a page with no exclusive facts	→	Everything kept is a fact buyers will use to decide

Figure: With only the ban gate, the editor's easiest move is to delete everything; the positive gate forces you to prove first that the page kept something

Gate	What to check	What to do if the answer is "no"
Positive gate	First run the ban scan; sentences it removes don't count. Then list, one by one, the checkable facts on this page that can't be found anywhere outside your own site, and check that there are at least 5, each of which would change a buyer's choice between candidates (generic information every peer has doesn't count). You can name which independent source (an official register, a government page, a third-party review, a ranking body) confirms what you are a fit for or your credibility, with its name and URL written out	Fewer than 5: send it back to step 1 to change the intent cluster; no patching facts in during review, and no keeping a sentence the ban gate hit just to make up the count. Can't name an independent source: fill in the off-site source plan before launch
Form	The conclusion sits right under the H1 or in the first sentence of the first H2, and both sentences are in the first 30% of the body (on the strict side, and whenever switch E (peer comparison banned) is on, check that the second sentence is an official-basis sentence, see → General Edition 5.6 Two legs, two kinds of sentence); the first sentence of every paragraph makes sense out of context, and none opens with "it depends" or "it's complicated"; one URL fully answers only one intent and one buyer scenario; length and how H2 questions are written follow → General Edition 5.7 Seven rules that hold for every page type, and the conclusion block's spec follows → General Edition 5.8 The conclusion block, the density gate and the source gate	Fix it on the spot. A missing conclusion block, or one that doesn't match the body, is not an on-the-spot fix: send it back to step 2 to be completed
Numbers	Every number carries its unit, threshold and how it is counted; the numbers in the conclusion block match the table, number for number; every fact has a source and a check date, and anything out of date is marked historical or deleted; Ctrl+F in the page source finds your own prices, results and specs; the fact-check file is complete, and prices, registration numbers, rule text and statutory deadlines have been read back word for word by a non-author; the organisation's operating figures have a ledger export and a written sign-off	Wording problems: fix on the spot. Any item missing from the fact-check file: does not go live

Gate	What to check	What to do if the answer is "no"
Compliance	First decide the side, then scan item by item down your column in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers and → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions, including a word-by-word scan against the Chinese–English banned-word list; scan the FAQ and the end-of-page Q&A as well, because they are not an exempt zone. Also scan for the general red lines: traffic multiples, ROI, conversion lift, guaranteed first place, exclusive, "independent review", and method leaks such as small numbers from a single question, how the probes are run, or the question list. The page passes professional review (pages that touch personal safety carry a reviewer's signature; remedy pages follow the order in → General Edition 5.19 Remedy and second-opinion page (pt08); statutory deadlines cite the original text). The compliance-reviewed draft carries the licence holder's written sign-off, and the banned-word scan shows zero hits; before launch, grep the whole site once against the banned-word list	Delete every hit and fill the same slot with the replacement wording from 5.2–5.3; the delivery note states "deleted X, replaced with Y"; after deleting, go back to the positive gate and recount. No written sign-off: does not go live
Launch	A non-cached load shows the new H1; view source shows the body text (no reliance on JS); fetches as the four identities match; IndexNow has been pinged and the receipt kept; the visible "Updated" line matches `dateModified` ; the page register has logged the intent cluster, buyer type, URL, split day and version; the off-site target is logged; a test enquiry has been recorded	Roll back
Chinese pages: extra checks	The three Chinese-page hard gates (see 5.9); whether each H2 is in the Chinese buyer's own words or a literal translation of the English heading; whether the answer sentences were written independently or translated	Any hard gate missing: send it straight back, no fixing on the spot. An H2 or answer sentence that was translated: send it back to step 2 to be rewritten

> **Example** (education and training): grep the whole site for the six "approval"-type words "approved / accredited / endorsed / supported / validated / recognised", and replace each instance with a sourced factual sentence or with "registered".

The red line on lying

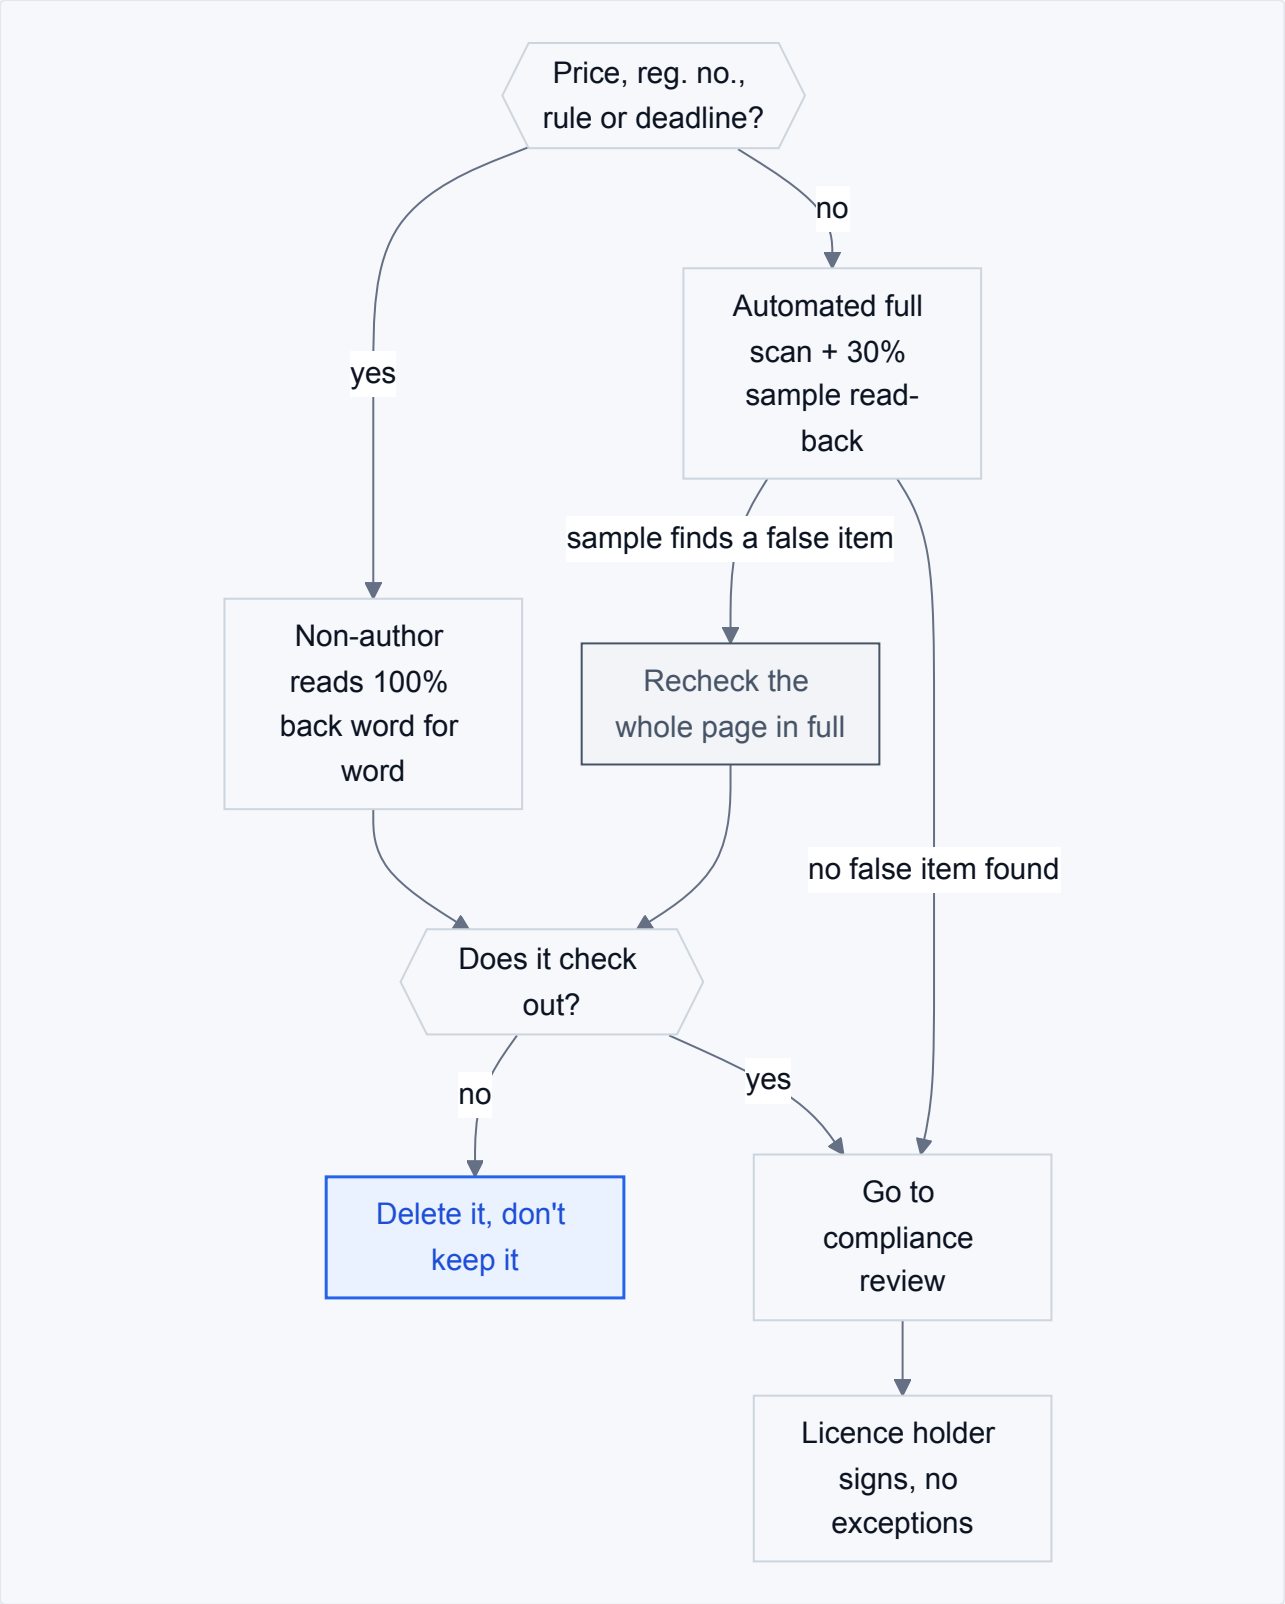


Figure: The red line on lying: better to leave a fact out than keep one that hasn't been checked

Neither the non-author fact check nor the compliance review may be skipped under any scheduling pressure, and not a minute of either may be cut. The sign-off has limits too: **A written sign-off**

resolves only the three stop-and-escalate situations; it never turns something banned as Statute text into something publishable (see → General Edition 0.4 The three labels for compliance sentences, and the stop-and-escalate rule).

5.32 Master table of common mistakes

What you'll do in this section: go through this master table before handing over. Fix, one by one, the mistakes in the figure below that aren't written up elsewhere; for mistakes already covered in the page-type cards and neighbouring sections, this section keeps only one line and a link and does not repeat them.

Wrong way		Fix
The conclusion block says "this article will introduce", and the first screen is all lead-in	→	Give the own sentence directly under the H1, and a number in the first sentence of the first H2
The first-screen summary and the table numbers don't match	→	Take the conclusion-block numbers from the table, and check them number by number before handing over
Prices and results appear only in images, JS or a booking widget, or switch by IP	→	Put all of it in the first-screen HTML
The product detail page body gives no actual price	→	Show the real price on the first screen, matching the feed, structured data and checkout
Writing the page without logging an off-site source target	→	Every page states the third-party page that should cite it, who handles the correction or submission, and the compliance approval status
A separate page for every SKU and every phrasing variant (not measured)	→	One intent, one URL; variants go into H2s; one question, one page only for genuinely independent questions
Changing only the update date, not the content (rule of thumb)	→	Keep the price check date, regulation check date and date of substantive change separate; only a substantive change moves `dateModified`
Machine-translating the Chinese version directly	→	Validate Chinese-language demand first; rewrite the question phrasing and answer sentences, and put it through compliance separately
Calling it done because ChatGPT, asked once on day 7, can quote the page back (rule of thumb)	→	Freeze the questions, run baseline and retest over several rounds, and log times named, position and cited URLs separately

Figure: The mistakes that recur most often when writing pages, and how to fix them

The two rows marked rule of thumb are not measured, the one row marked "not measured" is inferred, and the other six rows have measured evidence behind them. Evidence: see → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests.

Mistakes covered elsewhere: one line and a link

Mistake	Fix	See
Padding the word count to reach a minimum	Set the length by page type; stop when the page is complete	→ General Edition 5.7 Seven rules that hold for every page type
Writing the H1 as a question, or every H2 as a question	The H1 is a statement; only the first H2 is a question	→ General Edition 5.7 Seven rules that hold for every page type
A comparison page that compares only your own models, or puts the verdict mid-page	Competitors must be included; the verdict goes in the first sentence of the first H2	→ General Edition 5.15 Comparison page (pt04)
Turning one question, one page into one page with ten questions	One question, one URL	→ General Edition 5.17 One question, one page (pt06)
Several H1s on one page	One H1 per page	→ General Edition 5.19 Remedy and second-opinion page (pt08)
A remedy page that starts with the cause, is written as a self-test checklist, merges the two skeletons, or has an H1 blaming "another business did it badly"	Personal-safety type: order it "what to do now → red flags → what not to judge yourself → examinations → pathway → cost"; split the personal-safety type and the rights-deadline type into two pages; the subject is the buyer's situation	→ General Edition 5.19 Remedy and second-opinion page (pt08)
301-redirecting a byline after the person leaves	Keep the person profile page and mark it "formerly", or 301 the whole page to the directory; on business pages, only change the name to the person currently responsible	→ General Edition 3.3 Person pages and six-trace alignment
Writing your own "best" list to answer "who's best"	Neither leg cites a self-built list, and on the regulated sides it also breaks the rules; use business profiles, third-party rankings, the product feed and the organisation facts page instead	→ General Edition 5.6 Two legs, two kinds of sentence
Putting a market-range sentence into the feed or the Offer schema (not measured)	The feed and schema carry only your own real price	→ General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers

Chapter 5

Writing pages (extra): producing hundreds of pages, from studying a competitor to launching in batches

Picking who to learn from, choosing topics, the writing pipeline, checking the numbers, launch and review

5.D1 The seven steps at a glance, and what to tell clients

5.D2 Pick the right study target: the competitor that wins on on-site form

5.D3 Take the architecture apart until you can write from it

5.D4 Correct with measurement: the competitor sets the skeleton, measurement sets the details

5.D5 Topic selection: turn every head and long-tail keyword into a list with one main query per page

5.D6 The production pipeline: three agent passes per article, plus automated lint gates

5.D7 Six traps we hit: build the safeguards in before you start

5.D8 Source audit: trace every number in data reports to its source

5.D9 Launch in batches: from launch day to the keep-or-cut decision on day 45

5.D10 Add new keywords every month: new keywords come in, a person reviews before anything goes out

5.D1 The seven steps at a glance, and what to tell clients

What you'll do in this section: see the order of the seven steps for producing hundreds of pages in one go, understand how this chapter divides the work with the earlier writing-pages sections, and get an approved line on how far you can go when you talk to clients.

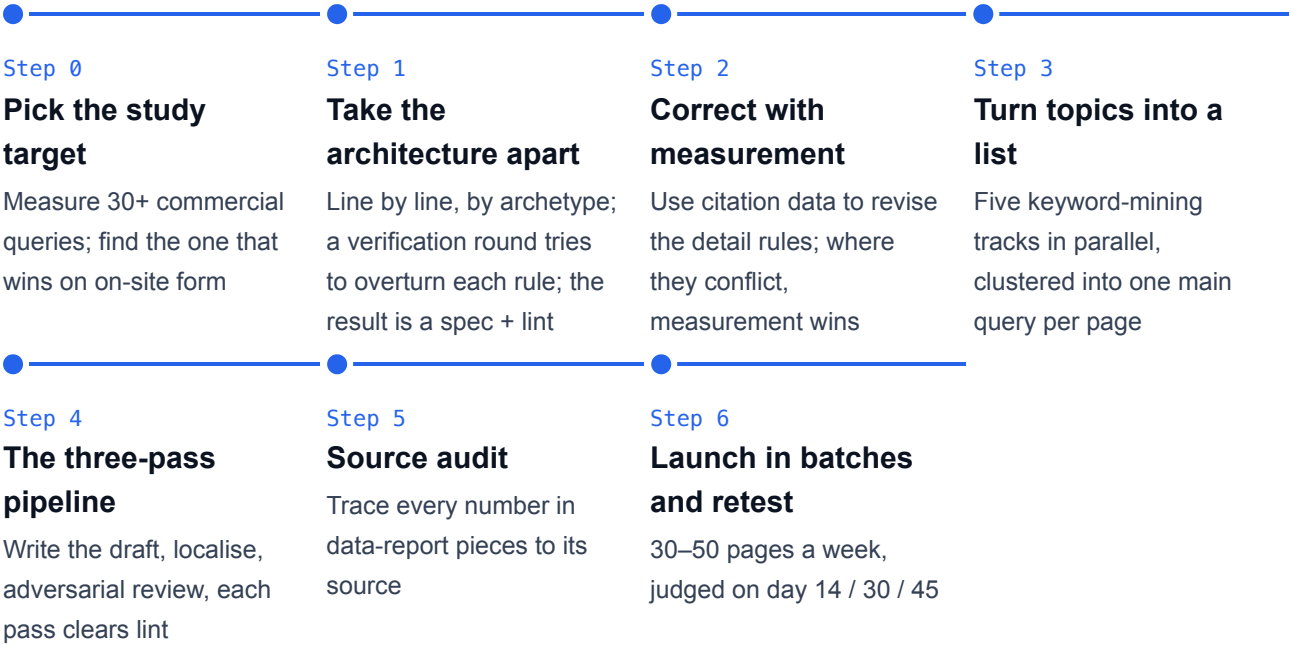


Figure: the seven steps are a one-way chain: each step's output is the next step's input; skip one and the mistake gets copied into hundreds of pages

Step 6 is followed by a monthly loop: add a new batch of keywords every month, have a person review the finished pages, then release them in batches — see 5.D10.

How this divides the work with the earlier sections. Which page type to pick, what each side may write, and the general rules for every page are still governed by → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side through → General Edition 5.10 Shared parts for all page types: six checks for numbers, no public price, billing units, dates and schema. This chapter handles one thing only: when you need to produce dozens to hundreds of pages at once, how to make every single one follow the same rules without drifting. The first batch of landing pages still clears its gates one page at a time under → General Edition 5.4 The 30-day writing order and the three mechanical gates; this chapter is for after that, for when you need to cover the long tail — it does not replace the ranking gate in 5.4.

How far you can go when you talk to clients. What this method has proven is the "production" half: in one replication project on a business's own blog, it was used to finish every page, and every page cleared its gates. The other half — "results", meaning rankings and AI citations after launch — has no retest data yet. So tell the client only "we produce using this method, and keep-or-cut is judged by the day 14 / 30 / 45 retests", never "this method has already been proven to earn citations".

> **Example** (a GEO agency's own blog) Using these seven steps, it finished 323 articles × en / zh / zh-tw (293 new articles + 30 rewrites), all of them passing the architecture gate.

5.D2 Pick the right study target: the competitor that wins on on-site form

What you'll do in this section: use three checks — measured rankings, domain age and backlinks — to filter down to one competitor whose "way of winning is copyable", write down which layer of keywords it wins on, which layer it loses on, and whether the page type it wins on is one your side can build. When you're done, you'll have a study target and one sentence stating which part of it, and only that part, you will copy.

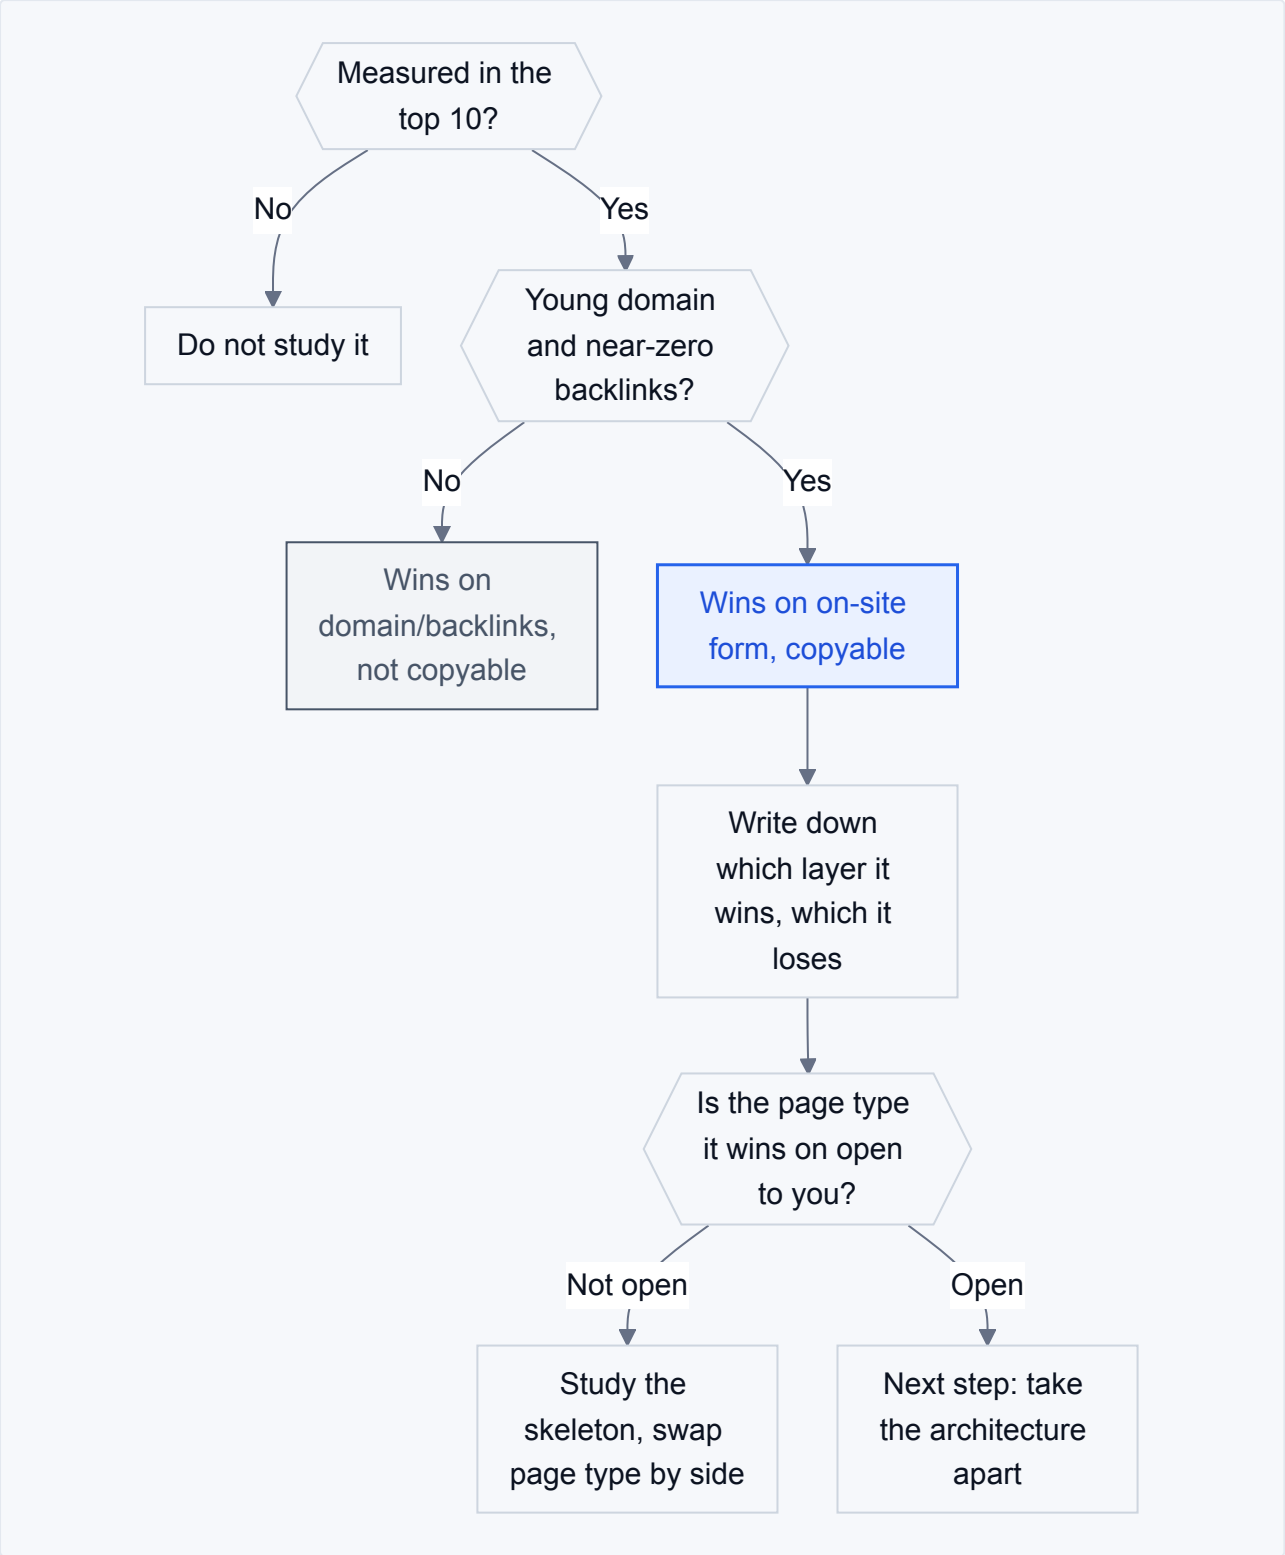


Figure: a young domain with near-zero backlinks still taking first place shows it wins on on-site form — only this kind is worth studying, and only study the layer it wins

How to check the three items. For rankings, use a SERP API such as Serper to measure 30 or more commercial queries, set `gl` to your target market, and see who is in the top 10; for domain age, check whois; for backlinks, look at third-party mentions, directory listings and inclusion in best-of lists. The criterion is not "whose content is better" but "who wins on your queries, and whose way of winning is copyable".

Why only "young domain + near-zero backlinks + first place" is worth studying. Domain age and backlinks are the two things you cannot copy in the short term. Once you rule those two out, if it still takes first place, the only thing left that can explain the ranking is on-site form — architecture, title, first sentence, tables — and these can be taken apart line by line. The reverse is also true: an old rival that wins on 20 years of domain age and backlinks will not reach its ranking even if you copy its writing — do not study it.

See clearly where it loses. A study target does not necessarily win on every layer of keywords. If it only wins on ultra-long-tail queries stacked with multiple modifiers, and never makes the top 10 on head terms, then what you are copying is a "long-tail matrix play", not "the ability to win head terms". This sentence decides which layer of keywords topic selection (5.D5) covers.

Check availability first for the page type it wins on. For the page type the study target wins its rankings on, check the availability column in 5.1 for your side first. If it is not open to you, only study its skeleton and where its sentences sit, and swap the page type to the one your side can build under the same intent, per → General Edition 0.3 Which page types are limited for you, and what to build instead.

> **Example** (a GEO agency) The study target's domain was 6 months old with backlinks ≈ 0, yet it took 5 #1 rankings, and Google AI Overview copied its first sentence word for word. But it only won ultra-long-tail queries stacked with 3 or more modifiers (best + industry + location + year); it did not make the top 10 on a single head term, and every page that took a ranking was a list page — exactly the type the strictly regulated side can use least. Evidence: see Appendix A.4.

A hard rule: the study target is never named. The study target is only used to study form; its name never goes into any article, list or comparison page, and no "us vs. that company" page gets built either; comparisons against a tool category or an approach are fine. This rule goes into a hard lint block: a banned-name list, plus a ban on any `<our-brand>-vs-*` slug.

5.D3 Take the architecture apart until you can write from it

What you'll do in this section: take every article of the study target apart into rules, line by line and grouped by archetype, then send a verification round to check those rules against the original articles and try to overturn them. When you're done, you'll have two files — a writing spec and a lint gate — neither one optional.

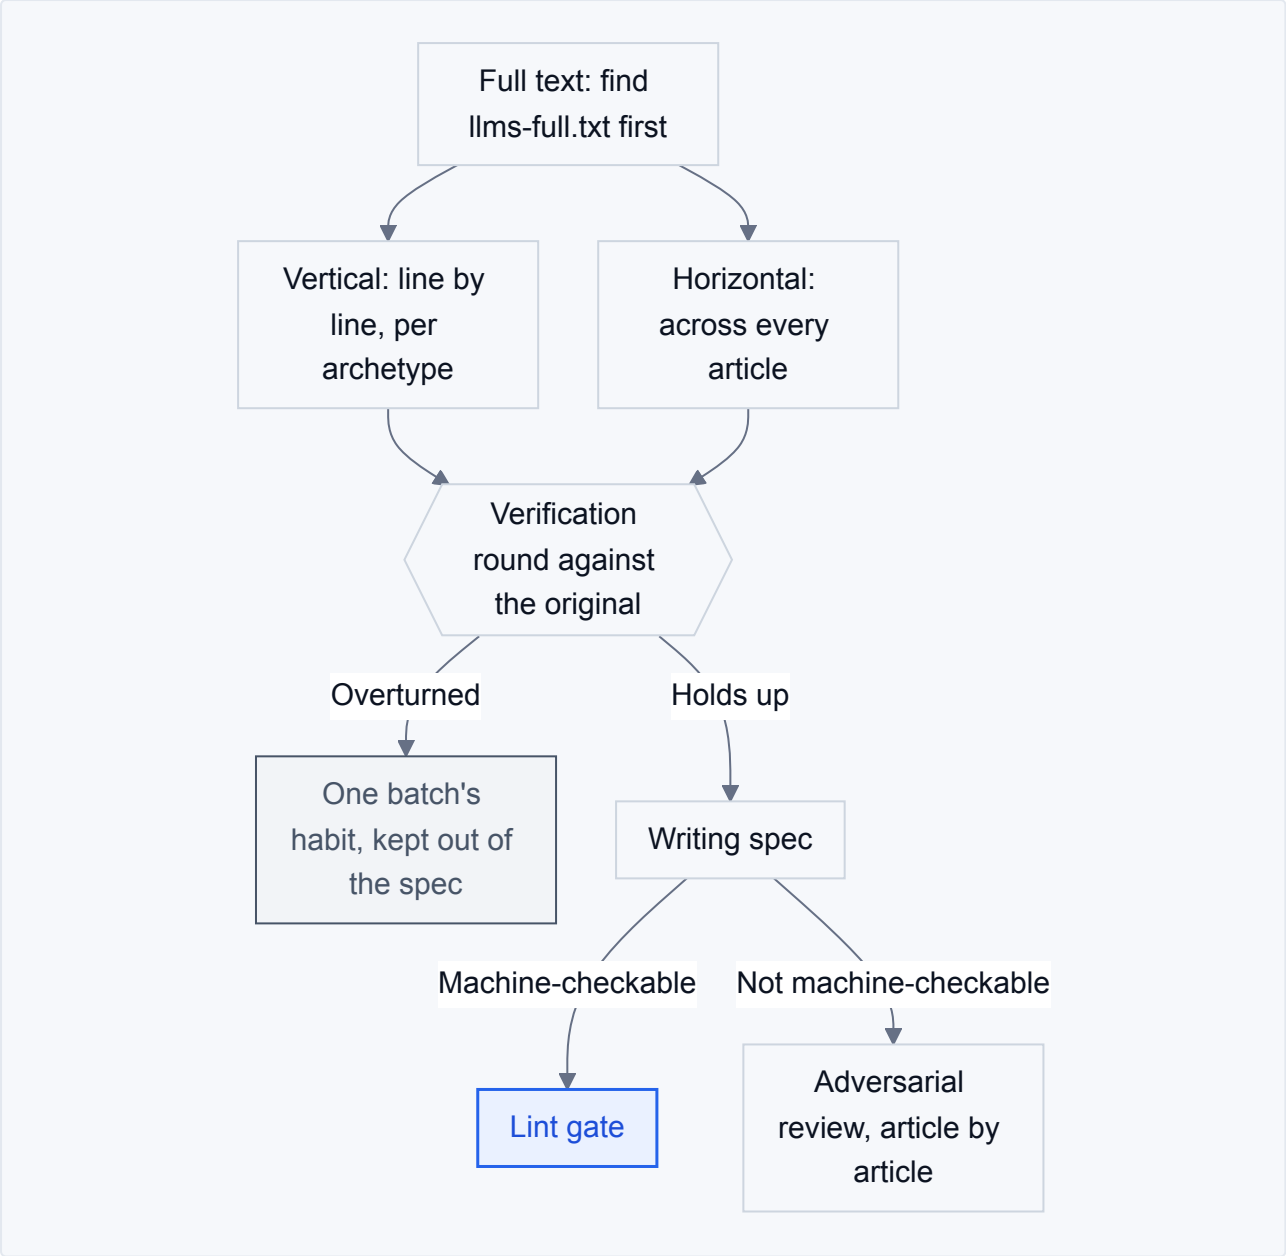


Figure: every rule taken apart goes through the verification round, and only the rules that hold up enter the spec; whatever in the spec can be machine-checked goes into lint, whatever cannot goes to review

Find `llms-full.txt` first. A competitor writing for AI citation often exposes the full Markdown of every article in this one file, so you get everything in one go without crawling page by page.

What each teardown covers:

Who takes it apart	Output
Vertical: group by archetype, one agent per group reads every article line by line	title formula, H2 skeleton, opening sentence, paragraph rhythm, table usage, FAQ, Method block, where the business mentions itself, verbatim boilerplate

Who takes it apart	Output
Horizontal: a separate agent looks across every article	the title and slug system, recurring fixed blocks, Chinese/English localisation differences, the technical layer (schema / llms / hreflang / publishing cadence)
Verification round: 3 agents	check against the original articles, with the sole aim of overturning the rules the first two teardowns produced

Why the verification round cannot be skipped. The person doing the teardown sees the habits of one batch of articles, and a competitor's writing can change from batch to batch. Skip verification, and "one batch's habit" gets written up as "a global iron rule", then copied into hundreds of pages by the pipeline.

> **Example** (a GEO agency) The verification round corrected 9 rules; among them, Oxford commas, abbreviations, summary length and bold usage genuinely differed from batch to batch in the study target.

What each of the two files is for. The spec is what people and agents write from; its content is "archetype × title formula × body skeleton × shared blocks × style rules". Lint is what the machine enforces: everything in the spec that can be machine-checked is written as a check item, and an article that has not cleared lint does not count as done. At a scale of hundreds of pages, no one can check formatting article by article, so nothing that can be handed to a machine is left for a person.

Lint must compare titles by word root, not by substring. "agencies" and "agency", "tracking" and "track" count as the same word. Compare by substring, and writers trying to avoid a repeat will twist the title into something awkward (e.g. "Agency and Service Providers").

5.D4 Correct with measurement: the competitor sets the skeleton, measurement sets the details

What you'll do in this section: take the measured citation data and go through every detail rule in the spec, change what conflicts with measurement, and add a qualifier to two of them. When you're done, every detail rule in your spec is clearly marked as either "the competitor's habit" or "what measurement showed".

What measurement showed	What the spec changed to
Copied sentences are almost always the first sentence of a paragraph	Every H2's first sentence must stand alone as an answer, carrying a number
About 73% of copied facts sit in the first 30% of the page	Lint checks the first 30% has at least 8 numbers (an empirical figure)
Tables get copied header row and all	Table headers carry a verdict word: Best for, Winner, etc.
ChatGPT looks for the primary source, AI Mode for second-hand pages	Every page carries two sentences: your own fixed price + the market range (unregulated side only)
Price-comparison and vs. pages are cited most; the competitor has neither	Add two archetypes: price comparison and comparison
Cited pages ran 62–9,000 words	Word count is only a hint; only list pages keep a floor
When the summary number and the table disagree, AI uses the table	The summary number must equal the table

Figure: each of the seven measured findings changes one detail rule in the spec; the skeleton follows the competitor, but when a detail conflicts, measurement wins

Why measurement wins when there's a conflict. The competitor's writing only tells you "this is how it happens to write", not "this is the writing AI most likes to cite". It ranks because its whole form works together; pull out any one detail rule on its own and that rule may not be why it gets cited. The correction draws on teardowns of cited pages across several industries and both AI legs; evidence: see Appendix A.3.

Qualifier 1: "two sentences per page" only holds on the unregulated side. The two legs look for different pages, so a page needs two kinds of sentence (see → General Edition 5.6 Two legs, two kinds of sentence); but whether the "market range" sentence can be written on a regulated side, and what to swap it for, follows the price row in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers and the comparison row in → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions — check your own column, then check your industry edition. Table headers carrying a verdict word such as "Winner" or "market range" should likewise be checked against 5.3 first on a regulated side.

Qualifier 2: "at least 8 numbers in the first 30%" is an empirical figure. It is based on that ~73%: of the 75 copied sentences that could be located, about 55 fell within the first 30% of the body. The denominator is the 75 that could be located out of 526 total citations, not all 526; the positions were

estimated by eye, so treat them only as a strong trend. The "8" is lint's threshold, not a number computed from this data set. Do not tell a client "put 8 numbers in the first 30% and you'll get cited".

The rest do not need a separate rulebook. First sentences of paragraphs, table column headers, and word count set by page type: the book already has general rules for these in → General Edition 5.7 Seven rules that hold for every page type and → General Edition 5.8 The conclusion block, the density gate and the source gate, and the spec simply cites them. The two new archetypes, price comparison and comparison, follow the blueprints for the Price family page types and for → General Edition 5.15 Comparison page (pt04) respectively; whether the comparison page is open on your side also needs checking against 5.1 first.

5.D5 Topic selection: turn every head and long-tail keyword into a list with one main query per page

What you'll do in this section: run five keyword-mining tracks in parallel, cluster the results into "one main query per page", then run three passes to catch what's missing. When you're done, you'll have an article list: one main query per row, a group of near-synonym secondary keywords, and one archetype.

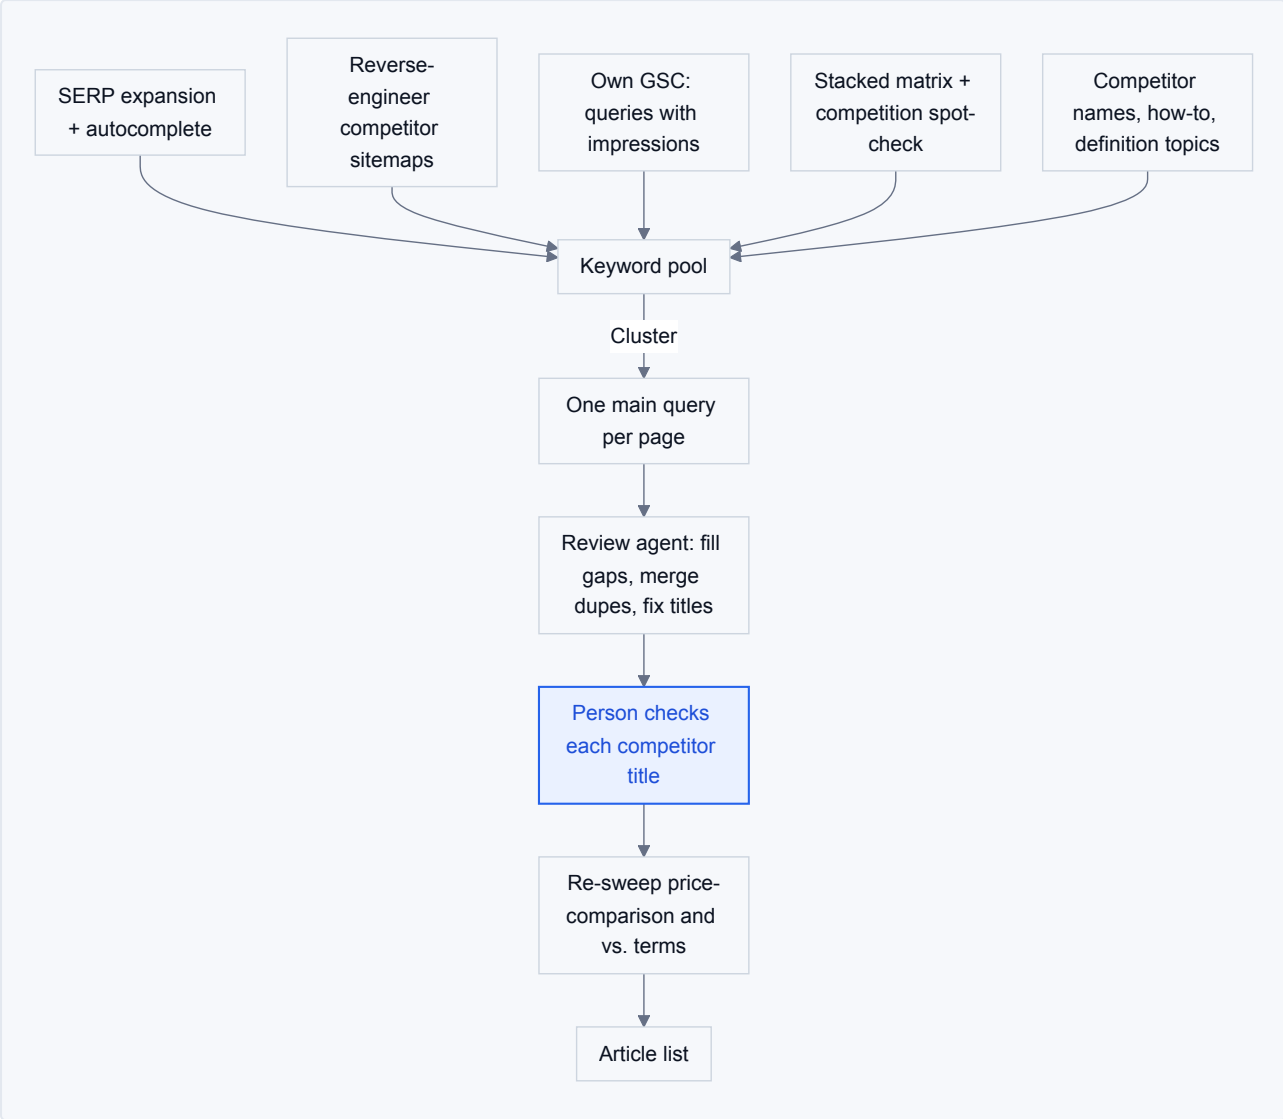


Figure: the five keyword-mining tracks cannot see each other and each mines on its own; after pooling and clustering, three more passes fill the gaps before it becomes the list

How each of the five tracks mines:

Track	How it mines	How it filters
SERP expansion	Seed keywords × Serper's relatedSearches + PAA; Google autocomplete (seed + a–z + market name / best / vs / cost / tools)	—
Competitor slugs	Crawl the sitemaps of competitors and tool vendors, and reverse-engineer from each blog slug the keyword it is targeting	—
Own GSC	Queries with impressions already recorded — the most genuine signal of demand	—

Track	How it mines	How it filters
Stacked matrix	[best/top] × category × industry × market × year × optional engine	Keep only combinations a real person would search; sample 60+ to test competitiveness (green / amber / red)
Name and definition topics	`{brand} review / alternatives / pricing / vs` , how-to, definition topics	Use autocomplete to drop terms nobody searches for

The two hard rules of clustering. Near-synonyms go into secondary; one intent gets only one page. This is the same rule as → General Edition 5.5 The page unit: one intent cluster, one page, and it matters even more at scale — otherwise, across hundreds of pages, you end up competing with yourself for the same keyword.

Why the human title check cannot be skipped. Clustering and deduplication are both done by an agent, and a wrongly deleted item throws no error. Having a person go through the competitor's title list line by line and ask "do we have a matching page?" is the last chance to catch a missing page.

> **Example** (a GEO agency) 1,476 keywords ended up as 293 new articles + 30 rewrites. During deduplication, an article that matched the study target 1:1 was wrongly deleted; the human title check recovered 9 articles in total.

Why price-comparison and vs. keywords need a separate sweep. The clustering agent tends to file them as "opinion pieces", when they should actually be price-comparison pages or comparison pages. Get the archetype wrong, and the three downstream pipeline passes will write the whole page on the wrong skeleton.

Check every row's archetype against 5.1's availability column for your side before you start writing. In the stacked matrix, combinations starting with best/top map to list pages, and vs. keywords map to comparison pages; both types have restrictions on regulated sides.

5.D6 The production pipeline: three agent passes per article, plus automated lint gates

What you'll do in this section: give every article a "write → localise → adversarial review" pipeline, with lint attached after every pass, and use the strongest model for review. When you're done, you'll have a production line that can run dozens of articles in one go and resume after an interruption.

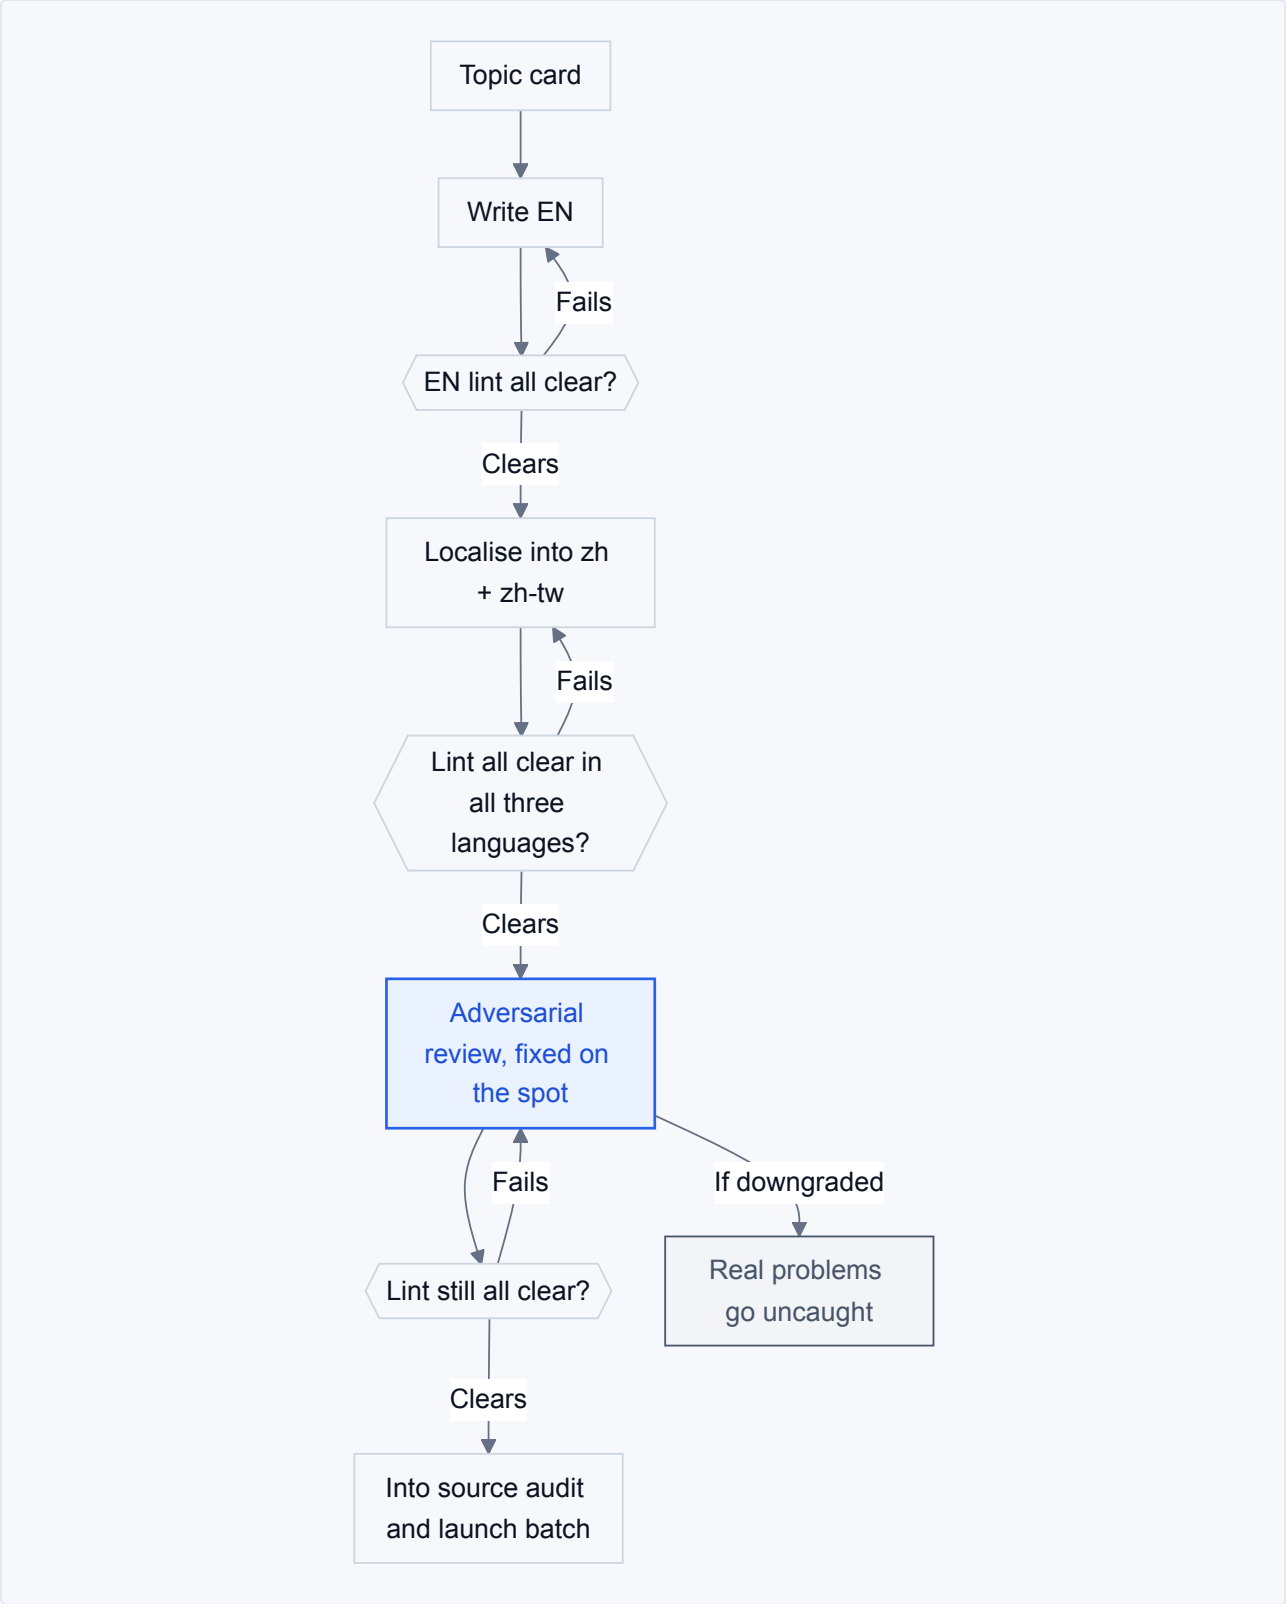


Figure: every article runs through the three passes independently, each pass followed by a lint stop; real problems are caught almost entirely by the third pass, so it must not be downgraded

What each pass reads, and what model it uses:

Pass	What it does	Model	Gate
Write EN	First read the spec, the template for the same archetype, **your own reference article for the same archetype** (one already-live article assigned per archetype) and the measurement report, then research public sources and write	Sonnet	The EN part of lint fully clears
Localise	Write zh + zh-tw, adapting rather than translating literally, sections and numbers matched one to one	Sonnet	Lint fully clears (all three languages)
Adversarial review	Check the 11 items lint cannot check, fix them on the spot	**Opus, never downgraded**	Lint still fully clears

Why writing does not read the competitor's original text. First, the pipeline needs to run in the cloud and cannot depend on a local archive of the competitor's pages; second, reading the competitor's original text lets its sentences leak into the draft. The study target's form has already been captured in the spec and lint, so writing only needs to look at your own reference article for the same archetype.

The 11 review items include: the subject and qualifiers of the first sentence, the business's own limitations, competitor wording, absolute dates, table notes, FAQ, a closing line telling readers to verify for themselves, the first sentence of each paragraph being the answer, and summary = table.

Why review cannot be downgraded. Every rule the spec cannot lint-check is pushed onto this one pass; lint can only check format, never facts or logic. Almost every real problem in the pipeline is caught by review.

> **Example** (a GEO agency) Problems review caught: a press release's first sentence attributed an award to the company itself when a different entity had actually won it, contradicting the page's own table; a summary number that did not match the table; a price with no currency stated; a number cited with no source; and the competitor described as "does not support a feature", which should have read "not found on the public page".

That last kind of mistake steps straight onto the red line on lying: not found is not the same as does not exist. Review corrects it to the standard set in → General Edition 5.31 Handing over: five steps per page, the red line on lying, and self-checks.

On a regulated side, review has one more job. First decide the side, then go through the banned-word and required-field checks for your column under 5.2–5.3, item by item. When you hit one of the three stop-and-escalate situations in → General Edition 0.4 The three labels for compliance sentences, and the stop-and-escalate rule, stop and escalate to the client's compliance officer for written sign-off; carry each compliance sentence's label over exactly as the manual gives it, and never let the agent fill in a label or

upgrade one on its own. **A written sign-off resolves only the three stop-and-escalate situations; it never turns something banned as Statute text into something publishable.**

How it runs.

- Every article runs through the three steps independently, none waiting on another: use a pipeline, no batch barriers, never make a fast article wait for a slow one.
- 45 articles per batch, 2–4 pipelines running in parallel.
- Support skipping by stage (e.g. `stages: ['zh','review']`): after an interruption, only the missing steps are run; steps that already cleared their gates are not redone.

5.D7 Six traps we hit: build the safeguards in before you start

What you'll do in this section: before you start the run, install one safeguard for each of six traps. When you're done, your pipeline configuration has six new settings: model split, writing in segments, GSC cross-check, slug freeze, source tracing and batching.

Trap	Install before you start
Running the strongest model everywhere in parallel blows through the weekly quota	Use Sonnet for writing and localising — saves about two-thirds of the calls
About 5% of Chinese localisation gets blocked by content filters	Switch to writing in segments; rewrite whatever gets blocked
A new page targets the same keyword as an already-ranking page	Cross-check against GSC before launch; a new page must not compete with an old page already ranking for that keyword
Rewriting an already-indexed old page	Keep the slug unchanged; a keyword that already has impressions must not be lost
Raw results of an agent's live measurement never saved to disk	Trace every number back to its source before launch (5.D8)
Hundreds of pages on the same template all go live at once	Launch in batches; make every page's list, market and evidence different

Figure: each of the six traps has a safeguard you can install before you start; fixing it after something goes wrong always costs more

Quota can only be saved on the first two passes. Real problems are caught by review (5.D6), so only writing and localising get downgraded; review keeps the strongest model.

Where the impression keywords stay when rewriting an old page. A keyword that already has impressions in GSC must remain in one of: title, first sentence, an H2 or the FAQ. The rest of the constraints on rewriting an old page are in 5.5.

Why saving to disk gets missed. When an agent writes a data-report piece it measures things live — crawling, calling an API — but the raw result is left only in the chat transcript. A number you publish must be reproducible, so this kind of draft must clear 5.D8.

Why every page needs to differ. With hundreds of pages on the same template, the risk in content at scale is exactly that the pages look alike. The batching cadence is in 5.D9.

5.D8 Source audit: trace every number in data reports to its source

What you'll do in this section: give every data-report piece and press release two agents (one traces every number to its source, the other re-checks a sample and is suspicious by default), then give the whole piece one of three verdicts: publish as is / fix then publish / hold back. When you're done, you'll have one more step that every such piece must clear before publication, so a number an agent measured live but never saved to disk cannot go straight out.

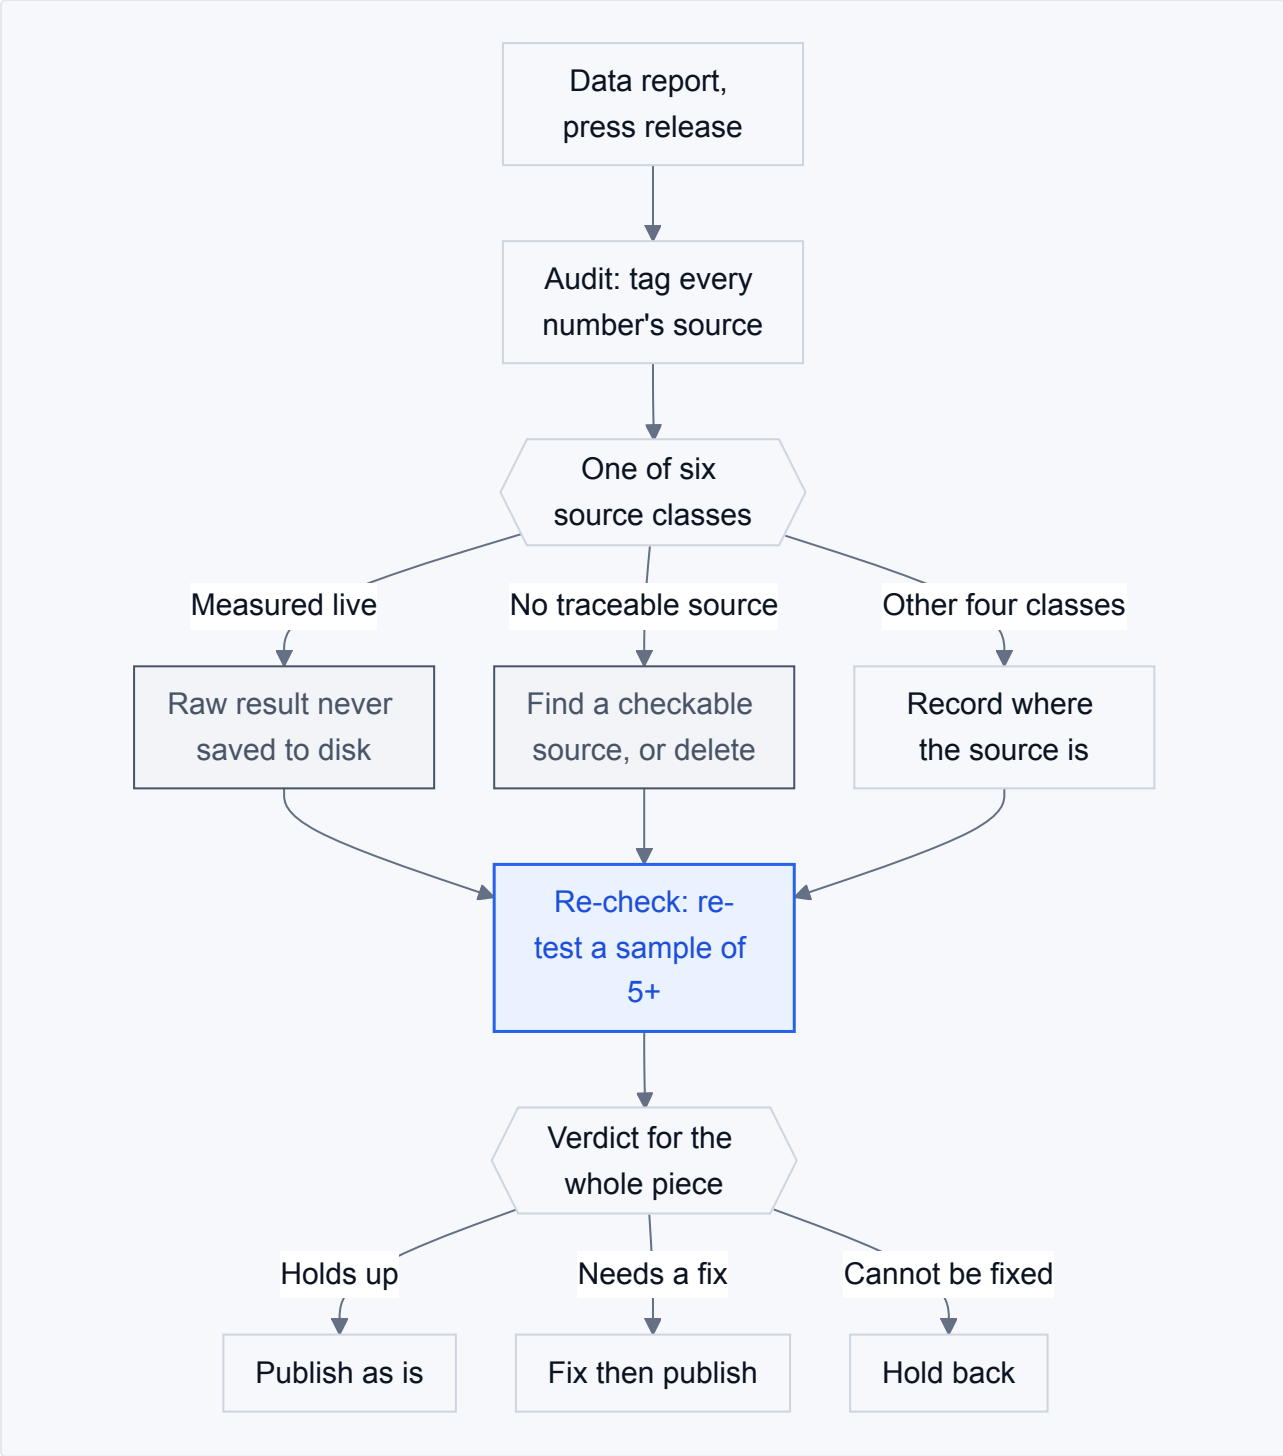


Figure: every number is first tagged with one of six source classes, then a second agent, suspicious by default, re-tests 5 or more; only after both steps are done does the whole piece get a verdict

The six source classes. From the writing record, the audit agent files every number on the page into one of: third-party citation, measured live by the agent, own evidence base, official website page, calculated, no source.

Why this is a separate step. Adversarial review can catch the odd sourceless number (5.D6), but it checks writing article by article, not every number's source; and the raw result of an agent's live

measurement is left only in the chat transcript (5.D7, trap five). Data reports and press releases have many numbers from mixed sources, so for this kind of piece the step is mandatory.

Why the re-check is suspicious by default. The re-check agent does not trust the audit agent's tags: it takes a sample of 5 or more, re-measures each one or re-opens its source, and looks for itself.

What to do with a sourceless number. Follow the red line on lying in → General Edition 5.31 Handing over: five steps per page, the red line on lying, and self-checks: if a checkable source can be found, add it; if not, delete the number. Never leave an unchecked number in, and never give one a source no one has actually opened and read.

5.D9 Launch in batches: from launch day to the keep-or-cut decision on day 45

What you'll do in this section: decide how many go out per batch and which types launch first, complete two tasks on launch day, then retest at three checkpoints — day 14, 30 and 45 — to decide keep or cut. When you're done, you'll have a timetable running from launch to verdict, so hundreds of pages are never released all at once, and never released with no one watching afterward.

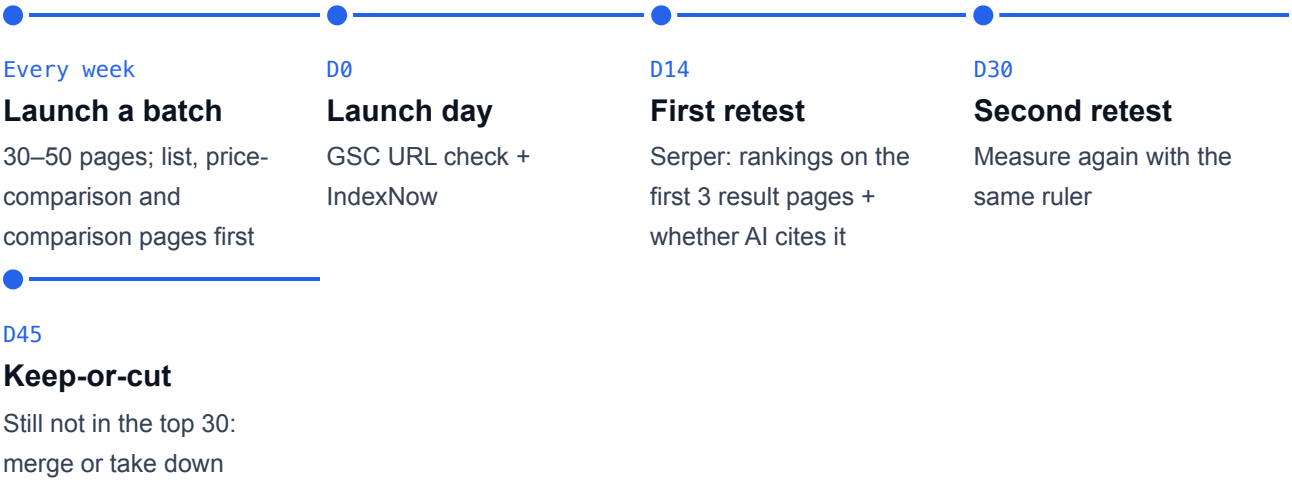


Figure: from launch to the keep-or-cut verdict, a batch of pages passes through only five action points; still outside the top 30 on day 45 means cut losses

Check your side first before deciding which types launch first. Comparison pages and price-comparison pages are the commercial page types cited most in measurement (5.D4), and list pages match the best/top combinations in the stacked matrix (5.D5). All three types have restrictions on the regulated sides: check the availability column in 5.1 for your side first, and for any type that is not open to you, put the type your side can build for the same intent into the first batch instead.

Why the two launch-day tasks cannot be skipped. They are not a retest — they only make sure the page can be indexed. Skip them, and a later retest that finds "not indexed" gets misread as "not cited".

Both retests use the same ruler. Day 14 and day 30 use the same tool and the same set of queries, and both check the first 3 pages of results and whether AI Mode / AIO cites the page; switch mid-way

and the two sets of numbers are no longer comparable (→ General Edition 7.5 Ruler discipline: change the ruler and nothing is comparable). This ruler measures a new page's organic ranking and whether AI cites it, and it is not the same ruler as the monthly one in → General Edition 7.1 What a retest produces, and the rulers — do not merge the two sets of numbers, and do not compare them to each other.

Where a merge goes. A page still not in the top 30 on day 45 is merged into the one page for its intent cluster, following the rule in → General Edition 5.5 The page unit: one intent cluster, one page; if there is nothing to merge it into, take it down.

Two protections for syncing at launch.

- An already-live article that a person has edited directly is never overwritten by the next sync: before syncing, compare against the content version recorded in the last batch, and if they do not match, keep the person's edited version.
- A link pointing to an article not yet live is removed first and added back automatically once the target goes live; every article keeps at least 2 related-reading links.

Judging results is not just about rankings. Success is judged by whether the AI's answer cites you and whether leads come in. Both of these only have data once the retest checkpoints arrive — this chapter stops at "production", as stated in 5.D1.

5.D10 Add new keywords every month: new keywords come in, a person reviews before anything goes out

What you'll do in this section: turn "find new keywords → write → launch" into a loop that runs itself once a month, but keep one human review step right before "launch". When you're done, you'll have a task that opens a PR automatically on the 1st of every month, so new keywords never pile up unattended and never go out unreviewed.

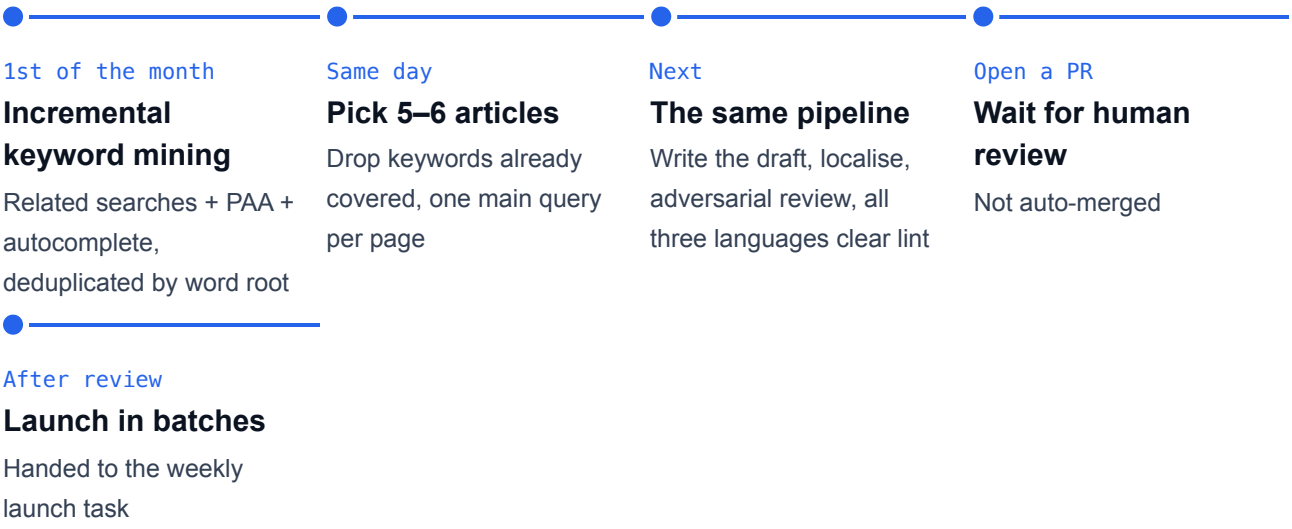


Figure: one round a month — the machine finds keywords and writes, the person only reviews; only what passes review enters the weekly launch batch

Why it is not auto-merged. Finding keywords and writing can both be handed to a machine, but "is this keyword worth doing, should this page go live" needs a person to look — get it wrong once across hundreds of pages and the mistake gets copied across the whole batch. The monthly volume is capped at 5–6 articles, a load a person can actually review.

The cost of incremental keyword mining. One round is roughly twenty-odd search-API calls: related searches, "people also ask" and autocomplete each run once, deduplicated by word root, then keywords already covered by an existing page are dropped.

Chapter 6

Beyond your website: get onto the 20 pages AI already cites

Directories, media, associations, reviews and lists, checked every month

6.1 Where the effort goes, and two numbers that must not sit side by side

6.2 The send gate, paid listings and the citation-slot table

6.3 Six dispositions, bylined contributions and outreach letters

6.4 Three source groups, and the Group B actions

6.5 Reviews: how to ask, how to reply, how to fill the empty slot

6.6 Annual rankings and off-site asset decay

6.7 Long videos narrated by the named expert

6.8 Off-site verdicts, common mistakes and rituals to cut

6.1 Where the effort goes, and two numbers that must not sit side by side

What you'll do in this section: Put this quarter's off-site person-hours on the "someone else writes about you" bucket, and remember that two similar-looking percentages each do one job — the share figure only decides whether to work a page-type category at all, the density figure only orders individual pages within a category you have already decided to work. The moment the two appear in the same table, the whole allocation table goes back to be rewritten.

Share figure · decides whether to work a category	Density figure · orders individual pages by priority
Ahrefs 43.8%: listicle citation share	deltaV: comparison page type highest at 1.87 per search
Banned: ordering page priority, setting how many letters to send	Banned: ruling out a whole category, applying across industries

Figure: the share figure and the density figure measure two different things; you cannot use one table to decide the category and order the pages

If you are on the strictly regulated side, read this first: on this side, list-type third-party pages get no listing and no payment, and if you are already listed on one, you ask the publisher to remove you — so this share figure does not apply to you. On the other sides, whether the listicle category goes into this quarter's work depends on your side: check → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. Do not copy the conclusion straight from the share figure.

The order of where effort goes follows the share breakdown in → General Edition 1.1 How buyers ask, and what AI reads: 84% of AI citations come from earned media, and paid placements and advertorials account for only 0.3%. The correct reading of 84% is "someone else writes about you", not "you buy the placement" — misread it as "84% comes from third-party sites" and you put your effort into directory listings, paid listings and list submissions, landing exactly in that 0.3% bucket. One more number to note separately: even for brand-name questions, owned content is only 23% — "the website is the main battlefield" does not hold up against the real data.

Default order	Industries with heavy platform control
① what you can edit yourself ② what someone else writes ③ what you pay for	① first fix the wording on platform profiles that need no one's approval ② then do the ones that need approval
Basis: earned 84%, paid 0.3%	Basis: back-office edits take effect immediately, no outreach cycle
Total workload unchanged, only the target changes	The container is not yours, but the words can be

Figure: the default order ranks by who decides this slot; for industries with heavy platform control, first fix what needs no one's approval

> **Example** (e-commerce) Store and product profiles on a platform whose door is open are places where back-office edits take effect immediately, so get the wording right there first. Judge the platform's door one retrieval leg at a time, by running curl against its robots.txt; if the platform blocks a leg, the platform counts as not reachable on that leg.

Getting yourself onto a list AI already cites, and ranking #1 on it, is currently the only lever that comes with a figure for the size of its effect; that number and its qualifier appear only once in the whole book, in → General Edition D.1 Number discipline: how to label numbers, and what stays internal hard rule 3 — use it to order the work, never to promise any result.

Within off-site effort, one group alone takes 60%: institutions, government, associations and manufacturers. One more fixed calendar rule: **when each page launches, log an off-site target — which source should cite this page — and within 14 days that must turn into either a sent letter or a "not reachable + basis" entry**; anything not converted in time gets listed item by item in the monthly review. Write a page without pushing it off-site, and its output just gets pushed back into the "owned content is only 23%" bucket. How to do this group's work in four steps: see → General Edition 6.4 Three source groups, and the Group B actions.

Putting facts into a container that is already cited shows results in days to weeks; building a new page takes 2–8 weeks (inferred, 65% confidence, not measured) — which is why "already present" dispositions are ranked ahead of self-built pages. Review counts and star ratings do get copied into the body of AI answers, but only for recommendation-type questions such as "who's best" — see → General Edition 6.5 Reviews: how to ask, how to reply, how to fill the empty slot for how. Reddit's share of ChatGPT answers has dropped sharply before; the last drop of the same size fully recovered later, and ChatGPT has kept reading Reddit the whole time — what stopped was putting it into the answer. So we do not fix a conclusion here; read your own share every month, rule in 6.4. For the numbers, samples and sources above, see → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests.

6.2 The send gate, paid listings and the citation-slot table

What you'll do in this section: Nail three things — the publishing authorisation, the materials sign-off sheet, and a zero-hit banned-word check — into a single gate; send no outreach letter if even one is missing. For paid slots, first ask how billing works and only take a fixed rate card. Then log every citation slot into the *Citation-slot table*, and on the monthly review tick off five binary features in passing; once a structure clears the threshold, write it into next month's page spec.

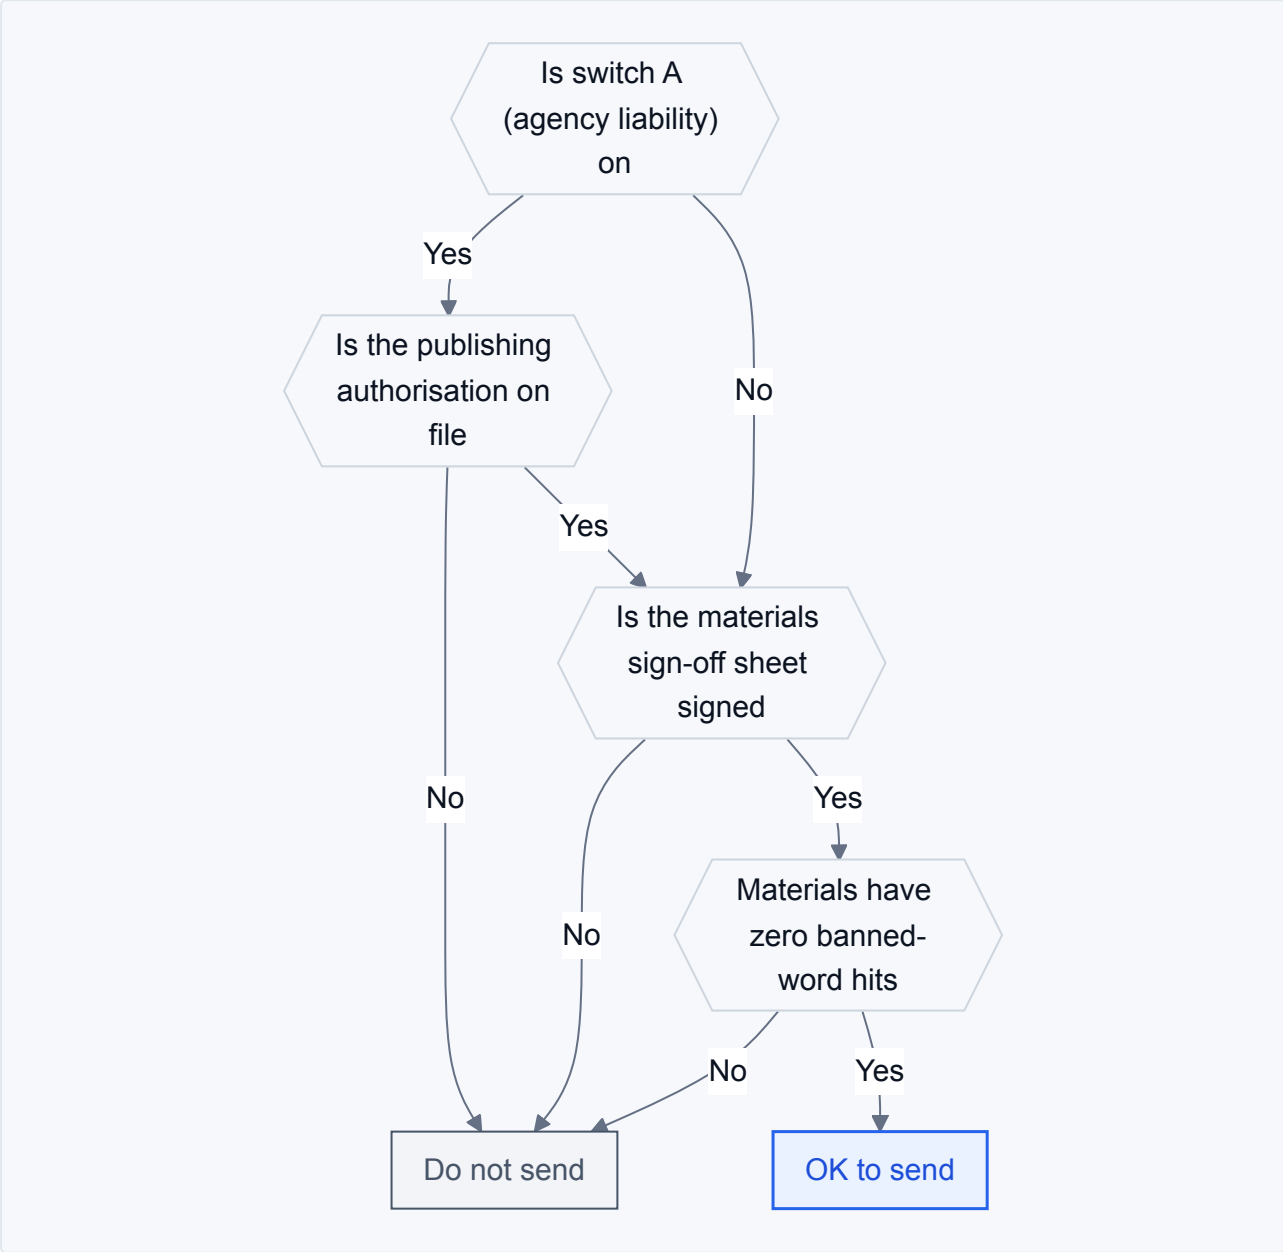


Figure: every outreach letter passes this gate; miss one item and it does not go out

Two documents; if either is missing, nothing goes out. The **written publishing authorisation** is only required on the side where switch A (agency liability) is on: the licence holder signs it, stating the licence number, the scope of authorisation, the term, and how it can be revoked — without it, publishing under an agent's name means the publishing party bears the liability directly; see the rule and the penalty in your industry edition. The **materials sign-off sheet** is needed on every side, one per batch: the

signatory confirms, line by line, that the facts sent to the third party in this batch (price, registration number, an 80-word description, photo permissions) are true and authorised for publication — without it, if the third party publishes them as given and something goes wrong, the liability lands on you. The compliance memo must open with the liability and penalties page. Check the banned-word list for your side in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers and → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions. → General Edition 6.3 Six dispositions, bylined contributions and outreach letters uses this same gate: every letter there must clear it before it is sent.

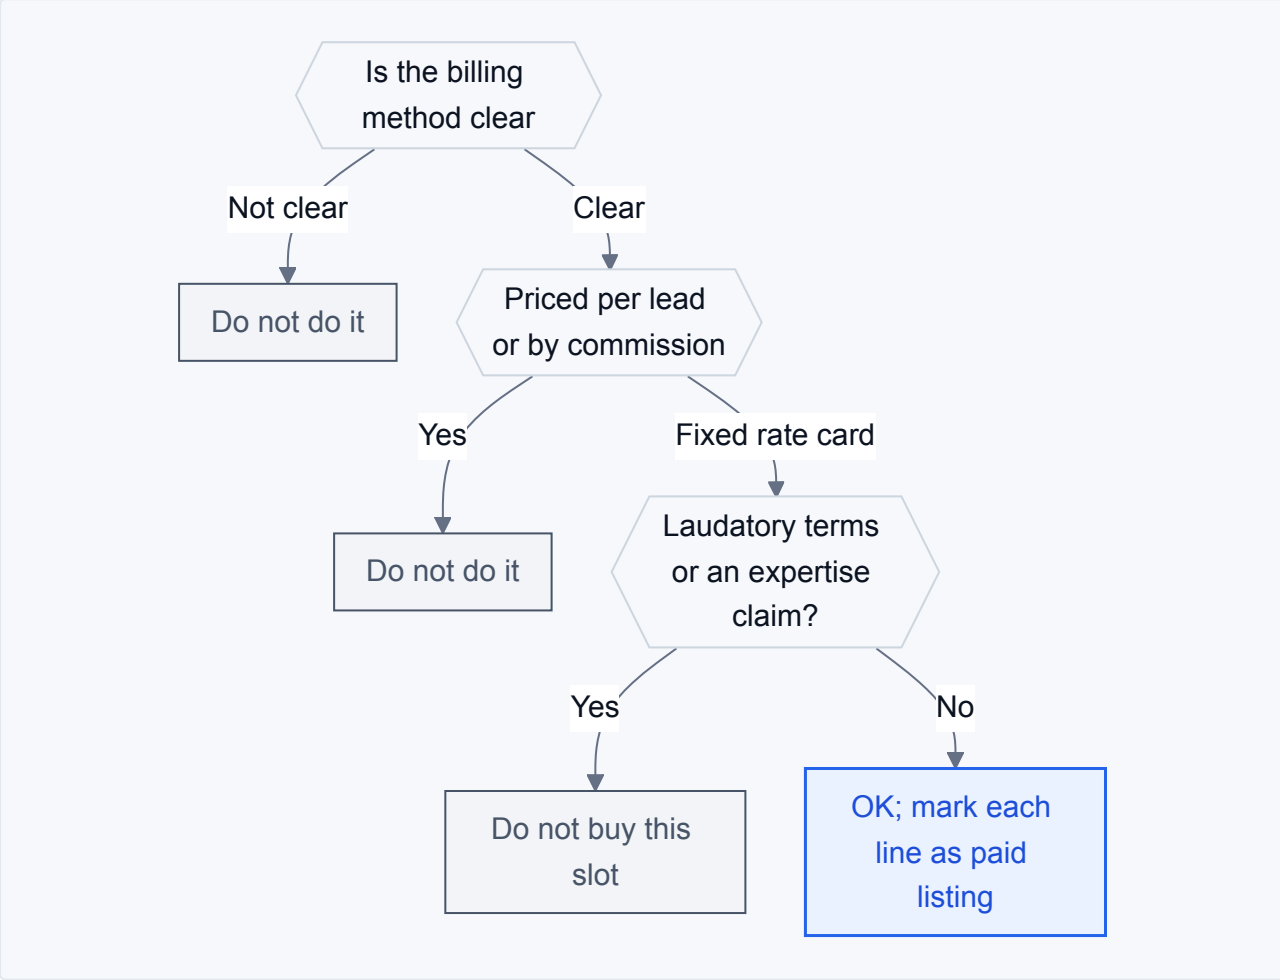


Figure: for a paid slot, first ask how billing works; do not proceed if it is unclear; never buy per-lead billing or a page carrying laudatory terms

Strict side: a list carrying laudatory terms (words like best, top choice, leading or number one in the title, a section name or an entry description) is never done, neither as a paid placement nor as an authorised publication. Disclosing the sponsorship does not make it lawful either; if you are already listed on a page like this, ask the publisher to remove you, and do not send a correction letter. All you can still buy is a pure directory or a price-comparison page whose title and entries carry no laudatory terms, and these two must not be budgeted side by side as equals: a price-comparison page is Group A, a pure directory is a Group C logging action (see 6.4). The lightly regulated side may buy fixed-rate-card listings and ad placements; if the paid page's wording includes comparisons or a claim of expertise, do

not do it there either. If switch F (referral commissions banned) is on, pricing per lead, per sale or by commission counts as fee-sharing for clients, so buy only fixed rate cards.

Every paid item must be marked "paid listing" line by line, with its cost listed separately, in every task and every monthly report, and **must never be written as "recommended" or "included in an independent review"**.

The *Citation-slot table* is off-site's main ledger: its fields are fixed, and none may be added or removed; one row = one URL (not one domain); the 16 fields and how to fill them are in → General Edition B.4 Headers for work orders, ledgers and monthly reports. The input must be the raw `domains_cited` you get from probing — the aggregated version strips out the target sites and cannot be used as input. Do not leave the "competitor's way in" column blank: a competitor has already tested, for you, which action works on this shelf — leave the column empty and all you know is which page to get onto, not how to get onto it.

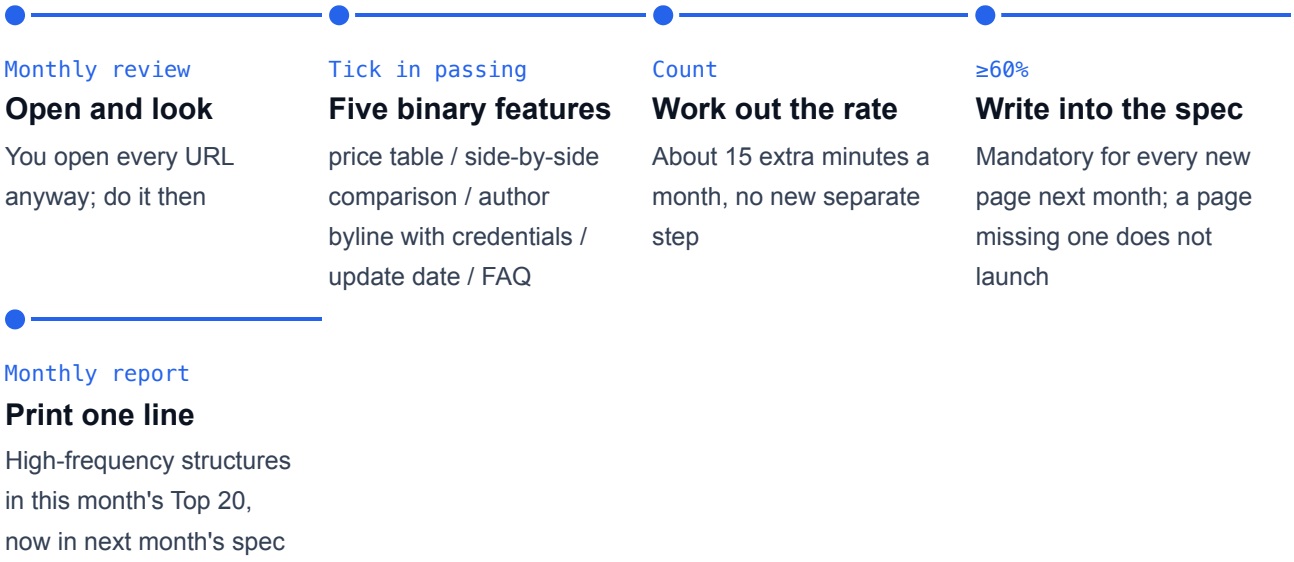


Figure: the citation-slot table's second use — reverse-engineer the structure AI already favours in your own question pool

Look only at visible HTML — not source code, not schema — and score anything you cannot tell as 0, never guess. This sample is what AI actually picked from your own question pool, so it fits better than averages borrowed from someone else's market. It cannot be used to override the density gate or the red line on lying: a structure that exists on a competitor's page is not the same as content you are allowed to write — the compliance verdict always comes first. The five-digit string for each row is for your own eyes only; outward-facing materials show only the summarised high-frequency feature names.

6.3 Six dispositions, bylined contributions and outreach letters

What you'll do in this section: Once past the send gate in 6.2, sort every URL in the baseline Top 20 into one of six dispositions, by who decides whether you can get into that slot. Bylined contributions give only checkable facts, under the contributor's own name. One outreach letter fits all cases; send it

only to the two classes that should get it; follow up once at +7 and once at +21, then stop. Set no target for the number of letters.

Before sending, clear the send gate in → General Edition 6.2 The send gate, paid listings and the citation-slot table first (authorisation on file / this batch's sign-off sheet signed / zero hits on the banned-word list).

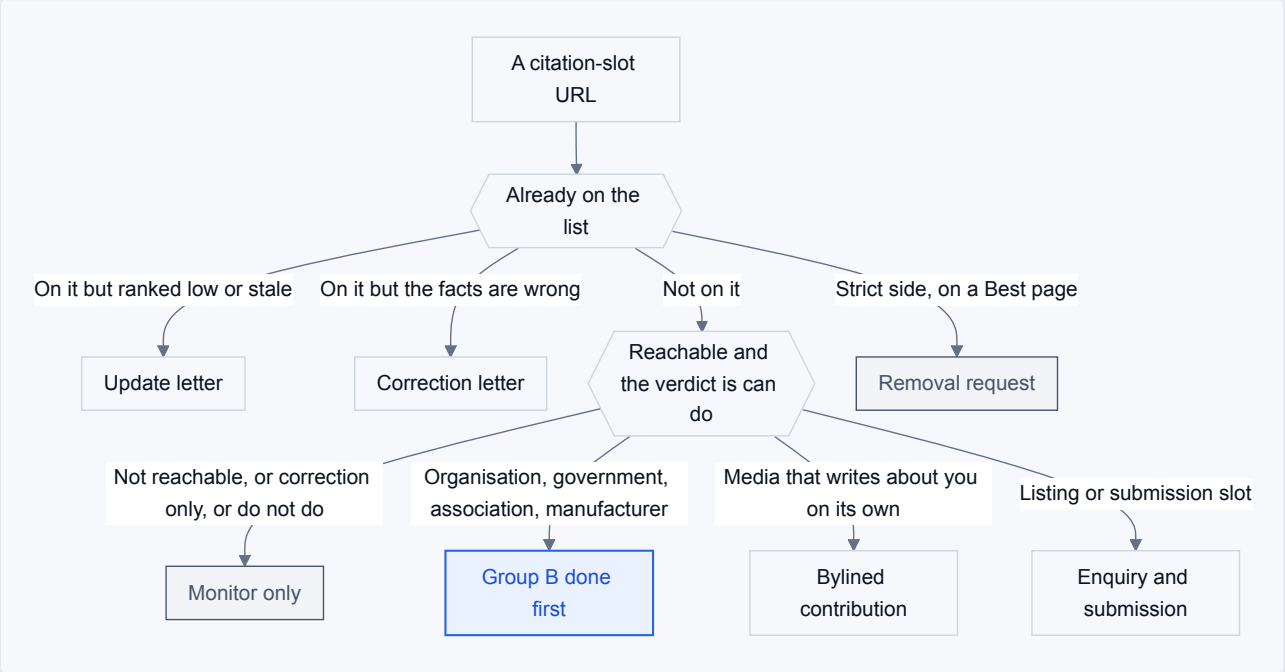


Figure: the six dispositions ranked by who decides this slot — what you can edit yourself first, what you pay for last

The order runs on "who decides this slot", not "which costs least effort": order by effort and you mix paid listings, directory submissions and bylined contributions into the same bucket, when earned media is 84% and paid placements and advertorials are only 0.3%. Total workload is unchanged, only the target changes. A correction-only item still gets a correction letter, and a request-removal item gets a removal request and a check on the due date; then both join the not-reachable items in the monthly review list.

Disposition	Time per item	Expected cycle
Update / correction letter	20 minutes	1–2 weeks (inferred, 60% confidence, not measured)
Group B institutional	15–30 minutes each	Claiming: immediate–3 days; registers and associations: 1–2 weeks
Bylined contribution	20–30 minutes per letter	2–6 weeks (inferred, not measured)
Enquiry and submission	Finding the contact and rate card: 2.5–4 h + writing each swap-in paragraph: 2.5–4 h, sent in two batches 1 day apart	Free editorial slot: 2–4 weeks; paid slot: 4–8 weeks (all inferred, not measured)

Disposition	Time per item	Expected cycle
Monitor only	0–20 minutes	—

Within enquiry and submission, order individual pages by page-type citation density: send comparison-type and price-guide pages first, the rest after — this density figure is only for ordering within this class, and must never be used to rule out a whole category. A batch outreach reply rate of 15–25%, and a 2–6 week lag from launch to a change in AI citations, are both inferred (55% confidence, not measured) and used only to schedule internal work hours.

① **Pick the target**
Local media → association newsletters → government registers → professional journals, in that order

② **Find the contact**
Byline → contact page → public email; write "could not be found" if you cannot

③ **Give the facts**
Exclusive checkable facts: a fixed price table, conditions, timeframes, model numbers

④ **Write the letter**
Reference their most recent piece; attach a source link to every fact

⑤ **Sign with a real name**
Own-name byline linking to the person page; never the organisation as author

⑥ **Follow up**
At +7 and +21, then switch to "reachable (no response)"

⑦ **Acceptance**
Open the live URL and verify it before entering it into the ledger

Figure: how to make a bylined contribution — the contributor's own name, checkable facts only

One letter takes 20–30 minutes. The facts you give out must match the official website line for line, and must pass your side's gate checklist before sending; **never pay in exchange for publication** — paid placements go through enquiry and submission instead, and the two must never be mixed into one letter. 84% is the global prior across industries; your own industry's measured data overrides it — order the targets by your own shelf's page-type mix.

> **Example** (B2B software) In a snapshot of this industry's category questions, news-type earned media comes in at about 1%, so it drops to the bottom of the order (the sample is too small to support "do not do it", only enough to support "do it last"); free, self-publishable developer-community long-form contributions are measured to get seats, so they go first.

Main letter	+7	+21	After that
One letter fits all	**Follow up once**	**Last follow-up**	**Send no more**
The middle section swaps by disposition: enquiry / update / correction / contribution	One question only: editorial decision, paid (billed how) or not open?	No reply counts as not open; move to the monthly review	Send again only in the next listicle rewrite season

Figure: the rhythm of an outreach letter; nothing more is sent after +21

Send only to two classes: **remaining reachable Group B items**, and **Group A items with all three of a specific URL, a usable contact and a written billing basis**; never send this letter to a pure directory. The full letter text is in → General Edition B.3 Outreach emails and review replies. Externally, say only "every reachable one gets sent, each one logged" — never give an expected reply count, and never dress up a small number of letters as precise targeting; how many letters to send is an internal scheduling decision, and the approved wording is in → General Edition B.2 Approved wording and scripts.

Cannot give		Give instead
Result numbers such as success rate or win rate	→	Process duration, cost breakdown, billing method
Result numbers with no stated method or basis	→	Numbers with a denominator, a year and a stated method
Before-and-after photos, a client's identity	→	Spec tables, model numbers, test conditions
Claims like "guaranteed" or "fall behind if you don't"	→	Statements of fact with the conditions spelled out
Negative claims about competitors	→	Only your own exclusive checkable facts
A fabricated original or reference price	→	The real selling price, with the comparison date noted
Adjectives, recommendations: best, top choice	→	Factual description, no adjectives
An export of the citation-slot table showing competitors' rankings	→	For your own eyes only, never sent out

Figure: what cannot go into outreach material, with a checkable substitute for every one

6.4 Three source groups, and the Group B actions

What you'll do in this section: Clear the compliance verdict first, then allocate person-hours by Group B (government, associations, manufacturers — do first) → Group A (opportunistic slots) → Group C (pure directories, log only). The four Group B steps run in a fixed order; do not touch Group A submissions until they are done. The day a price page launches, send the same price table straight to the price-comparison slots.

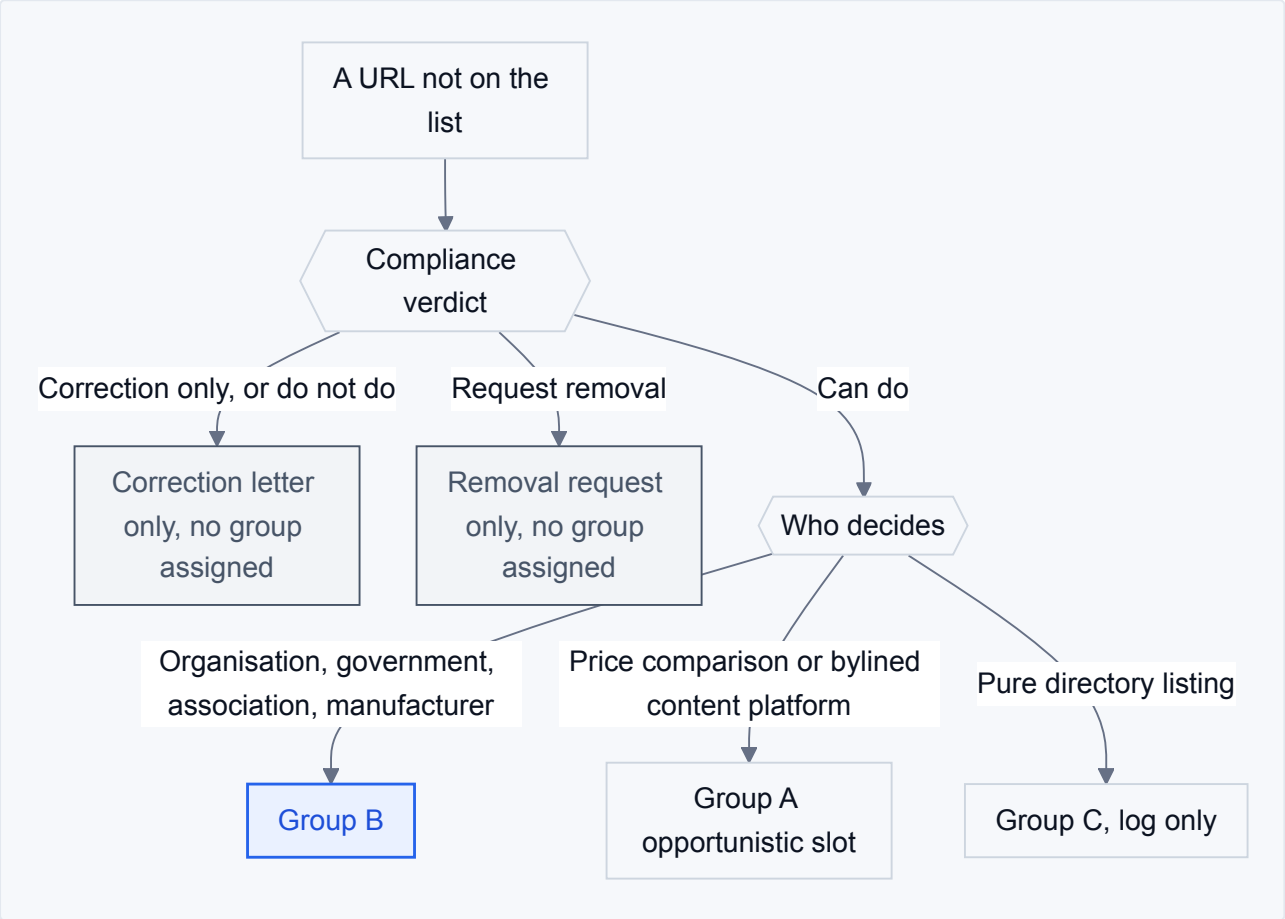


Figure: the compliance verdict comes before the group; an item verdicted correction-only, request-removal or do-not-do enters no group

The three groups decide whether person-hours go in; the compliance verdict decides whether you can touch it at all.

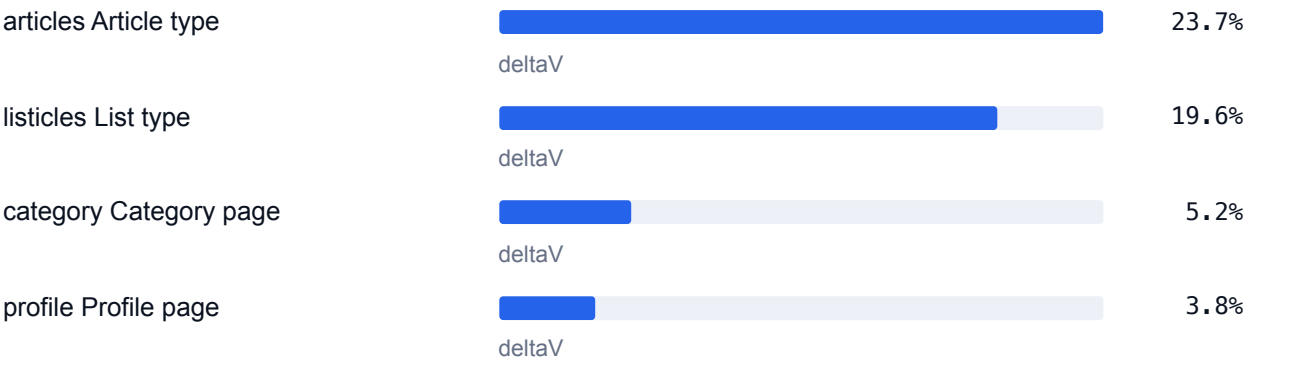


Figure: in deltaV's page-type share, pure directories and profile pages rank lowest, so person-hours are allocated by group, not by domain size

Group B (government and statutory registers, professional association profiles, official links from manufacturers) is mostly free, stays once done, does not depend on a third-party editor's nod, and turns effort into on-the-list pages with almost no loss. Do not treat it as "cheap, so do it in passing"; the reason for the effort is the opposite: in the ChatGPT source reshuffle of 2026-08, institutional and government sources rose from about 1/6 to close to 1/3 — this is the group whose share is rising, so it is scheduled ahead of Group A; every free item must be done, none skipped. Group C is done only when it is free or very cheap and can be finished in passing; log it in the citation-slot table once done; it takes no person-hour budget, gets no enquiry letter and no separate follow-up.

Group A (price-comparison and price-guide pages, bylined professional content platforms) is kept to opportunistic slots only: send one letter along with the Group B batch only when all three of a specific URL, a usable contact and a written billing basis are present and the verdict is can do; anything missing one is logged to the reserve list — not chased, not followed up, not reordered — and reviewed once a quarter in passing. The number of Group A items is never a threshold for any verdict, and a thin Group A is never grounds to downgrade this quarter's off-site work. On the strict side, Group A is a structural gap: most local third-party comparison slots carry laudatory terms in their titles, so you cannot submit to them, and where one lists you, all you can do is request removal — do not wait on it; move the effort to a self-built, all-products price guide page instead (comparison-page-type density is counted by page type, not by domain ownership, so a self-built page earns it the same way). Do not submit to, or pay for, a laudatory-titled page to fill out Group A either — marking it "paid" does not make it compliant.

> **Example** (dental and aesthetics) Almost every real local third-party price-comparison page carries Best / Top in its title, so none of them gets a submission, and any that lists the business is verdicted request removal. A second search turned up the same batch of titles, so letter counts and schedules are no longer set on the premise of "filling out the Group A list".

① **Claim profiles**
The five business profiles, every fact field matched word for word

② **Check registers**
Government and statutory registers, check only, do not submit; fix any mismatch the same day

③ **Complete association profiles**
Add credential fields to professional association profiles, link back to the person page

④ **Manufacturer listing**
Apply for a locator with proof of purchase; it links to you (inbound)

Figure: the four Group B steps, in a fixed order; no free item may be skipped

Do not touch Group A submissions until every free Group B item is done; leaving even one free item blank means you may not mark this quarter's off-site work "complete". Step 4 applies only to a brand actually in use; one application pays off in two places at once: a link from a third-party site, plus a third-party endorsement that you use the genuine brand and model. It is an inbound link; on the strict side, your own pages state only the model number and the registered facts, and do not link out to the manufacturer. How to fill in the free-text field, the business post and the URL field for the five profiles: see → General Edition 3.4 Third-party credential tiers, Wikidata and the five business profiles.

The day a price page launches, send the public version of the same table as an attachment straight to price-comparison articles, a directory's price column, and price guides — they all want the same thing, one checkable price table, at close to zero cost. Red line: any price table sent out must match the official price page line for line, and the price wording must follow your side's price rules (→ General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers); send an update letter the same day a price changes.

Build your own list of sources: write down a price only when a page actually states one, mark unstated prices "quote on request", and mark anything you could not find "unverified"; the compliance-verdict column always outranks the priority column — an item verdicted correction-only, request-removal or do-not-do never gets a submission letter, however high its priority. One set is shared by every side: the five business profiles (Google / Bing / Foursquare / Apple / Yelp) come first, with Bing at P1 and the rest at P0; no fixed conclusion for Reddit — read your own share of `domains_cited` every month, and **only schedule it as a task above 2%**; at or below 2%, log it but do not commit effort. See your industry edition for the specific sites and contacts.

This quarter, off-site is judged on only two things, both fully within your own control: ① has everything due to Group B been sent? ② has the all-products price guide page (English + Chinese) gone live and been entered in the page register?

6.5 Reviews: how to ask, how to reply, how to fill the empty slot

What you'll do in this section: First decide, by your side, whether you may ask for reviews at all — the strict side never may, the other two sides may. On sides that may ask, draw from the whole body of paying customers under a fixed set of hard boundaries, keeping the incentive decoupled from the

content. On every side, reply within 72 hours and never confirm a reviewer's identity. On the strict side, fill the empty review slot with the highest-strength substitutes first.

Whether each side may ask for reviews, and may display them on-site, is set by your side — check → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions; this section covers the actions only. On every side, a public reply never confirms that the reviewer is your customer and never restates the service they received — this is a privacy and confidentiality obligation (Conservative line (not statute text)); a customer writing it themselves in a review does not give you authorisation to confirm it publicly.

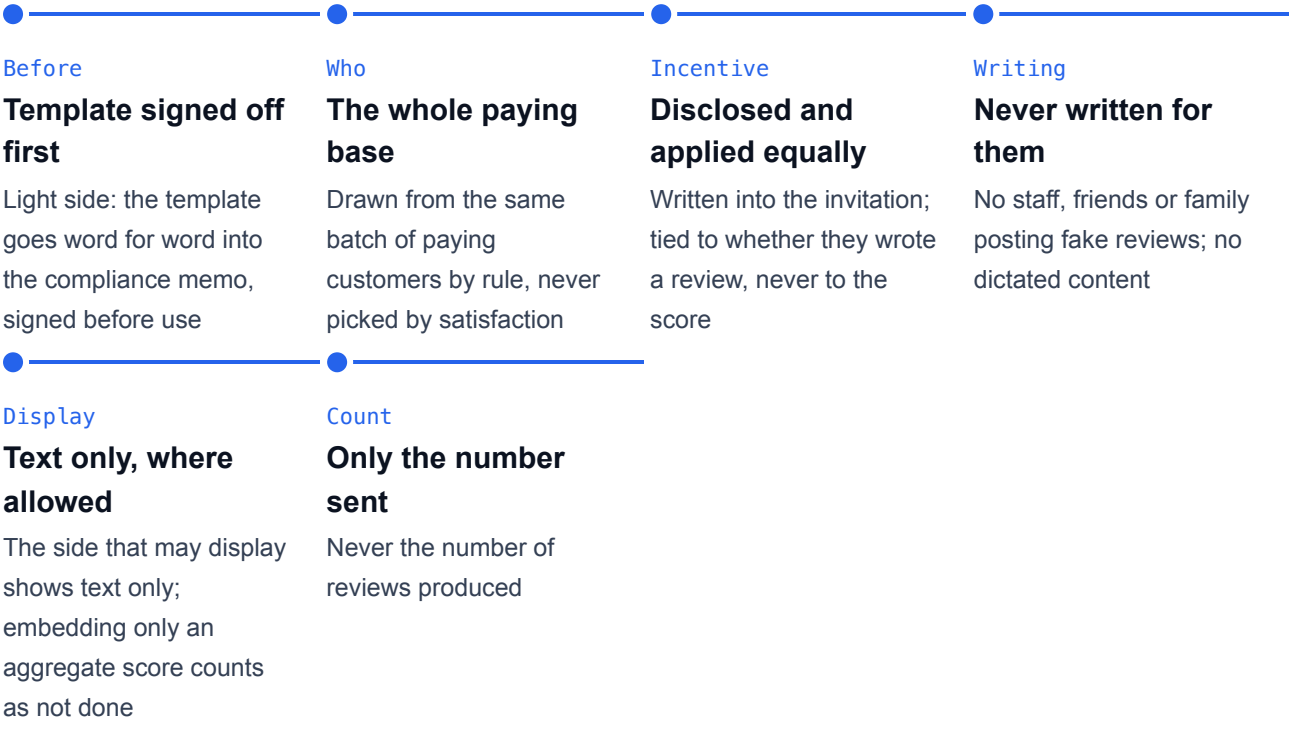


Figure: on the side that may request reviews, send the request like this; count only the number sent

On the lightly regulated side, until the sign-off is back, all you may do is "ask them to leave a review", with no steering on what to write. Whether asking for a review on this side falls under the restrictions on professional advertising is uncertain: our confidence in the original text is under 60%, and it is not measured. Do not decide this on your own. The strict side may not ask even with a sign-off. A written sign-off resolves only the three stop-and-escalate situations; it never turns something banned as Statute text into something publishable. See → General Edition 0.4 The three labels for compliance sentences, and the stop-and-escalate rule.

> **Example** (e-commerce) Asking for reviews is allowed and must be done: send the invitation to the full set of that batch's paying buyers, not "ask if they're happy" first and decide from that; state any incentive in the invitation, and give it the same regardless of the score.

There are two layers of gatekeeper: the platform's seller policy (which can suspend your store outright) and how fair-trading enforcement treats false and misleading reviews. Follow whichever is stricter; platform policy is usually more detailed than the law, so follow the platform's first. Who the gatekeeper actually is — the law, a professional code, a platform rule or platform terms — has to be measured industry by industry, never inferred from the side.

Wording	Red line
Route to the right role, leave a contact	Never combine a project name, a branch name and personal details
Reply within 72 hours, negative reviews included	Never state a price, an outcome, a success rate or a peer comparison
Use the same template for every negative review	Never argue, never restate what happened

Figure: replying to reviews is the same on every side — never confirm a reviewer's identity, never restate what they bought

① **Price and comparison pages**
Tiered fixed-price pages + third-party price-comparison pages, highest density

② **Named long-form content**
Institution-bylined professional educational long-form; not settled yet, not scheduled for now

③ **Profile credentials**
The five profiles + /facts + checkable third-party credential tiers

④ **Links from institutions**
Group B, free, must do: government, associations, manufacturers

Figure: on the strict side, the four substitutes that fill the empty review slot, in order of strength

Pure directory listings do not count in these four — measured data shows it is the lowest-citation-density page type, and belongs to Group C logging actions. For the third item, state only the credential tier's name and issuer; the strict side never writes a manufacturer-granted tier (Conservative line (not statute text)); never write "highest", "only" or "leading" — even if true, it is always banned. "AI copies the credential tier, not the review" rests on only a single-engine general-search snapshot, and until locally retested it is used only for internal ordering and direction, never in outward-facing materials — evidence in → General Edition A.1 Evidence for mechanics, the door and identity. Reply templates in Chinese and English are in → General Edition B.3 Outreach emails and review replies; the monthly three review numbers are in → General Edition B.4 Headers for work orders, ledgers and monthly reports. The strict side's monthly report always adds one fixed line in its review section: we do not request reviews. The figures above are a passive record of reviews that arose naturally and are not used as a work target or acceptance measure.

6.6 Annual rankings and off-site asset decay

What you'll do in this section: For an annual ranking, do only three things — build a calendar and ask about the next edition's window, send only the participation link and never fill anything in for others, and write a factual statement once the result is out. Write the monthly review of off-site assets into the deliverables — they decay, and left unmanaged they quietly leak away. The ledger tracks only three numbers.

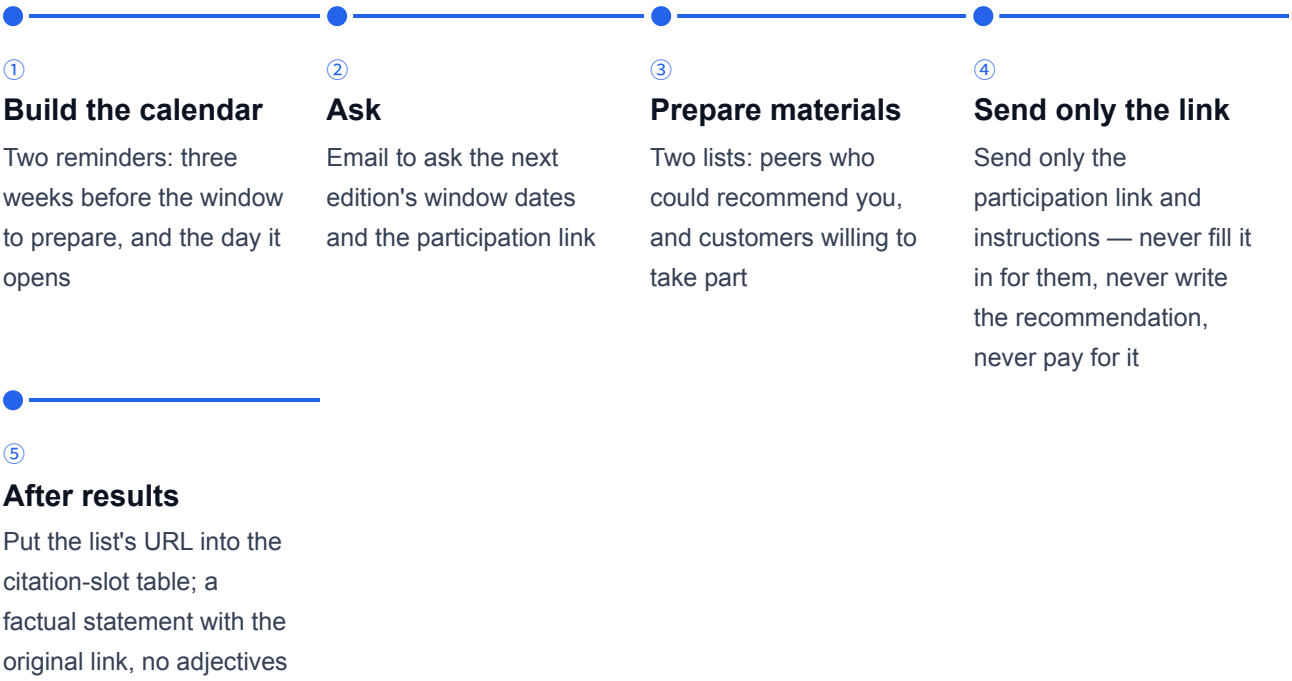


Figure: for an annual ranking, do only these steps; being selected is never treated as something you control

> **Example** (legal) An annual best-of ranking co-run by a newspaper and a statistics body: the survey window has clear start and end dates, self-nomination is explicitly banned, and the page states no fee anywhere (zero-fee entry, 85% confidence). It is worth a standalone entry because it is free, annual and from an authoritative domain, and there are third-party list-type pages that openly use it as a ranking input — but who gets in is decided by the organiser and the voters; all you can guarantee is preparing, taking part and keeping a record, on schedule.

For peer-voted or panel-judged rankings: the annual cycle runs 3–9 months, entry is free (unverified), self-nomination is banned, and the calendar reminder and preparation process are the same as above; see your industry edition for the specific ranking names. Once selected, citing the placement in your own advertising should still only be a factual statement with the original ranking link — no adjectives, never write "the best"; on the strict side, your own pages do not quote the placement (see → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions), and whether the website mentions it goes to the compliance officer first.

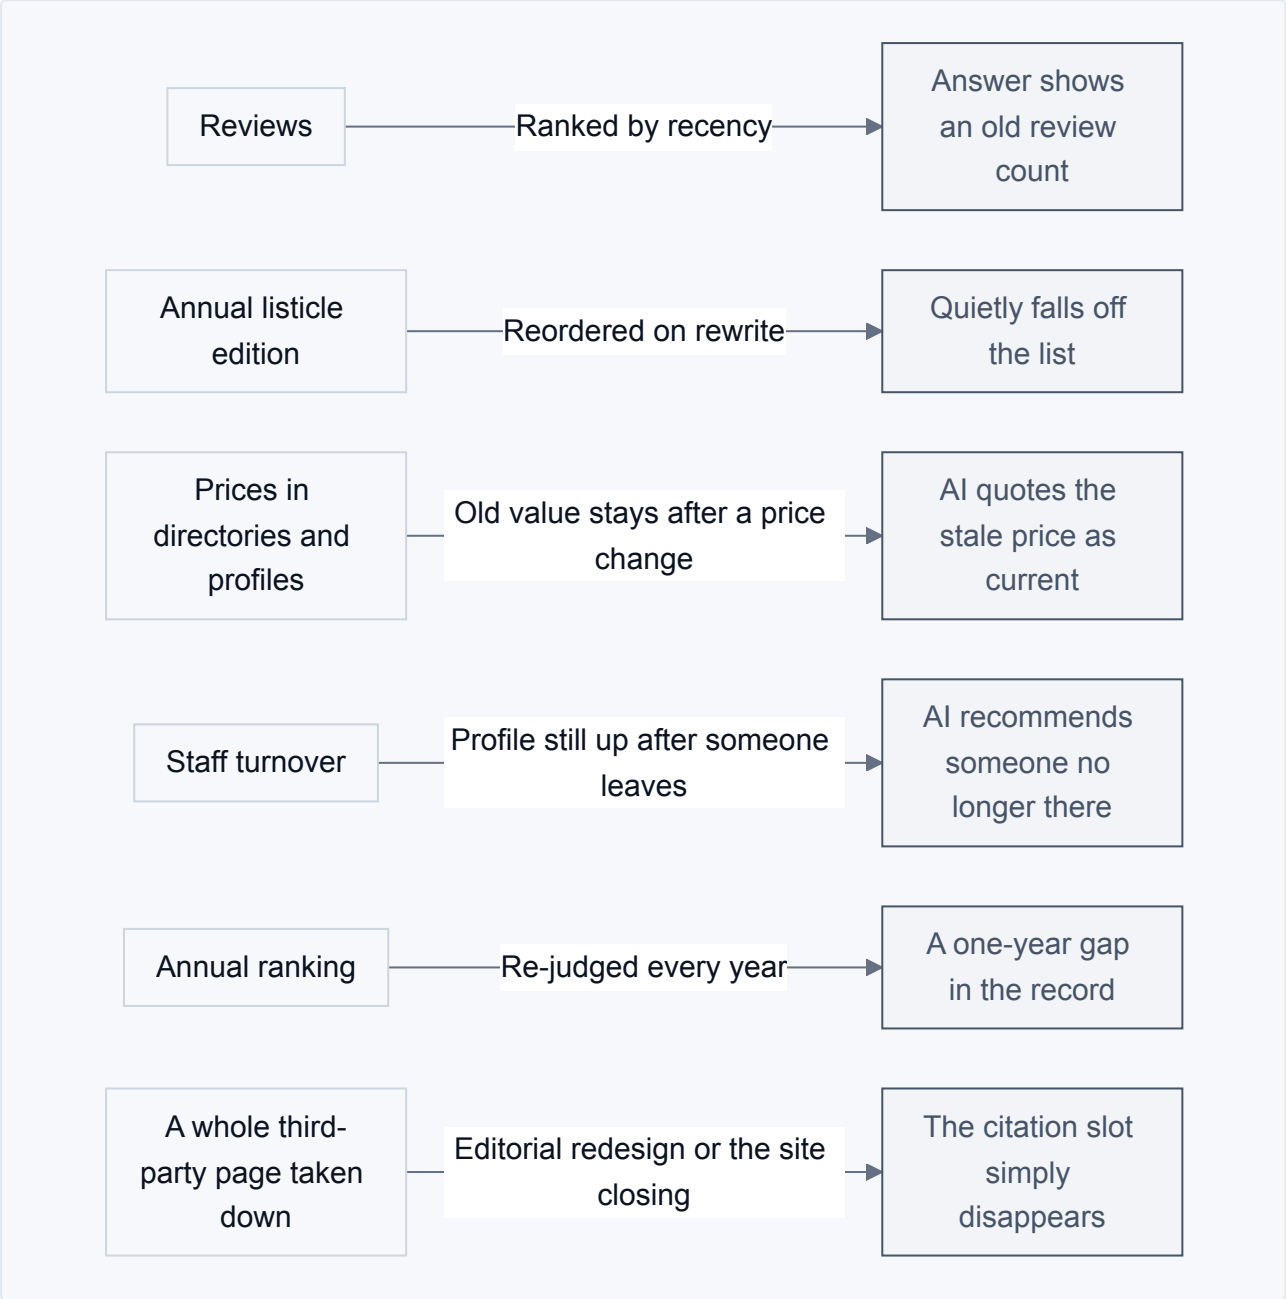


Figure: the six kinds of off-site asset decay, checked one by one in the monthly review

Review frequency: check listicles monthly from Q4 each year; check annual rankings 3 weeks before the annual window; check the other four monthly; tick each one "present / gone / page deleted". Leaving the price item unmanaged results in AI quoting a wrong price, and complaints.

D1 afternoon
Produce the citation-slot table
Raw domains_cited, three states plus a verdict, URL by URL

Week 1
Updates, corrections and removals
Send every letter for already-present, correction-only and request-removal items

Weeks 1–2
Group B first
Profiles → registers → associations → manufacturer locators

Week 2
Annual rankings
Send the enquiry email, build the calendar reminder

Weeks 2–3
Contributions and submissions
Bylined contributions first, Group A enquiries follow, 1 day apart

Weeks 3–5
Follow-up and close-out
+7 for the whole batch in weeks 3–4, +21 for the whole batch in week 5

From week 4
Monthly review
Tick three states, produce three numbers, tick the five binary features in passing

Figure: the work order for off-site's first two months

The first two months total about 19–22 hours, then about 1.5 hours a month after that. No new tool is needed; the only machine input is the raw `domains_cited` — never the aggregated version, which strips out generic domains, and business-profile and review-site domains are exactly this step's target sites.

The ledger tracks only three numbers (none depending on AI's own fluctuation):

1. **Pages on the list X/20**: the denominator is the frozen baseline Top 20, opened and screenshotted by hand, URL by URL;
2. **New listings**, tracked separately, not counted in number 1: third-party URLs newly on the list this quarter that were not in the baseline Top 20, each one marked whether it has been cited in an AI answer this quarter — one not yet cited is only a record of work done, never a result number;
3. **This month's new and dropped citation slots**, URL by URL.

Three deliverables — not produced by month-end means not done (headers in → General Edition B.4 Headers for work orders, ledgers and monthly reports):

Deliverable	Content	Frequency
Citation-slot table	Baseline Top 20, every row filled, "could not find" where you could not find it	First version produced from the quick-start coarse screen on D1, marked "quick-start basis, to be reproduced on D6"; reproduced to overwrite once the D6 baseline is frozen; monthly after that

Deliverable	Content	Frequency
Outreach ledger	Three states — sent / replied / live — with the email date and the live URL	Weekly
Off-site asset ledger: three numbers	The three numbers above	Monthly, and into the monthly report

A one-page *Maintenance notes*, kept always at hand: which items need re-submitting every year, which rankings get redesigned, that a price change in the profiles must be synced across all five, and that the strict side does not proactively request reviews. Paid-listing discipline is an integrity red line: for any listing obtained through a paid slot, mark it "paid listing" line by line, with the cost listed separately, in the plan, the task list and the monthly report; never pay for any specific ranking, never buy reviews, never organise fake recommendations, never pose as an independent review, and never use a platform billed per enquiry or by commission share. **Disclosing a payment fixes an integrity problem, not a regulatory one**: an item verdicted correction-only, request-removal or do-not-do must not go in, even when it is marked paid. How these two disciplines are worded externally: see → General Edition B.2 Approved wording and scripts.

Off-site assets decay: that is the only answer to "why keep doing this" you can give without lying.

6.7 Long videos narrated by the named expert

What you'll do in this section: Within 14 days of every thick page going live, pair it with an 8–15 minute video narrated by the named expert — the script is simply reading out that page's H2 questions in order, one by one; the description holds the page URL, the full transcript and a link to the person page, with timestamps matching the H2s. Do not chase view counts.



Figure: correlation with AI Overviews visibility; for video mentions it is more than three times that for backlinks

This leg needs no viral hit: the correlation between view count, likes, subscriber count and citation frequency is about −0.03 (a different study, measuring citation frequency, not the same measure as the chart above), so buying views, chasing plays and fancy editing return nothing. It is scheduled into off-site, and not last, because it is the strongest single correlation among all measured factors, and YouTube accounts for 23% of citations in Google's AI answers: on the Google leg, nearly a quarter of citation slots are videos. The benefit lands mainly on Google and Perplexity; do not use ChatGPT's seat

count to check it. **Note** 0.737 is a correlation coefficient, not causation; the working hypothesis is that a YouTube page is already inside the cited set, so whatever text sits on it gets a chance to be copied — the causal mechanism is not measured. Sources in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests.

Title	Main question + location + year, in the buyer's own words that match the page's H1
On camera	The thick page's named expert, in person, narrating throughout
Length	8–15 minutes; under 8 minutes does not count
Description · line 1	The full URL of that thick page
Description · body	The complete transcript, in full
Description · end	The named expert's person-page URL
Timestamps	One per H2 of that page, in order, text matching the H2 word for word
Byline	The one authoritative legal-name spelling, linking back to the person page
~~**AI voiceover or a slide deck instead of a real person**~~	~~do not do~~ Don't
~~**Buying views, boosting, chasing plays**~~	~~do not do~~ Don't
~~**Jump cuts, background music, effects, a hook-first opening, cutting into Shorts**~~	~~do not do~~ Don't

Required Don't

Figure: the blocks of a video page — every one is text for retrieval to read; no gimmicks of any kind

94% of YouTube's AI citations go to long-form video; on this leg, short clips are effectively invisible. Timestamps follow the H2s because **78% of videos with timestamps get cited repeatedly**; the

description holds the full transcript because description length correlates with being cited (r = 0.31) — the engines crawl the text around the video, and the description is the main crawlable text on this leg.

The script is simply that thick page's H2 questions read out one by one, one section per H2, in the page's own order, adding no fact that is not already on the page — everything comes straight from the fact table. The transcript is public advertising text; the outline and the full transcript go through the compliance gate in the same process, with the same signatory, as the thick page, and nothing is published before sign-off. On the strict side, the video covers only process steps and duration, device or product model numbers, the cost breakdown and who it is not suitable for: never outcomes, never before-and-after images, no clients on screen, no mention of peers. Who appears on camera follows the anchor layer: the organisation lead, the named professional, or the product owner.

> **Example** (training) A public course clip is already a long video with its own chapters; timestamps follow the H2s of the course outline.

Division of labour and time are fixed: the script takes 1 hour (by the page writer, produced in the same batch as the thick page), filming takes 30 minutes (one take, never a full reshoot), post-production takes 1 hour (only four things: trim the head and tail, add captions, export the transcript, and set timestamps by H2), compliance review happens in the same batch as the thick page, and logging takes 10 minutes (one line into the outreach ledger). Do not publish the video before the thick page is live — the description's first line needs that URL, and publishing first just leaves a 404. This is a mechanical gate within the page-writing flow: see → General Edition 5.4 The 30-day writing order and the three mechanical gates.

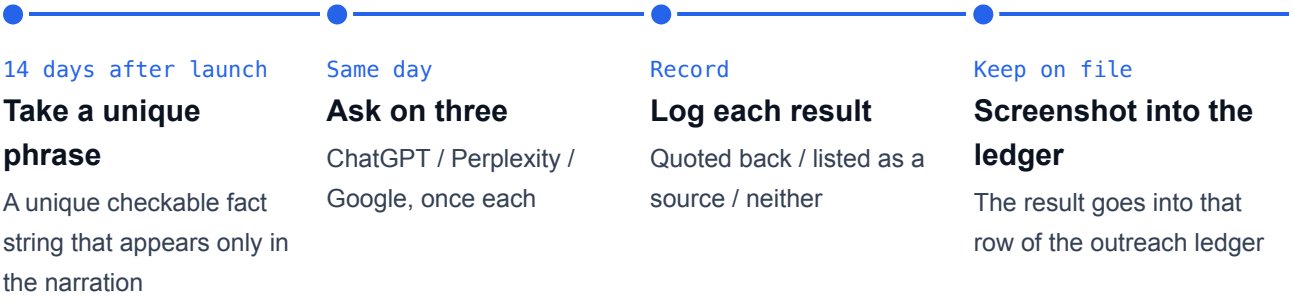


Figure: video acceptance, valid even at n = 1: it measures whether it was read, not which position it ranks

Never track view count, subscriber count or completion rate — these three numbers have almost no relation to citation frequency; writing them into a report just reports a number already known to be irrelevant.

Verdict (six items)

1. Every live thick page has a matching video, or a line in the page register stating why not;
2. Length ≥8 minutes, the named expert on camera throughout;

3. All three description items present: the thick page's URL (line 1) + the full transcript + the person-page URL;
4. Timestamps set for every H2, text matching the page's H2s word for word;
5. The transcript has passed the ban table for your side and has been signed by the signatory;
6. The 14-day quote-back test has been run, all three outcomes recorded, screenshots on file.

Common mistakes

- Making a 2–3 minute promotional clip instead;
- Putting only a one-line blurb in the description and skipping the transcript, effectively hiding the body text;
- No timestamps;
- Jump cuts and fancy on-screen text added for looks, breaking up the narration so the transcript gets chopped up;
- The named person does not appear, replaced with a slide deck plus AI voiceover;
- The video states a fact that is not on the page;
- The transcript published before it clears the compliance gate;
- Publishing the video before the page is live, leaving a 404 in the description;
- Reporting view count, or stopping the effort because view count is low.

6.8 Off-site verdicts, common mistakes and rituals to cut

What you'll do in this section: This quarter's off-site work is not done until every one of nine verdicts is ticked. Hit the abstain line (12 or more "not reachable" items in the baseline Top 20) and off-site work stops for that intent — state the reason and switch to monitor only. Check the ritual table, cut the actions that never move rankings, and put the saved time back into week-1 identity fixes and weeks 1–4 off-site work.

Nine verdicts

1. Every URL in the baseline Top 20 has had its three-state check and compliance verdict done, no gaps;
2. Every update letter, correction letter and removal request for the three classes "already present", "correction only" and "request removal" has been sent;
3. Every Group B item that is reachable and verdicted can do has been sent out, none left, each logged in the outreach ledger with a date and a live URL; a Group A item with all three present is sent in passing, anything missing one goes to the reserve list — no floor on the number sent, and this is never counted as a gap; Group C is not counted in this item. One more item counts equally: the all-products price guide page (English + Chinese) is live and entered in the page register, all three of its specs ticked — one page covering all products / set alongside the official benchmark

(how it is set alongside varies by side: on the strict side the official benchmark stands as its own section, with none of your own prices in that section — see → General Edition 5.12 Price guide page (1): why it is always the first page (pt02)) / the H1 carries the year and `dateModified` is up to date. Neither of these two items depends on whether the third party accepts your submission — they are judged only on whether it was sent and is live; whether the third party takes it up is still never promised;

4. All five business profiles claimed, the URL points to the landing page, not the homepage, and every free-text field has passed the banned-word list;
5. Every off-site cost marked "paid listing" line by line, no item billed per lead or by commission;
6. Review actions run by side: **zero review-request actions on the strict side; on the light side, any request must first have a signed compliance memo, with an actual sending record**;
7. The annual-ranking calendar reminder has been built;
8. Both off-site-asset maintenance and paid-listing disclosure are written into the process;
9. Every off-site material has zero hits on the compliance memo's banned-word list, with a signed materials sign-off sheet; on the side where switch A (agency liability) is on, the publishing authorisation is also on file. **Missing any one of the three means do not send** (see the send gate in → General Edition 6.2 The send gate, paid listings and the citation-slot table).

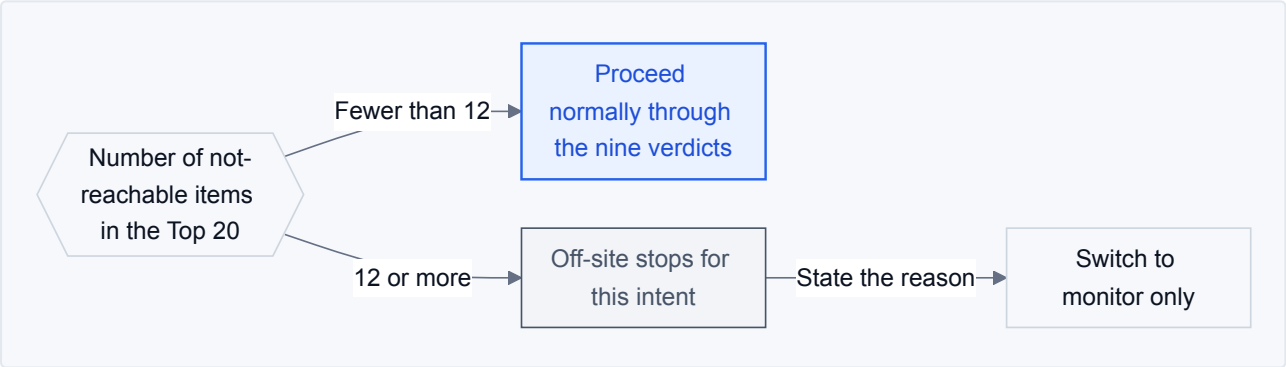


Figure: hit the abstain line for not-reachable items, and this intent's off-site work stops, switching to monitor only

Common mistakes

- Using the aggregated `domains_cited` as input — it strips out the target sites;
- Ticking the three states by domain rather than by URL — the same domain can have some pages open and others not;
- Testing outreach one letter at a time — the whole Group B batch should go out in one go;
- Sending Group C pure directories to pad the letter count — the count was never the target; stop once Group B is done and Group A's opportunistic slots are sent;
- Waiting for the Group A list to build up before starting — Group A is a structural gap and never fills up; the right move is a self-built all-products price guide page;
- Treating comparison-page type as something only others can give you — density is counted by page type, not by domain ownership, and a self-built page earns it the same way;

- Reporting an expected reply count externally;
- Writing a paid listing as "included in an independent review";
- Paying for a laudatory-titled ranking slot on the strict side, thinking that marking it "paid listing" makes it compliant;
- Citing any third-party directory that claims to be certified by a regulator — the only citable official credential is your own licence/registration number plus a link to the official lookup page;
- Using a per-enquiry or per-commission traffic platform when switch F (referral commissions banned) is on — this is not a purchase, it is fee-sharing for clients;
- Filling a business profile's URL with the homepage, or posting promotions or reposted reviews as a business post;
- Not sending an update letter after a price changes — the stale price left in a third-party profile gets quoted by AI as current;
- Nominating yourself for a peer-nominated ranking — the ranking's own rules explicitly ban self-nomination, and a single self-nomination can count against you for the whole edition;
- On the side where switch A (agency liability) is on, publishing material externally under an agent's name without a publishing authorisation on file;
- Sending a review-request template on the strict side;
- **Setting long-term disposition from a single month's fluctuation** — see → General Edition D.1 Number discipline: how to label numbers, and what stays internal hard rule 5;
- **Treating off-site as a one-off project** — see → General Edition 6.6 Annual rankings and off-site asset decay.

The rituals to cut are, every one of them, actions that "get done but never move rankings"; they share one table with the target-picking layer, and the five overlapping items are written only in → General Edition 1.4 Three kinds of things not to do:

Cut	Replace with
A full baseline of 30 questions × 5 rounds × 2 engines	Two-stage probing + one same-week retest round
Page-by-page audits + ledgers + calendars + a strategy deck	A one-page work order on D1 afternoon
Multiplying four scored factors, divided by cost	Annual opportunity value = price per order × gross margin × monthly capacity cap × 12
Always running URL-level and answer-level double counting	Deferred until the first retest
A deep review run on your own site	Aimed only at competitors; deep-review questions never enter the question pool
Positioning-statement checks + fields logged but never used	Delete them; turn negative/excluded tendencies into same-day queue-jumping alerts
Paid listings and submissions as the main line	Opportunistic slots only — by their mechanism they take none of the three paths
Six-way classification / JSON-LD / permutation testing / notarised screenshots / pure-symptom questions	See the fixes in 1.4

Figure: the off-site and target-picking rituals to cut, and what to replace them with

Paid listings cannot move rankings, and the reason is mechanical: what money buys is only exposure or an intelligence product — not one word about you on that page changes. Of the three paths (putting facts into an already-cited page, expanding the candidate set, raising confidence that it is the same business), it takes none; see → General Edition 1.3 Three gates and three paths. The 0.3% figure is the citation share of paid advertorials: it is only supporting evidence, not the basis for this rule, because paid listings and ad placements have never been measured separately. Cutting the first six rows of this table plus the five items in 1.4 **saves about 81 person-hours per quarter** (summed line by line, estimated; the same-week retest is already netted back in as an added cost).

Chapter 7

Monthly retest: how to tell whether it's working and what to do next month

Ask the same questions every month, and tell real change from random noise

7.1 What a retest produces, and the rulers

7.2 The noise band, the page-level signal and the ten monthly steps

7.3 Not-moved triage and next month's three points

7.4 Holding position and the 90-day settlement

7.5 Ruler discipline: change the ruler and nothing is comparable

Ask again every month with the ruler you froze for good. The verdict is never "it worked" or "it didn't" — write **which three points to hit next month**. Alongside it, keep three firm numbers that don't depend on AI's own randomness, to prove that this month's work really landed.

7.1 What a retest produces, and the rulers

What you'll do in this section: turn the monthly retest's verdict from "did it work" into "which three points to hit next month". When you're done, you'll have five rulers, each answering one question, the limits each one must print exactly as they are, and four sentences you'll never say again.

The order is fixed: **read the free first-party data first, then look at the probes that cost money.**

Why we do not judge with statistical tests

Statistical judgement (internal reference only)	Raw counts + noise band (used for the verdict)
Wilson three-state: moved / seems to have moved / not moved	Seat count, pages on the list, factual errors (count)
The threshold for "moved" is out of reach on a small sample	Retesting 3 times in the same week is enough to measure the noise band
Default output: "not demonstrated"	Output: next month's three work orders
Only used to trigger not-moved triage	Used for the verdict, and to schedule the work orders
Permutation test: at n ≤ 10 it changes no decision	Don't run it — that's just ritual

Figure: On a small sample, statistical judgement can only output "not demonstrated"; raw counts + a noise band are what actually produce next month's work orders

Run it on a statistical basis, and week 13's default output is necessarily "not demonstrated" — **you walk in already carrying the conclusion "this didn't work"**, whether or not anything actually moved (for how the threshold is worked out, see the evidence in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests). The fix isn't to loosen the statistics — it's **to change what you accept as proof**. Rankings are moved by work orders, not by p-values.

If you're doing this yourself, **the minimum viable ruler set = three numbers + one page-level signal + one web control leg**. You don't need a Wilson interval, a control page, or a permutation test — but you do need to honestly separate "actually moved" from "noise". The three numbers are the seat count, on the list and cited, and factual errors (count) below; the page-level signal and the noise band are in

7.2; the full specification of the web control leg appears only in → General Edition 4.4 The web control leg and the frozen baseline (the book's only full spec).

Name the legs correctly first

The two API legs' official names are ① **OpenAI API leg** and ② **Gemini model leg**. The diagram mapping the three products to their three rulers is in → General Edition 1.2 Two legs: ChatGPT looks for the source, AI Mode for second-hand summaries; here we only add the three mistakes people most often make at retest:

- Treating the Gemini model leg's number as "visibility in Google's AI Overviews". **Gemini API ≠ AI Overviews ≠ AI Mode** — the fetch path, the way the answer is generated, and the source pool all differ; passing one leg's number off as another's is measuring A to sign off on B. Visibility for AI Overviews / AI Mode **uses only the impressions from the Search Console generative AI report**.
- Checking `Google-Extended` to judge whether AI Overviews can crawl you. **You should check `Googlebot`.** `Google-Extended` governs training and the Gemini app side; blocking it does not affect whether AI Overviews indexes you — it's `Googlebot` being blocked that means the door is shut.
- Describing the API legs' results as "what users will see in ChatGPT". Where the only evidence is from an API leg, the wording must always be qualified as "on the model API".

Five rulers, one question each

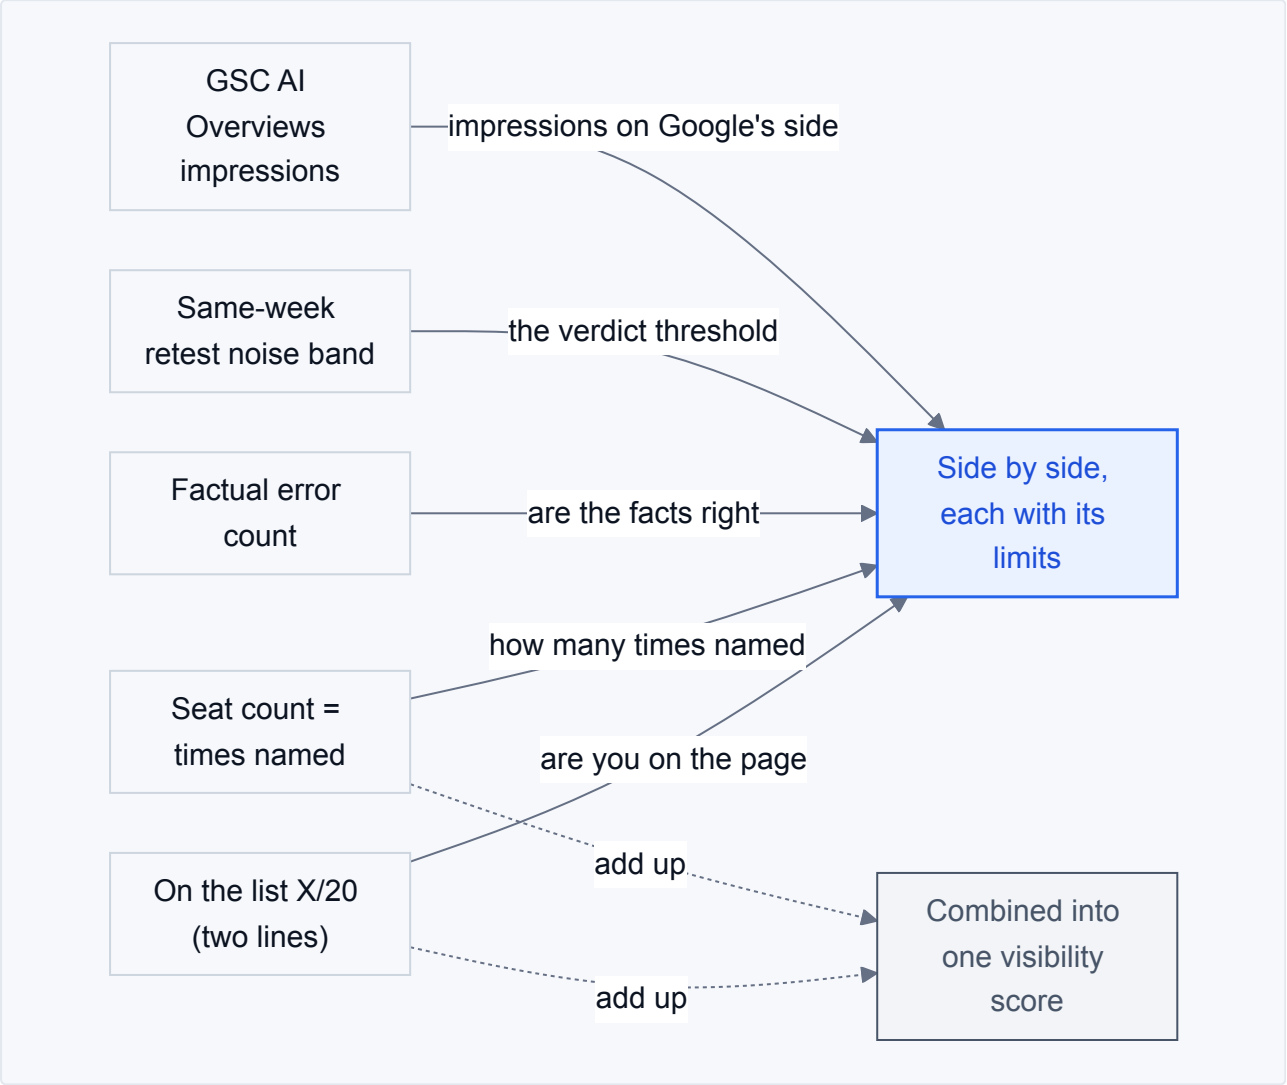


Figure: Five rulers, one question each — print them side by side only, each with its limits attached; add them into one score and any one going bad gets covered up by another

How each ruler is counted, and the limits that must be printed exactly as they are on the report:

Ruler	How you count it	The limits that must be printed exactly as they are
Pages on the list (report as two lines: **on the list X/20** and **on the list and cited this month Y/20**)	Of the frozen baseline Top 20 sources, how many pages now actually feature you (open each one by hand + screenshot). **The denominator, 20, does not change all quarter**	**On the list ≠ cited**: AI may cite this page without mentioning your line — being on the list is a necessary, not a sufficient, condition. The reading is fixed: on-the-list rises but on-the-list-and-cited doesn't, **for two months running** → you're investing in a container that isn't being cited; next month, invest only in sources that were actually cited this month

Ruler	How you count it	The limits that must be printed exactly as they are
Seat count	How many times your name was named across this month's N question-and-answer runs; also log **this month's total names** (the sum of every business's name count) and each competitor's seat count	It's a sampled count under your own question pool, engine and number of passes — it does not represent the whole market; total names shifts with how long the answer is, so **you must print total names and seat share together**; it **scales linearly with the number of passes**, so if the number of passes changed, the two numbers must not be shown side by side or subtracted from each other
Factual errors (count)	A person reads the brand six (6 questions × 3 rounds × 2 engines = 36 runs / month, frozen for the quarter, defined in → General Edition 3.5 The brand six questions, and what to do when you find an error), copies down every fact AI got wrong, and attaches a screenshot of the original	The hardest cell in the whole ruler set; it measures only whether the facts are right, not ranking. **If you can only watch one number, watch this one**: the error is copied by AI from the source page it retrieved, so fixing the source moves both legs together — of all the rulers, it is the least sensitive to skew between the API and web paths
Same-week retest noise band	See 7.2	The noise band itself is also a sampled value
AI Overviews / AI Mode impressions	GSC generative AI report impressions, exported by page and by country (free, first-party)	Only impressions — no names, no position; it's already included in total impressions, so **it must not be added to total impressions**; flag it when a blank value is exported as 0

Seats are simply times named, and the measure has only one name. Do not write "the raw seat count" and "the answer-level named count" as two separate lines — they are the same number, and the only effect of splitting them into two is that the noise figure you're banned from reporting on its own sneaks back into the acceptance test under a different name.

How you split changes by industry; the definitions do not

The definition of a ruler doesn't change by a single word; the only thing you can change is how you split it:

- **Seats can be split into two**: when a category has both an answer shelf and a shopping shelf, count brand seats (answer shelf) separately from SKU card placement and top-spot share (shopping shelf), each with its own noise band. **This answer's total names is fixed by measuring the category itself** — never reuse another category's number.
- **The denominator can be split by work unit**: when two fronts' named sets barely overlap (by product line, or by school stage × subject), report the denominators separately; merge them into one denominator and the ups and downs of two independent fronts will cancel each other out.

- **Factual errors can be split into two columns**: log "errors in your own source" separately from "third-party errors" — the two have completely different repair paths and time-to-effect (7.3, layer 2).
- **How many rounds go into the noise band is an industry variable** — measure it yourself in the first quarter. The fourth reference number (things like enquiry source) is chosen per industry; **it's the only cell that varies by industry, and it never enters acceptance**.

> **Example** (e-commerce) Seats split into two: brand seats (answer shelf) + SKU card placement and top-spot share (shopping shelf); the noise band is also measured as two, one for brand seats and one for SKU card placement.

The source skew between the web version and the API is systematic, not noise, and adding more rounds doesn't fix it — the number appears exactly once, in → General Edition 4.4 The web control leg and the frozen baseline (the book's only full spec).

Four things never to say

Never say	Say instead
"We're on the list, so AI will recommend us"	On the list X/20, of which cited this month Y/20
"Seat share 1.29% = market share 1.29%"	Asked K times, named X times, noise band ±N
"Up 2 seats, heading the right way" (within the band)	Within normal variation, cannot be determined
Adding two rulers into a combined "visibility score"	Each ruler on its own line, side by side, with its limitations

Figure: Four things never to say, each paired with a reproducible way to say it instead

Writing a report, briefing your boss or a client, or even summing up in your own head — the four sentences on the left must never appear. Why: on the list ≠ cited; seat share is a sampled count under your own question pool, not the market; a rise that hasn't cleared the noise band is treating noise as a result; and two rulers with different denominators and different failure modes, once added together, let either one go bad while the other one covers for it. The noise-band statement that must be printed word for word in the monthly report is in → General Edition B.2 Approved wording and scripts.

Retesting also has three more criteria; the full statement is in → General Edition 1.4 Three kinds of things not to do — here they're just applied:

- An observation that is single-run, single-engine and n ≤ 3 can only serve as directional evidence — **you cannot declare "zero visibility"**;

- The wording of the question pool can produce a false zero: when bare category words dominate, suspect the ruler first (7.3, layer 4); evidence in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests;
- **Even if pages-on-the-list has risen, still run triage** — the trigger looks only at seats (7.3).

7.2 The noise band, the page-level signal and the ten monthly steps

What you'll do in this section: get three more things every month — a measured noise-band value, a list of page-level-signal hits, and a monthly schedule with people and hours assigned. With the noise band you know which rises and falls you're allowed to write into your conclusion; with the page-level signal you know which page was actually read and taken.

The noise band: the cell people most often get wrong

What to do: run the same batch of questions, same engine, same number of passes, again **3 times within the same week**, giving X′ / X″ / X‴. **How long**: about 1.5 hours/month (the machine runs it; a person only checks it lands on disk). **How to check**: three independently saved answer packs, all timestamped within the same week; the noise-band value is written on page 1 of the monthly report.

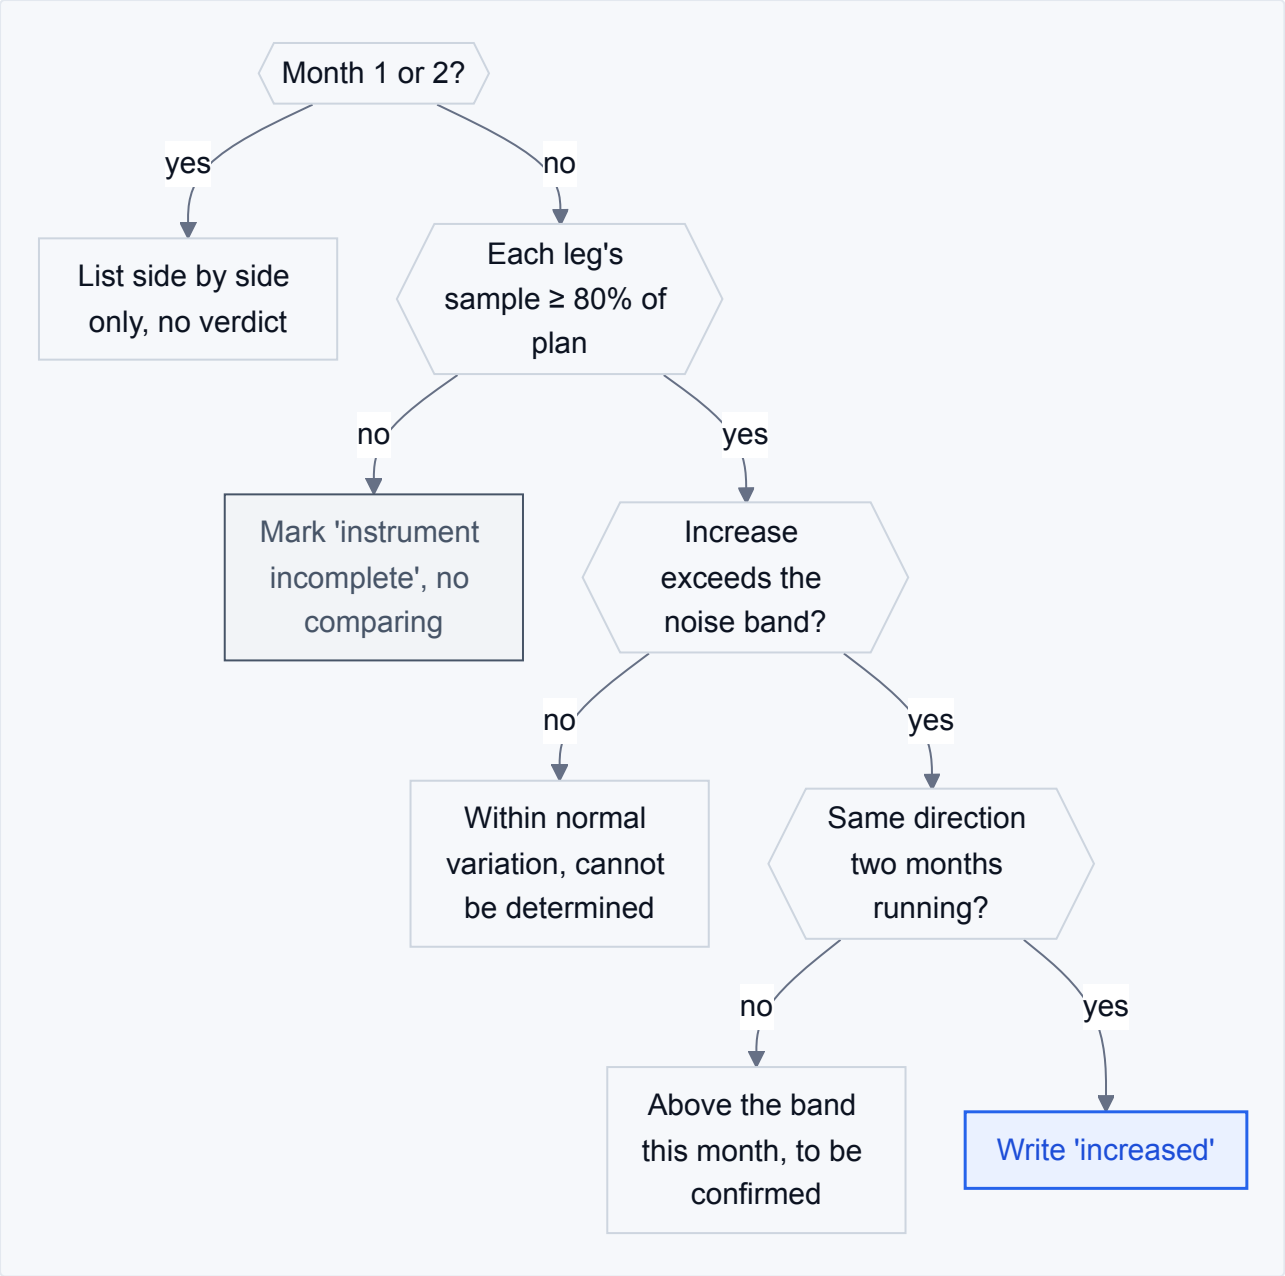


Figure: A rise in seats doesn't mean you can write "increased" — it has to pass four checks in a row, and each one you fail has a fixed way to write it

The wording at each check in the figure is fixed: for months 1–2, mark "noise band baseline not yet stable (K/3 months recorded)"; in a month where the sample falls short, print each leg's figures separately, **do not print a total, and do not compare with last month**. What the figure can't show is how the noise band itself is defined:

Rule	Why
Run it 3 times, not once	If you use one sampling difference as the threshold for another, then under the null hypothesis that "nothing happened" **the two have the same distribution, so the false-positive rate is about 50%**. Running it 3 times brings the false-positive rate down from 100% to 50%, not down to 0

Rule	Why
Noise band = the range (max − min); the floor is fixed at ±1 seat — even if all three runs come out identical, record it as ±1	When the range is 0, any +1 would be judged "increased" — the gate fails in reverse
The noise band baseline is the median of three months, not the maximum	Taking the maximum is a **monotonically non-decreasing ratchet** — the longer you run it, the harder it becomes to declare any improvement, and you end up in a closed loop where "the worse the ruler, the more it excuses you"
If the noise band widens for two months running → add more passes and retest; you may not invoke "instrument defect, not counted this month" again	An excuse clause used more than once becomes an excuse machine. Adding more passes is changing the ruler — per 7.5, start a new baseline
If one leg's valid sample that month < 80% of the planned volume → mark that month "instrument incomplete"	When one leg is missing for the whole month, the absolute cross-engine count gets cut in half. This is the same problem as "adding passes doubles the seat count" — just the other side of it

The page-level signal: the only page-level evidence of cause

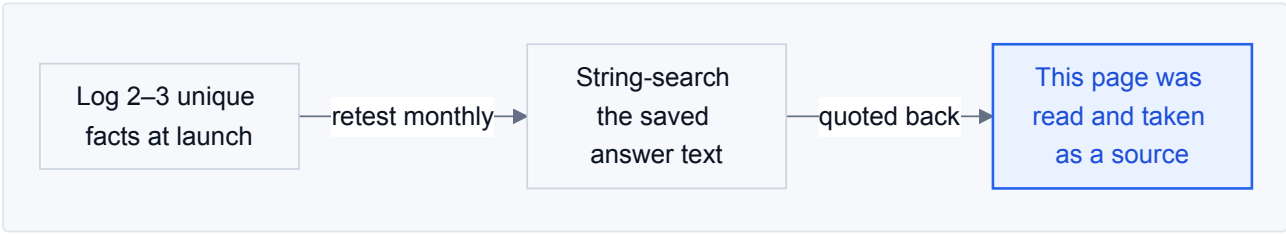


Figure: The page-level signal measures "was it read or not", not "which position" — so n = 1 still counts

What you log must be a **checkable fact string unique to that page** — for example, the `S$6,800` fixed into that page's price list. Adjectives, or sentences that also appear on other pages, don't count; finding them proves nothing about which page was actually read.

It has nothing to do with seat count or the noise band, and reuses the answer-text packs you're already storing — **the marginal cost is close to zero**. Seat count tells you whether things are rising overall; the page-level signal tells you which page is doing the work — record the two numbers separately, and never infer one from the other.

Ten monthly steps, grouped into eight stages

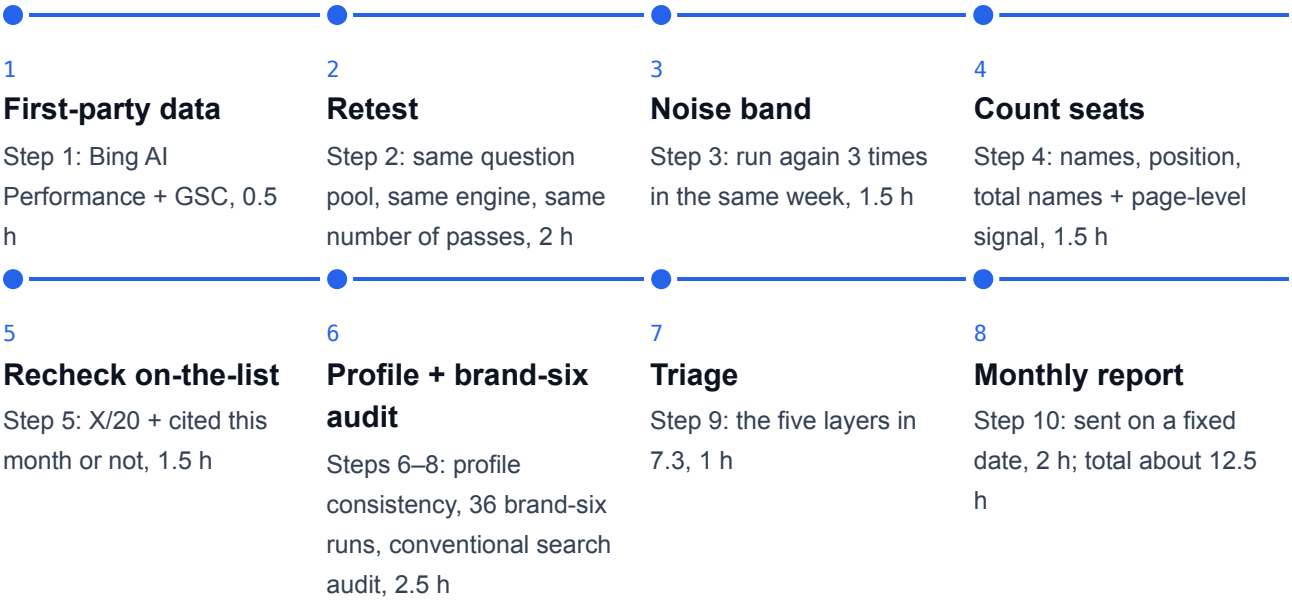


Figure: Ten monthly steps grouped into eight stages; the first six stages are just bookkeeping — what moves rankings is stage 7, triage

Steps 1–4 are only responsible for accurately measuring what the shelf looks like right now; what actually moves rankings is step 9 — **without step 9, the first eight steps are just bookkeeping**.

Jump-the-queue rule: if AI produces a false statement about you or shows a negative tendency → **jump the queue and fix it first; don't wait for the monthly report** (the five situations that jump the queue are in → General Edition 3.5 The brand six questions, and what to do when you find an error).

The web control leg's 40 minutes **is not in this table** — schedule it as a separate day at the start of each month; the full specification is in → General Edition 4.4 The web control leg and the frozen baseline (the book's only full spec).

Who does each step, what they hand over, and who checks it:

Step	Who	Output	Who checks	Points the figure doesn't show
1	The report owner	`first_party_<month>.json`	Self-check	Export this month's landing-page impressions and Citation Share
2	The person running probes	`probe_<month>/` , the raw text of every answer	Self-check	Print "not measured" for any engine not measured because the quota ran out

Step	Who	Output	Who checks	Points the figure doesn't show
3	The person running probes	`probe_<month>_retest/` + the noise-band value	The report owner	—
4	The person running probes	`seats_<month>.csv` + the page-level-signal hit list	The report owner	Which businesses were named in each answer, at what position, and **this answer's total names**
5	The off-site lead	`offsite_ledger_<month>.csv`	The project lead	Open every item in the baseline Top 20 by hand: are you on it, at what position, is the information out of date; while there, tag each one with the five binary page-type features
6	The profile owner	Same ledger as above	The project lead	The three review numbers are a **passive observation**; five-profile consistency
7	The person running probes	`brand-6.jsonl` + `entity_errors.md`	Read by a person, no skipping	6 questions × 3 rounds × 2 engines = 36 runs, producing the factual-errors count
8	The person who edits the site	Audit page (with the four identities' four receipts + screenshots of the three `nosnippet` greps)	The project lead	Non-brand clicks, the canonical Google picked, the index status of key URLs, that the four identities' fetches still match (defined in → General Edition 2.3 Gate 2: four identities, fetched live (read-only)), whether the three `nosnippet` controls have recurred this month
9	The project lead	Next month's work orders	Look it over again the next day	Run the 7.3 triage, produce next month's three points
10	The report owner	The monthly report + the raw answer pack	Read page by page by a person	Sent on a fixed date

7.3 Not-moved triage and next month's three points

What you'll do in this section: for a month when seats don't clear the noise band, come away with a triage verdict of "stopped at which layer", plus next month's three work orders, laid out mechanically by priority order; add to that a fixed-format page 1 for the monthly report, and a competitor board for your eyes only.

There is only one trigger

> **This month's seat increase did not clear the noise band → always start at layer 1; no skipping layers. Whether pages-on-the-list rose or not has nothing to do with this rule.**

Why you can't add "and pages-on-the-list did not increase": pages-on-the-list is **the output of your own work**, and a monthly +1 is almost routine. Fold it into the trigger as an AND condition and you've fitted the whole triage table with a **switch you can flip off whenever you like**, and no month all quarter will ever really go through triage. **No variable you control yourself may appear in the trigger condition** — any new trigger condition proposed later must be checked against this rule.

Five layers, cheapest first

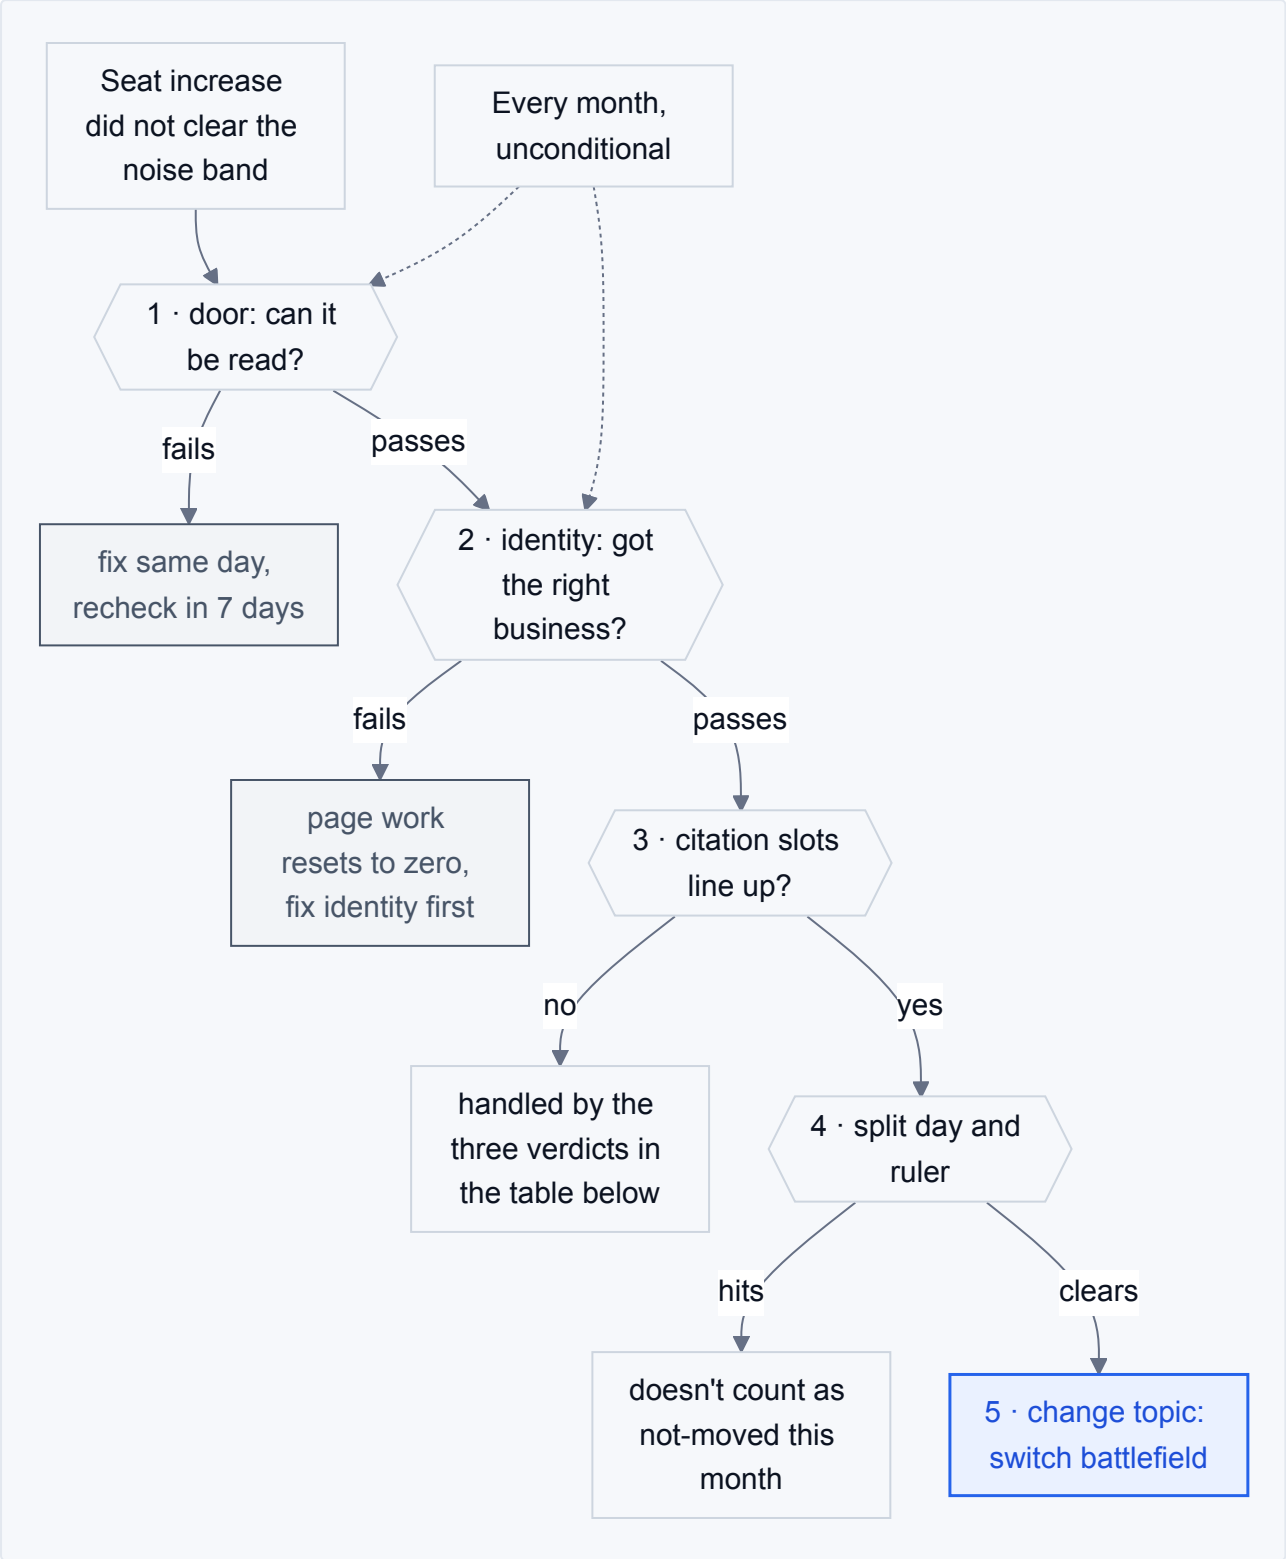


Figure: Not-moved triage checks layers cheapest-to-rule-out first, layer by layer, no skipping; the door and identity layers are checked every month unconditionally, even when seats rose

The door and identity layers are checked unconditionally because failure at either one is **silent**: when a page was never read at all, or AI has the wrong business, seat count can perfectly well be rising for

some other reason, and you won't find out until next month — by which point a whole month has been wasted. Which fix each door-layer symptom maps to is looked up in the two diagnostic figures in → General Edition 2.9 Door-layer troubleshooting and this chapter's acceptance checks; they aren't redrawn here. Three of these are the easiest to misjudge: the Google leg has seats while the ChatGPT leg is zero — **check CSR first, not your question selection**; don't use `GPTBot` as your evidence, it's a training crawler — for the retrieval legs look at `OAI-SearchBot` and `ChatGPT-User`; and don't draw a conclusion from what `curl -A "OAI-SearchBot"` returns — both directions of that test lead to the wrong verdict.

Layer	What to check that day	Test and action	How to say it in the monthly report
1 · Door	① Access logs: in the past 30 days, has a real `OAI-SearchBot` hit these pages, and what status code came back (**the only direct evidence**) ② have the three `nosnippet` controls been switched back on by the CMS or a plugin ③ do the four identities match word for word ④ `site:` search to check Bing indexing ⑤ paste the URL into ChatGPT and have it **quote back the price line word for word** ⑥ is the Foursquare category and business status right, is the Yelp entry there	Failing any one = **a door problem, not a targeting problem**: fix it the same day, recheck in 7 days, **don't change topics or add new content pages this month**	"Last month ChatGPT simply wasn't reading those pages at all; the reason was X, it's fixed today, we'll recheck in 7 days"
2 · Identity	Read the brand six by hand: is it talking about this business? Are the name, address, registration number and price right? Any namesake confusion? Any complaint or negative tendency?	Wrong business, or facts wrong → **all page work resets to zero and gets recalculated**; fix identity first: tick off and screenshot each of the six controllable sources (own-site `/facts` → Google Business Profile → Bing Places → Foursquare → Yelp → Apple Business Connect) + send a correction letter to the cited source. **Mark every error as either "third-party page source" or "business profile source"** — mixing the two makes profile-type errors run late systematically	"AI mistook you for another business (or got your price wrong); this month we deal with that first"

Layer	What to check that day	Test and action	How to say it in the monthly report
3 · Citation slots	The intent-level top 10 URLs by times cited this month, checked one by one against the four columns of the citation-slot table (`four states / judged on / basis / judged by`)	**Judged reachable at the time, entry point gone now =** the shelf really changed → switch to "hold", go after that new source, put it first on this month's off-site work orders **Judged "to ask" at the time, or the basis column is blank =** the judgement drifted → fill in the basis this month and recompute the abstain line; this month it doesn't count as a wrong target **Fewer than 2 reachable-enough URLs in the intent-level top 10 =** wrong target → move this intent out of the main-attack pool and swap in another; pages already invested in become long-tail landing pages and get no more investment	"A got into some directory last month and pushed you from 3rd to 5th; go fight for this this month" / "Last month's basis for this one wasn't kept on record; fill it in this month, then judge" / "We can't get into this slot for this question because we picked the wrong battlefield at the start; swap it"
4 · Split day and ruler	① the page has been live less than two weeks ② the URL changed after launch ③ are the frozen questions bare words ④ has the engine / model / number of passes been touched ⑤ is the noise band bigger than last month ⑥ is the web-control-leg overlap < 50% ⑦ was one leg's valid sample < 80% of the planned volume	Any hit → recompute the split day, or mark the question pool "instrument defect", and **it doesn't count as "not moved" this month**. If ⑥ hits, this month's seat conclusion always gets the non-extrapolation sentence printed alongside it (→ General Edition 4.4 The web control leg and the frozen baseline (the book's only full spec)), **and it must not be used as a reason to change topics**. The "instrument defect" excuse **may not be used in two consecutive months**: the second time, always add passes and retest instead	"This batch of pages has been live less than two weeks; the first comparable point comes next month"
5 · Change topic	Layers 1–4 all pass, and it still hasn't moved	Seats are saturated, or the demand isn't real → **this is the only case that's actually "change topic"**: go back to target-picking and pick again, and it must be **a different battlefield, not just different wording**	"The shortlist for this question is locked; swap in a battlefield we can actually get into"

Two disciplines:

- **You may not skip layers 1–4 and go straight to writing "demand doesn't exist".**

- **Stopping at the same layer, unfixed, two months running = an execution problem, not a shelf problem** — handle it ahead of the queue next month, and the same excuse may not be used again.

The definitions of the four states, the abstain line and the citation-slot table are in → General Edition 4.6 Four states, two denominators and the abstain line and → General Edition 6.2 The send gate, paid listings and the citation-slot table.

Turn the triage result mechanically into next month's three points

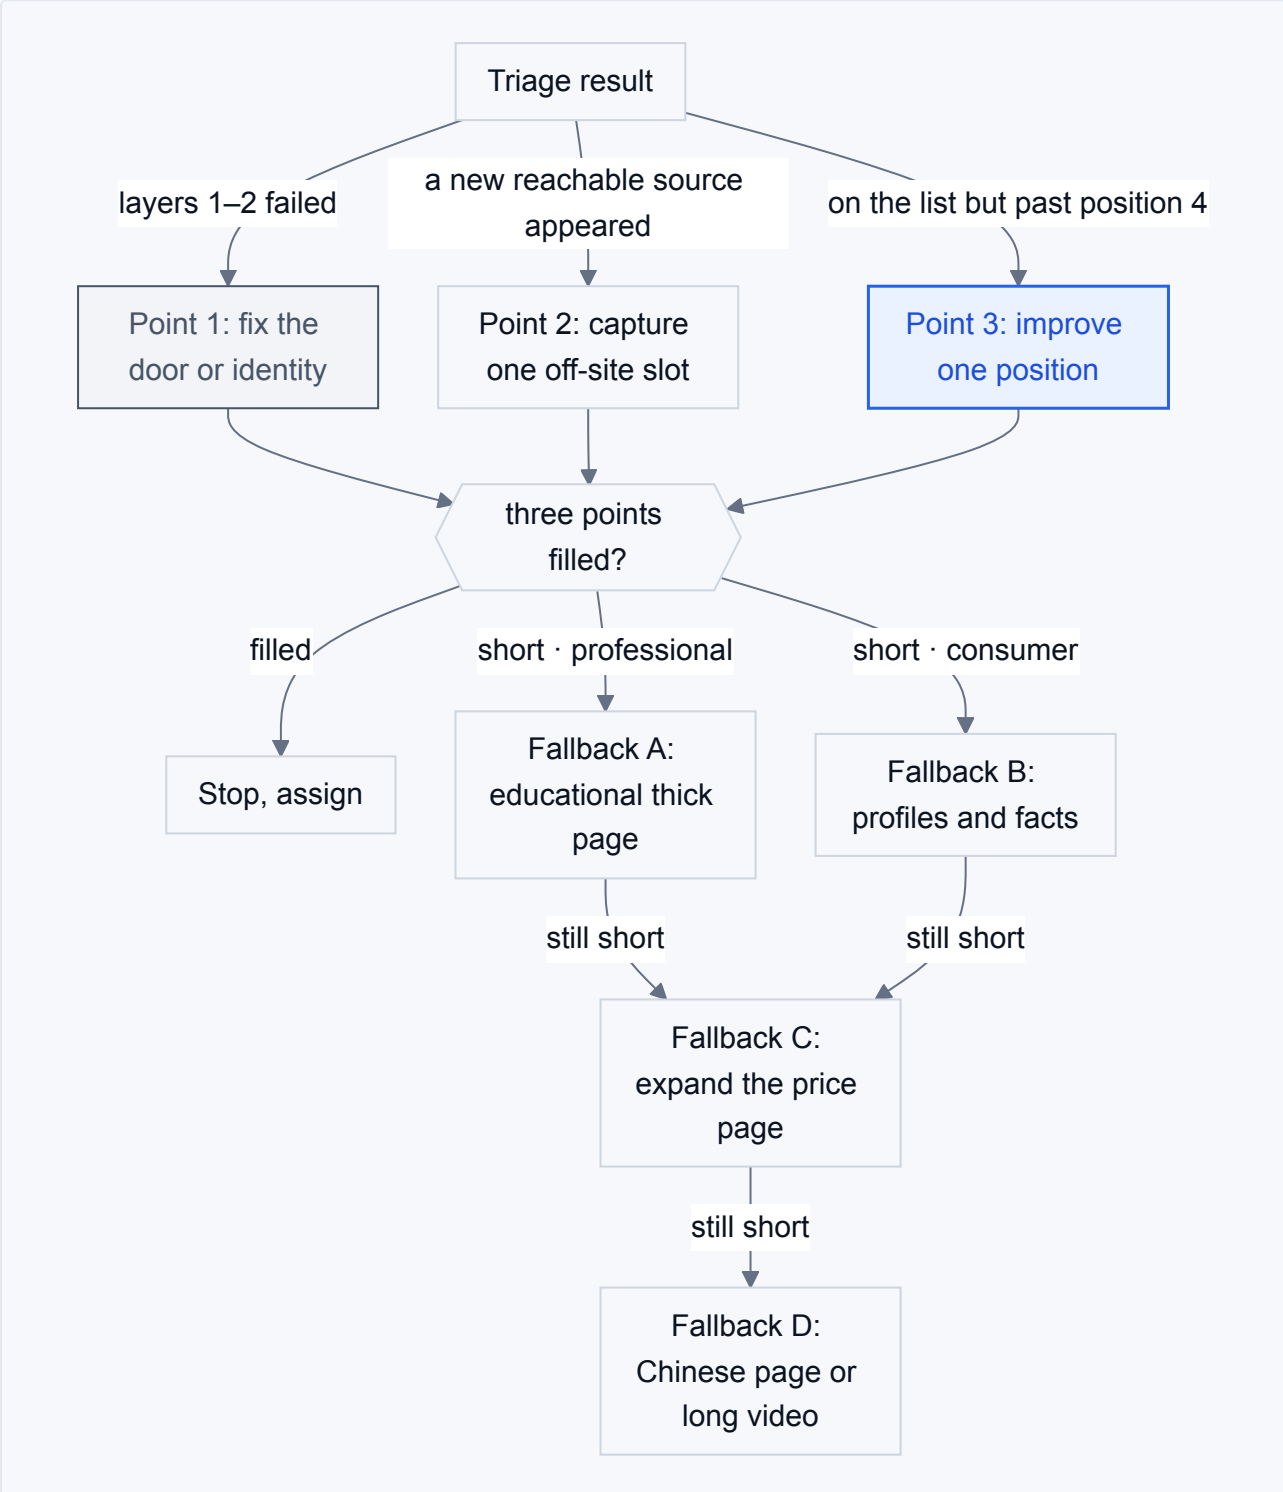


Figure: The triage result is mechanically converted into next month's three points, by priority order; stop once you've filled three, and if you can't, fall back by content form — no guessing

How each cell in the figure picks that one item, and what work order it assigns:

Priority	Which one to take	Next month's work order
Point 1	Triage failed at layer 1 or layer 2	Fix the door / fix identity. **Once triggered, it automatically takes point 1, and no new content pages are added this month**
Point 2	A source that has **newly appeared** among this month's cited sources, is judged "reachable" under the four states, is allowed on this side, and does not have you on it	Take the 1 source **actually cited this month**, with the highest citation count, where **the competitor's way in ≠ platform-edited entry**, and assign an off-site capture work order
Point 3	A source where you're **already on the list** but ranked past position 4	Take the 1 with the highest citation count, and assign an update-letter / supply-materials work order. **Improving position is cheaper than winning a new slot** (cross-industry pooled basis, not verified in this local industry — carry the qualifying sentence exactly when you print it; sample in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests)
Fallback A	Three points can't be filled, and the content form is **professional**	Educational-thick-page work order: fold the landing page for the billable intent with the lowest seat count into its intent cluster and thicken it. How thick follows the page-type spec — word count is not a criterion (→ General Edition 5.7 Seven rules that hold for every page type)
Fallback B	Three points can't be filled, and the content form is **consumer**	Profiles-and-facts work order: fill in the fact fields on all five business profiles, expand the `/facts` page, complete itemised prices on the price page (wording follows the side, → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers), and complete outbound links to official and regulatory bodies. **No review work order is assigned on the strictly regulated side**: that side may not actively ask for reviews
Fallback C	Still short	Expand a price page: price and fee queries trigger AI Overviews at a high rate; evidence in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests
Fallback D	Still short	Complete the Chinese pages, or pair an existing thick page with a long video narrated by the named expert (→ General Edition 6.7 Long videos narrated by the named expert)

Content form (professional / consumer) is the second judgement, made after the side has been decided; how to judge it is in → General Edition 5.9 Content form, FAQ, Chinese pages and other-language pages. The basis for splitting the fallback by form is a cross-industry query sample — **it does not include Singapore, at about 75% confidence** — the sample's basis is in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests.

Reddit: read the number monthly; do not fix a conclusion

Read once every month at retest: **Reddit's share of this month's cited sources** = the number of this month's cited URLs that are reddit.com ÷ the total number of this month's cited URLs (two decimal places, logged on page 2 of the monthly report).

This month's share	Action
> 2%	**Schedule one Reddit work order this month** (target whichever leg it's showing up in)
≤ 2%	No investment this month — just log the number; read it again next month
> 2% for two months running	Only then may Reddit be written into the regular work rhythm

Why we don't fix a rule of "never invest in Reddit for ChatGPT": ChatGPT's Reddit citations once fell sharply in a single month, but the same kind of collapse had happened earlier and fully recovered, and ChatGPT never stopped reading Reddit — what stopped was putting threads **into the answer**. That's "the answer surface temporarily narrowing", not "this source is dead" (numbers and dates in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests). **A single month's swing is never written into a long-term rule.**

Page 1 of the monthly report: three firm numbers and two measures printed beside them

Three firm numbers (don't depend on AI's swings)	Two measures that must be printed alongside them
Pages on the list X/20 (including on the list and cited)	This month's noise band ±N seats
Reviews and star rating	Web control leg · named-set overlap
Citation slots gained / lost	Web control leg · source overlap

Figure: On page 1 of the monthly report, the three numbers on the left prove the work landed; the two measures on the right decide whether those numbers can be read as a rise or fall

The reviews cell is explained by side: **on a side that may not actively ask for reviews, reviews are only an observation item**; on a side that may ask for them, reviews are printed in the report but are never the main effort (what each side may do is in → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions).

Page 2 of the monthly report is the competitor board: who AI is recommending this month, who's up and who's down, who pushed you out (down to which question, which source page), your seat share, and **the competitor's way in** (bylined submission / labelled sponsored or paid / platform-edited entry / user-

submitted profile / cannot tell — pick one of five, all observable from the page). Fill in this column each week while you're at it, 2 minutes per entry — your competitors have already run the experiment of "what works on this shelf" for you, and **without this column you only know which page to get into, not how to get in**. The page-1 sample and the competitor board's column headers are in → General Edition B.4 Headers for work orders, ledgers and monthly reports.

> Note **The competitor board names third-party organisations and is internal material**: it must not be used in any advertising, posted, forwarded, or made public. Material that can go external is marked item by item.

7.4 Holding position and the 90-day settlement

What you'll do in this section: for a client who already holds positions, come away with a monthly five-item holding checklist and one sentence you can say externally; at week 13 (D90), come away with an honest before-and-after comparison and a keep-or-cut verdict reached by the criteria, not by hope.

Holding position: add nothing new, keep what you have

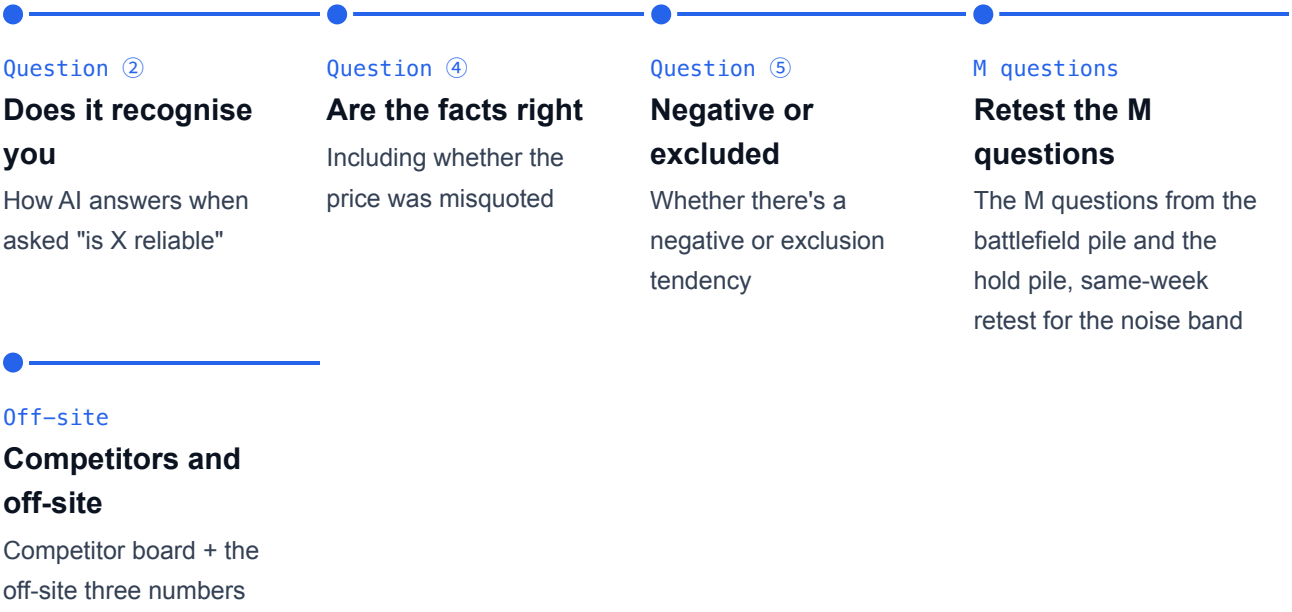


Figure: Holding position is five items a month; add nothing new, only keep what you have — skip one and you lose ground silently

The first three items are three of the questions in the brand six (the six questions are defined in → General Edition 3.5 The brand six questions, and what to do when you find an error). Four things the figure can't show:

- **What to run every month**: the M questions + the brand six (6 questions × 3 rounds × 2 engines = **36 runs/month**) × the same set of engines, archiving the original text of every answer; same-week retest for the noise band; recheck the five profiles' consistency. **Once the question composition,**

rounds and engines are set, they're frozen — no changes for the rest of the quarter — change them and the factual-errors counts before and after stop being comparable.

- **How to check**: ① factual errors (count): baseline X → this month Y, each one with a screenshot of the AI's original text, **anyone can reproduce it just by asking the same question** (X is taken only from that one freeze-week run of 36, → General Edition 4.3 Coarse screen and piles, full rounds, and the brand-six baseline) ② the number of times a negative or exclusion tendency appeared ③ profile consistency 5/5.
- **Correction**: fix everything on sources you control **within 3 working days** (the own-site facts page, the five business profiles) + send a correction letter under your own name to every third-party page that got it wrong, and keep a record. Whether and when a third party adopts it is up to them — **we make no promise there, only the two follow-ups at +7 and +21, plus a monthly recheck**. Notify the relevant parties within 1 working day of discovery.
- **Claiming profiles (one-off, about 2 hours)**: Foursquare + Yelp + Apple Business, all free. **70% confidence**: the approved line is that it's free and done in passing, **not that "filling it in gets you recommended by ChatGPT"**. The evidence (which data sources ChatGPT's local answers draw on, and which year the partnership was signed) is in → General Edition A.1 Evidence for mechanics, the door and identity; how this confidence figure is printed is in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests.

> **Say only the verifiable cell externally**: "At baseline, AI got X facts about us wrong; this month it's down to Y, each with a screenshot of the original, and you can reproduce it yourself just by asking." × Do not say "most of the 90-day increase will land on brand questions" — that is a **tabletop projection, not a measurement** (70% confidence), for internal reference only.

Why ninety days

The public line: visibility usually only starts moving after **two to eight weeks**, and a meaningful lift in citation share takes **three to four months**. Seats don't stay won — ask the same question three times, and only a tiny few of the sources cited last time are still there; **an answer has a limited number of names, and the moment a competitor follows suit, the gain is cancelled out**. So this is ongoing occupation, not a one-off build; off-site assets decay too (→ General Edition 6.6 Annual rankings and off-site asset decay) — this is the only answer to "why keep doing this" that you can give without lying.

Three things, each in its place

Settlement measure	What it is	Strength
Deliverables list	Thick pages launched and accepted, claimed profiles and statutory-register materials completed and checkable, paid slots live, retrieval-crawler door fixed with live-fetch evidence attached, month-by-month original-text comparison reports	**Can be treated as a hard commitment**: everything is in your own hands, self-verifiable, and involves no prediction of a third-party engine's or third-party editor's behaviour
Seat count (before-and-after side by side + noise band)	Three points — day 0 / day 60 / day 90 — **the same batch of questions, the same set of engines, the same number of passes**, counted item by item, printing total names, seat share and this quarter's measured noise band alongside them	**List side by side only, never promise a figure**
Pages on the list X/20 (split into two lines, "on the list" and "on the list and cited")	The denominator is the frozen baseline Top 20	**Can serve as hard evidence** (it's either on the page or it isn't, unaffected by sampling noise)

The Wilson three-state appears only on the internal review page; externally, at most, it's used to describe a trend.

> Note **Seat count scales linearly with the number of passes**: baseline 5 rounds, D90 10 rounds — **it would double even if you did nothing**, and "seat increase > noise band" would necessarily come true. So **the number of passes is frozen for the whole quarter; D90 adds not a single extra pass**. If you must add passes, the only option is to **start a new baseline**, writing "**not comparable with the previous baseline**" plainly on the same page; the old and new **must not be added together or drawn on the same trend chart**.

D90 settlement: the order of steps

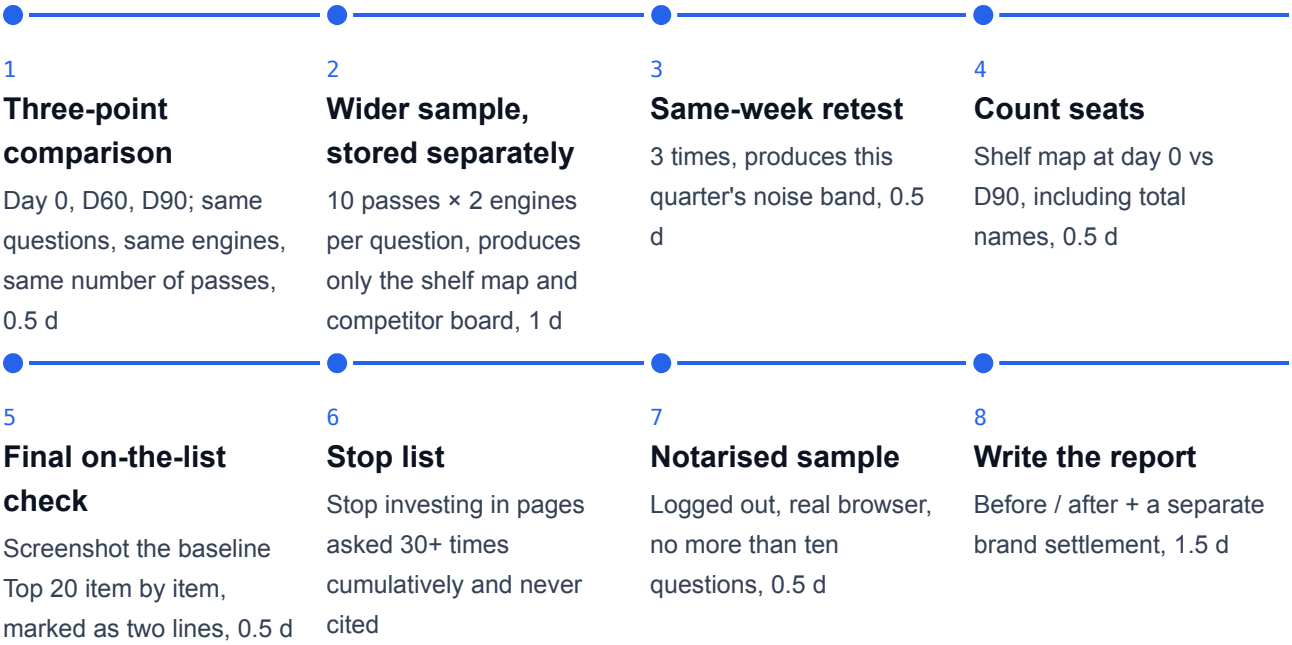


Figure: D90 is settled in this order; the main verdict comes only from stage 1, and the wider sample never enters the before-and-after comparison

Points the figure can't show:

- The wider sample is stored in `probe_d90_wide/`, **kept separate from `probe_d90/`; the two must never be added together or shown side by side on the same trend chart**, and the sampling spec is labelled separately in the figure. The shelf map is internal material.
- The stop list goes into the report with reasons given. The notarised sample **is never the headline number** — it only backs up reproducibility.
- Stage 8's "before / after" writes four things: the shelf three months ago / now; which questions went from absent to present, which didn't move and why (citing which layer triage stopped at); how competitor positions changed; the off-site three-number curve. A separate page for the brand: factual-errors count, baseline X → now Y, each with a screenshot attached, plus a *Try it yourself* card.
- There is also a separate internal three-state review (the difference between your pages and the control pages, before and after), **done once a quarter, not part of the monthly process, and never external**.

The settlement verdict

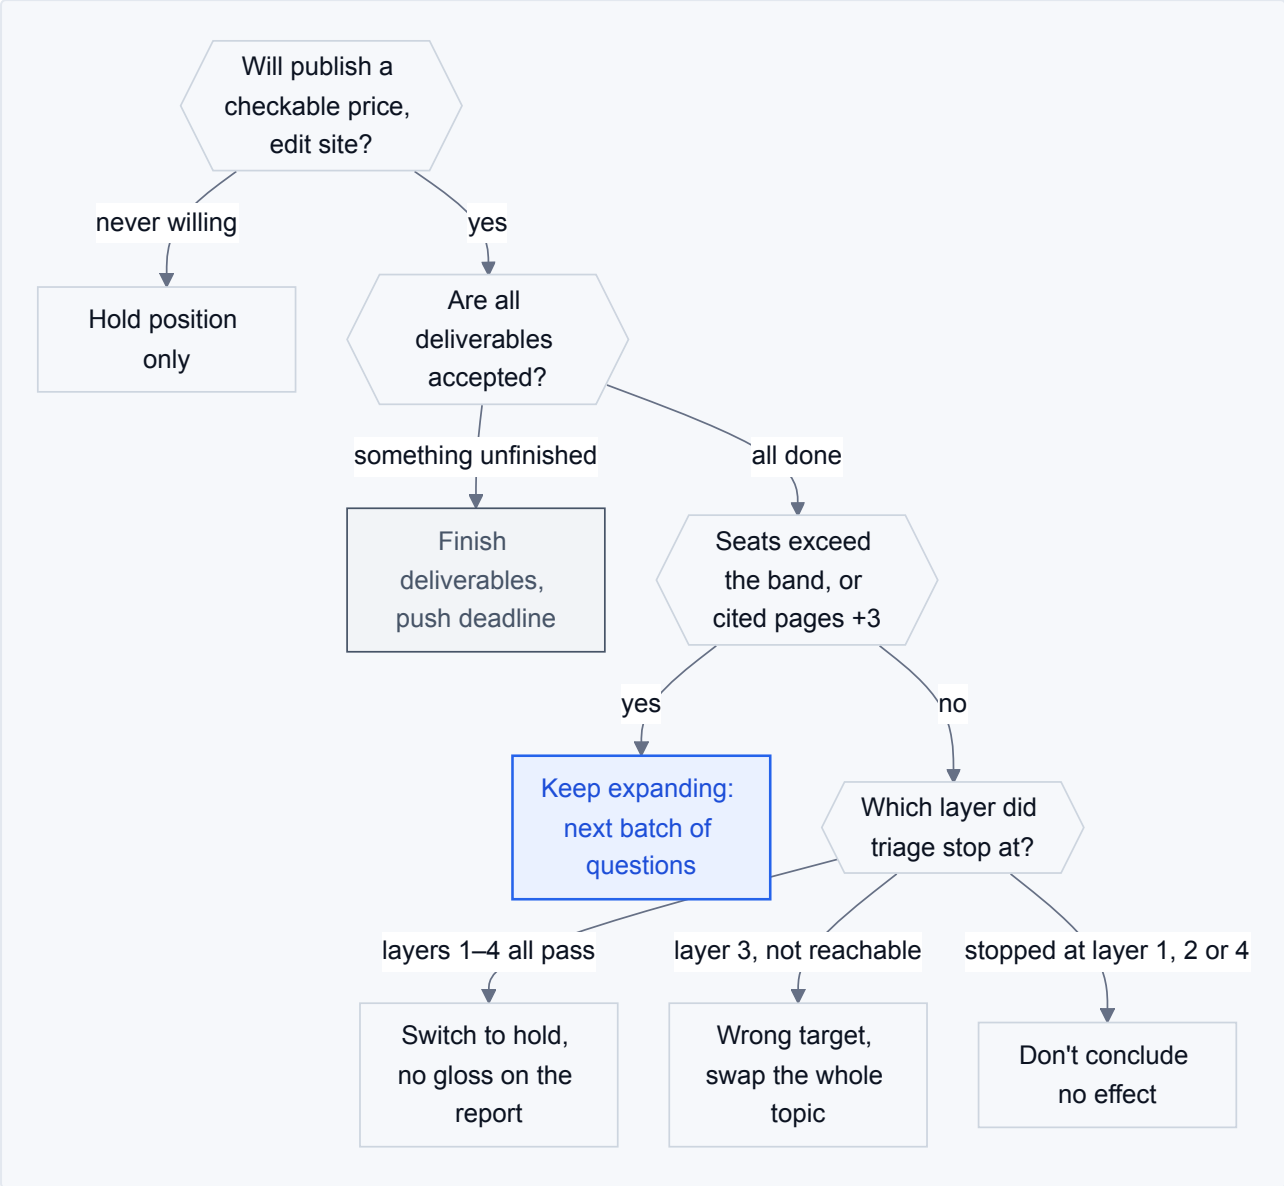


Figure: Settlement checks admission and deliverables first, then whether there's a determinable increase, and finally which layer triage stopped at; you may not conclude "no effect" without completing triage

The full statement of the criteria:

- **Keep expanding**: the seat increase > this quarter's measured noise band (**must use the same number of passes, same engine, same question pool**; if the number of passes was changed for any reason this quarter, this cell is judged "not comparable" across the board and may not be used — look at pages-on-the-list instead), or **on-the-list-and-cited pages increase by ≥3**. Keep expanding = move to the next batch of questions and redo target-picking and the shelf snapshot.
- **Switch to hold**: deliverables all passed, seat increase did not clear the noise band, on-the-list didn't increase either, and triage passed layers 1–4 all the way through. Stop adding new pages, and only do the five holding items above; write plainly in the report, "this quarter did not achieve a determinable seat increase on the main-attack questions" — **no glossing over it**.

- **Finish deliverables first**: **never offset a deliverables shortfall with a visibility number**.
- **Wrong target**: triage stopped at layer 3, "citation slots are all at unreachable sources" — move this whole intent out of the main-attack pool and actively change topic. **Don't count a targeting mistake as "this industry can't be done"**.
- **Admission condition**: the website must state a **definite, checkable price**, in the form set by the side (→ General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers's price figure); a client who's never willing to write one, or whose site can never be changed, gets hold position only. The price page is one of the few cells on an owned site that can actually get into an answer — **not giving a price is the same as giving up that cell**.
- Note **On the strictly regulated side, admission is judged by itemised fixed pricing**: on this side, writing a "price range" is itself non-compliant. Judging a client as "not giving a price" because it has no price range would wrongly disqualify a project that is actually compliant and simply can't write its price using the general template.

The decision rests with the data, not with hope. But "no effect" must be a conclusion reached only **after triage is complete and the seat change has cleared the noise band — not a side effect of the ruler's resolution being too coarse**. At settlement, these three sentences may not be said to yourself:

> × "Give it another three months and it will definitely rise" × "The industry average is six months" × "The data is trending up" (when the seat increase has not cleared the noise band and on-the-list hasn't increased either)

7.5 Ruler discipline: change the ruler and nothing is comparable

What you'll do in this section: pin down "what counts as changing the ruler" into a reference table, and pin down "whether the baseline should be re-frozen this month" into one mechanical test. When you're done, you'll have a way to handle a changed ruler, the rule for pinning the probe model, two ways to freeze the baseline when the shelf gets rewritten in bulk, and a list of which numbers can be printed and which stay internal only.

What counts as changing the ruler

Action	Consequence	What's allowed
Change engine	**The ruler has changed**	Start a new baseline, and write plainly "not comparable with the previous one"
Change the AI model (including a minor version)	**The ruler has changed**	Same as above

Action	Consequence	What's allowed
Change the phrasing / question pool / ratios	**The ruler has changed**	Same as above
Change the number of passes	Seat count **scales linearly**	Same as above, and the old and new must not be added together or shown on the same chart
Change the web control leg's 5 questions	**The calibration ruler has changed**	Do not change it for the whole quarter. If it must change, change it next quarter
Change the brand six's number of questions or rounds	Factual-errors count **before and after is not comparable**	Frozen for the whole quarter

Across versions, list the numbers only — do not compute a rise or fall. When numbers from two different baselines sit on the same page, they must be printed in two separate blocks, with a line "Below: a different ruler" between them. This table and the noise-band rule in 7.2 are two scales of the same test: the noise band governs sampling variation inside the same ruler; this table governs the moments when "you've actually already switched to a different ruler".

Probe models: pinned, never drifting with your main model

The model used to run probes **has a fixed version — it never drifts along with your day-to-day main model**. Name the legs as in 7.1; pin the model version per the table below:

Leg	Rule
① **OpenAI API leg**	**Must connect directly to the official API**. It measures ChatGPT itself — routing through any relay or proxy pool is measuring A to sign off on B
② **Gemini model leg**	May go through a proxy pool (when the free tier's direct quota isn't enough, the whole leg comes back empty — a "high-fidelity" path that can't get through is no more honest than a proxy path that can). **If you go through a pool, pin the exact model version**: a floating alias returns 503, and that whole month's sample is gone
Both legs	The version only ever moves on an explicit decision, and **every leg must move together**. Changing the version = changing the ruler, and it makes the new report incomparable with old reports already sent out

The cost has to be written into how the report is worded: when sampling through a proxy pool, region and the model's minor version can't be independently verified — print this sentence in the methodology note.

Whether to re-freeze the baseline this month: one mechanical test



Figure: Whether to re-freeze the baseline this month rests on one mechanical test only — never a gut feeling

Why this test is needed: put "the denominator doesn't change all quarter" (7.1) together with "the baseline isn't frozen during a window of concentrated shelf rewrites", and a gap opens up — the baseline freezes in some month, and afterwards most of the Top 20 pages get rewritten in bulk, so the denominator has already come apart from the shelf, yet discipline says you still can't change it. The test in the figure is what plugs that gap (not measured; the mechanism is clear).

Three points the figure can't show: "updated" only counts the **visible update date** on the URL, compared against the last freeze date; the old and new segments' X/20 **must not be drawn as one single line on a chart**; the monthly report for the re-freeze month **must state exactly which URLs triggered it**.

Two ways to freeze the baseline when the shelf has predictable rewrite windows

Some shelves have bulk-rewrite windows you can see coming: the same batch of sources gets rewritten around a particular point in time, and a rewrite = the candidate set gets swapped out (not

measured; the mechanism is clear). Freeze the baseline inside the window, or let the baseline and the retest fall into two different phases, and the ruler is effectively broken for that month — any rise or fall is entirely produced by the window, not by anything you did.

Option 1 · Avoid it	Option 2 · Record one per phase
Applies when: the window can be avoided	Applies when: it can't be avoided, e.g. the window covers half the year
Schedule both the baseline freeze and the retest outside the window	Measure once before the window and once after
If a measurement point lands on the window, move it earlier or later — never move the work schedule	Record two noise bands separately: normal-phase band / window band
Every comparison lands in the normal phase	After that, compare same phase to same phase only, never across phases

Figure: When you hit a bulk-rewrite window, pick one of two options and write it into the kickoff document — never choose on the day of the retest

Whichever option you choose, write the phase you're in at the time of freezing into the **first line of the baseline-freeze file**, and label the phase on every line of the monthly comparison; when numbers from different phases are shown side by side, state plainly "different phases, not comparable". Why this must be fixed in advance: getting this wrong doesn't throw an error — it just makes the whole month's report point the wrong way. Just like a misjudged door, the costs are asymmetric.

> **Example** (e-commerce) During a major sale, prices, stock status, platform ranking and buyer search terms all change at once, and both seat rulers' noise bands blow out. Normal phase = outside the two weeks before a major sale and outside the two weeks after one; for any month where the retest window and the baseline window fall in different promotional phases, always write "cannot be determined (cross-promotional-phase)" — you may write neither a rise nor a fall.

Different confidence, different printing

Low-confidence work still gets done — it just has to be **flagged when printed**, carrying the qualifying sentence exactly as given. Which of this chapter's statements get printed, and which stay internal only:

Statement	Printed or not	How to print it
This month's measured noise band	Printed	The value goes on page 1 of the monthly report; before the first measured value is obtained, seat counts are only listed side by side, with no verdict
Completing the five profiles benefits ChatGPT's local answers (from a public partnership)	Printed	State it on the basis of "free, done in passing" — never promise an effect
Improving position is cheaper than winning a new slot	Printed	Carry the qualifying sentence exactly: cross-industry pooled basis, not verified in this local industry
Fallbacks split by content form	Printed	Carry "does not include Singapore, about 75% confidence"
The rate at which a review is quoted word for word into an answer (external live search, a very small sample)	Not printed externally	Used only to decide direction of work; the monthly report prints your own measured values for the M questions, split into recommendation-type / non-recommendation-type lines, and the two lines must be shown together
The share of the 90-day increase that lands on brand questions	Not printed	A tabletop projection, not measured, internal reference only (the "do not say" line from holding position in 7.4)
An intent where "0 businesses are named, and someone is already fighting for the citation slot" (external live search, one-off)	The specific figure is not printed	Used only for internal ranking; the report prints the named-business count and citation-slot ownership from your own baseline numbers, with the retest date and engine stated

The reviews row requires the two lines side by side because printing only the recommendation-type line would lead people towards asking for reviews, which the strictly regulated side cannot do at all. The last row carries one more red line: × no material may ever say "no one is doing this" or "competition is near zero".

The percentage, sample and source behind each row are in → General Edition A.2 Evidence for picking targets, writing pages, off-site and retests; the general rule for how to flag a number and what stays for your eyes only is in → General Edition D.1 Number discipline: how to label numbers, and what stays internal.

Chapter 8

Using this in a new industry: four steps to work out your approach

Ten questions, how tightly you are regulated, ten factors and one decision table

8.1 Four steps for a new industry: ten questions → side check → ten axes → decision table

8.2 Five steps to decide the side, and three criteria

8.3 The ten axes (1): page types, containers, anchors, regulation, question shapes

8.4 The ten axes (2): decision-makers, capacity, attribution, half-life, language

8.5 The decision table and rulings

8.6 Seven tactics: how regulated and unregulated sides differ, and why each moves rankings

8.7 Cheap tactics on the unregulated side, and substitutes on the regulated sides

8.8 Unregulated ≠ easy: admission criteria J0–J7

8.9 Cross-border: whose law applies

8.1 Four steps for a new industry: ten questions → side check → ten axes → decision table

What you'll do in this section: work through the ten questions to get the raw input for the ten-axis table; then decide the side (8.2); fill in the ten-axis table (8.3–8.4); find the closest row in the decision table (8.5); and finally have this industry's ban list ready. Only when all four are in place may you open the *First work order*.

Get the starting point straight first: a one-line verdict like "this industry is hard to work" or "this industry is highly competitive" gives the people who execute nothing they can act on. What actually helps is to sort things into two groups: impressions that look like variables but change no action when they change, and real variables where a different value means rewriting tasks.

Looks like a variable, but isn't	What a real playbook must give you
Industry size / search volume	Measured share of page-type mix (axis 1)
How fierce the competition is	The four states, URL by URL, in the intent-level top 10 (axis 2)
Budget size	Entity anchor layer (axis 3)
B2B or B2C	Side + switches + substitutes (axis 4 + 8.7)
Whether you have brand awareness	The row in the decision table + the first 30-line work order

Figure: A one-line verdict gives you no playbook; a real playbook must give five things

Why each of the five items in the left column doesn't count: industry size / search volume changes no action — priorities are ranked by annual opportunity value, not search volume; how fierce the competition is is an adjective, and its actionable version already lives in axis 2 (how many reachable-enough items) and the abstain line J1; budget size changes how much you do, not how you do it; B2B or B2C is too coarse — two businesses that both wear the B2B label often differ more from each other than a B2B business differs from a B2C one, and what actually matters is axis 6 and axis 2; whether you have brand awareness doesn't change the order of actions, it only changes how high a starting point you get on the "get identified correctly" path (for the three paths, see → General Edition 1.3 Three gates and three paths).

The only test for a qualifying dimension (fixed): change one value, and at least 5 of the 30 lines in the *First work order* must be rewritten. If it can't force 5 lines to change, it isn't a dimension.

The order of the four steps cannot be shuffled — deciding the side comes before the ten axes: axis 4 (side and switches) and the decision table both use the side-check's conclusion, so deciding the side

first avoids referencing forward to something not yet worked out.

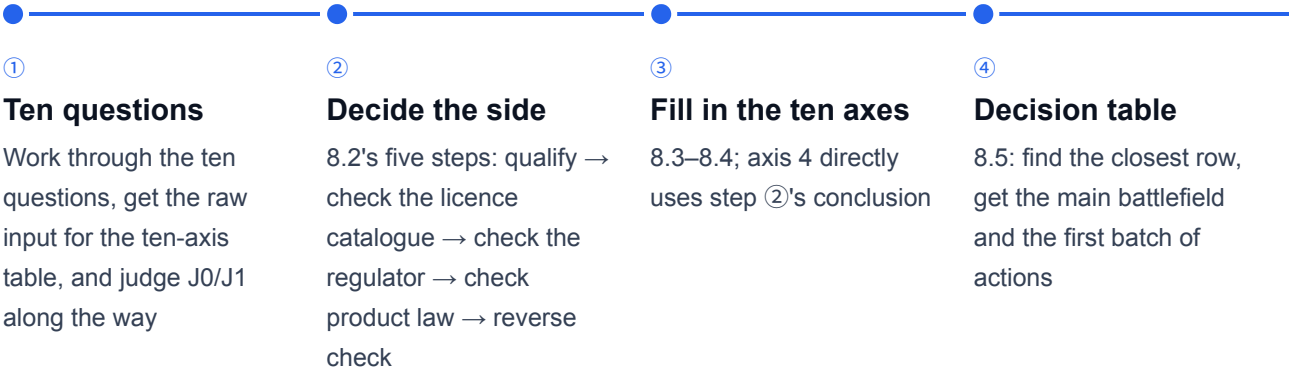


Figure: order for a new industry; decide the side before the ten axes

The ten questions: for each, where to look, what to check and which axis it fills

Take an industry you've never seen before: ask these ten questions and you can fill in the whole ten-axis table.

Question	Where to look, what to check	Which axis it fills
Q1 In this category's answers, does AI name any company? Which layer of noun does it name?	Gather 15 real questions (front-desk wording, customer-service transcripts, autocomplete suggestions from platform search boxes), run each on both engines; count whether any company is named at all, and check whether the subject of the naming sentence is an organisation, a person, a product or a course	Axis 3, axis 5; also judge J0
Q2 What page type do the citations for these 15 questions' answers come from?	Open every source URL one by one and hand-label the page type, judging only by the visible page, never by the domain's brand; work out the share	Axis 1
Q3 On the pages that get cited, can you get your own words onto them?	For every URL in the intent-level top 10, look for an entry point: a submission page, listing page, claim page, correction page, nomination page or seller back-office; if you can't find one, send an enquiry, and no reply within 7 days counts as not reachable; count the number of reachable-enough items	Axis 2; also judge J1
Q4 Of the pages that get cited, how many are "the container is not yours but the words are"?	Ask, one by one: "Am I the supplier of this page's content?" — your product page on a platform, your vendor profile on a review platform, an app store page, an integration marketplace listing, your own entry in an official directory, your speaker profile on a platform: all of these count as "yes". This is the question people most easily miss entirely	Axis 2
Q4b Does the platform itself let AI crawlers in?	For every platform you sell on, `curl robots.txt` and check the retrieval-leg crawler group; log it by date and rerun every quarter	Axis 2

Question	Where to look, what to check	Which axis it fills
Q5 Who regulates this business? Is there a written advertising regulation?	See 8.2's five steps	Axis 4
Q6 Which tactics get taken away? What replaces the ones that are taken away?	Go through 8.6 tactic by tactic and check each one against the rule text; for anything taken away, find a substitute in 8.7, and find an actual example of what that substitute looks like in this industry	Axis 4
Q6b Can someone else publish on your behalf? Who is the "authorised person"? Is the publishing party jointly liable?	Check this industry's rule text for three things: ① how "publisher" or "cause it to be published" is defined, and whether it names the publishing party, an agency, an affiliate or an influencer among the liable parties; ② whether there's a penalty clause for joint liability when publishing on someone's behalf; ③ whether there's an authorisation document you can actually sign (its name, who signs it, which media it covers). If even one of the three can't be found, the verdict is "cannot publish on the client's behalf" — not "can"	Axis 4 sub-item (form of agency liability, see 8.6 item 7)
Q7 How big is one order, how many orders can you take in a month at most, and what's the bottleneck?	Get the price list and the last 12 months' sales records; being unable to answer "what's the bottleneck" = no hard cap, switch the numerator	Axis 7
Q8 For the last three deals, who made the call, how many people signed off, and how many days did it take?	Get three specific deals; if cross-functional sign-off shows up even once, count it as a committee	Axis 6
Q9 How do buyers find you? Can the back office see the source?	Check whether there's a source field, whether sessions coming from AI carry a referrer, and whether you can add a required self-reported-source field; where the platform side can't show a referrer, just log "can only read the brand-term curve, for reference, not attribution"	Axis 8
Q10 How often does the price list change? How many buyers search in the local language?	Compare historical versions of the price list (don't take an estimate); run one pass of the local-language questions, compare the named sets and take the difference set — check whether the businesses in that difference set all have a local-language page	Axis 9, axis 10

Red line **Q6b must be answered within the ten questions — don't let it surface only after work has started**: when the answer is "cannot publish on the client's behalf", the whole off-site block of work (bylined contributions, outreach letters, submissions, filling in platform back-office fields) can only be

carried out, in this industry, by the licensee itself; putting someone else on it in the work order still won't get it executed.

Only when all four are in place may you open the *First work order*: one line of the *Side receipt* (8.2, written at the top of the work order), a fully filled ten-axis table (with one sentence per axis on "which point of leverage it lands on"), the matching row in the decision table, and this industry's ban list (for the strict and light sides) or platform-policy table (for the unregulated side).

8.2 Five steps to decide the side, and three criteria

What you'll do in this section: when 0.2's four questions don't give a firm answer, work through the full five steps — three questions to qualify, check the licence catalogue, check the regulator's three kinds of document, check product law, and reverse-check if you find nothing; steps 3 and 4 are each capped at 1 working day, and if either overruns or cannot be confirmed, always start work on the stricter side; then use three criteria to decide switch A (agency liability) and the cost of judging wrong; the side decision has an expiry date and is rechecked every quarter.

0.2 gives you the quick version: four questions that give you the side directly (strictly regulated / lightly regulated / unregulated) and switch E (how many tactics get taken away). For most industries that's enough. This section is the full version: five steps to gather all the documents you should read, use those documents to verify 0.2's four questions properly, and then use three criteria to work out two more things — switch A (who can publish on your behalf, and whether they're jointly liable) and the cost of judging wrong. The two sections use the same terms: the three sides plus the six switches, and no other tier gets invented. In this book, "regulated sides" means the strict side and the light side combined.

How to split the work: steps 1 and 2 can be finished in the same meeting (ask three questions + search the licence catalogue once); steps 3 and 4 don't belong in that meeting — reading the rule text for penalty and joint-liability clauses, and checking product law category by category for what's sold, are not things one meeting can finish; force them through in one sitting and what you get is a guess.

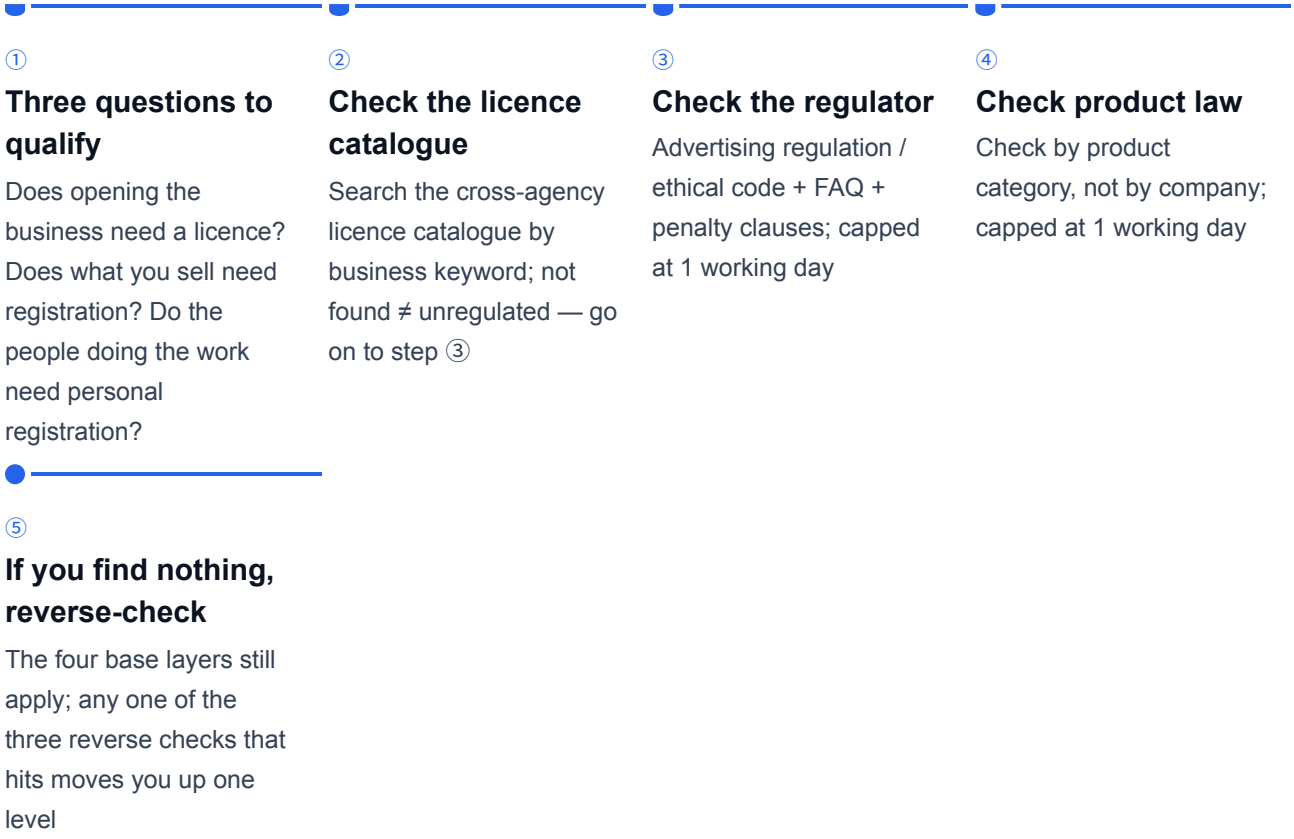


Figure: five steps to decide the side; steps 3 and 4 are each capped at 1 working day

Time limit and escalation (fixed): if step 3 or 4 overruns, or you checked but cannot confirm it, start work on the stricter side — never let it enter the work order marked "TBD"; leaving it as TBD means you simply can't schedule the first batch of tasks.

If all three questions in step 1 come back no, provisionally judge unregulated, but the verdict isn't complete until you've worked through step 5's three reverse checks. Not finding anything in step 2 doesn't mean unregulated: the licence catalogue is indexed by licence name, not by a business's common name, so not finding it means moving on to step 3 — never judge unregulated directly.

One addition to step 1: check the statutory hard definition first. Many rules mark out who's regulated with a string of conditions, and the conditions must **all hold at the same time** to count as a hit — miss one and it's not a hit.

> **Example** (education and training): one category's definition of "full-time" requires all three of "duration ≥1 month + 5 days a week + ≥3 hours a day" to hold at once before it counts as a hit; a tuition class held two hours a week doesn't hit.

Step 3: check the regulator's three kinds of document

General-pattern sentence: **every regulated industry has three things — a regulator, a dedicated advertising regulation, and an official register lookup — fill in all three before you decide the side**; the specific agency names and rule names are an industry variable and are left to the industry

edition. On the regulator's own website, check three kinds of document: the advertising regulation / ethical code (including its section on promotion), the agency's FAQ (its stated position is often stricter than the rule text itself), and the penalty and joint-liability clauses.

Step 4: check product law to decide switch C (regulated products) and switch D (spans two sides)

In some industries the company itself needs no licence, but **what it sells** is governed by a dedicated regulation — the copy is regulated all the same, and that's switch C (regulated products). Decide the side by the product category you sell, not by the company; if any cluster of product categories or business lines under your name hits, run the whole site to the strictest cluster — that's switch D (spans two sides).

> **Example** (e-commerce): once SKUs such as supplements, medical devices or infant food hit a dedicated regulation, the whole site must follow the strictest requirement in that cluster, even if the company itself needs no industry licence at all — you can't decide the side by company; AI treats one domain as one entity, so splitting subdomains into stricter and looser sides just breaks the entity apart.

Step 5: what to do when you find nothing

Judge it as the **unregulated side**, but the base layer isn't "no rules".

Base layer	Hard constraint
Fair trading / consumer protection law	No misleading statements; false reviews and misleading web design have already led to enforcement cases
Price transparency guidelines	Covers drip pricing, strikethrough prices and discounts, "free" claims; mandatory fees must be included in the stated price
Personal data and the Do Not Call Registry	Sending a message needs either clear consent, or the register checked within the last 30 days; **B2B messages are exempt from this rule** — a real loosening
Platform terms	Often stricter and faster to enforce than the law (store suspension, review removal, ranking demotion)

Before you judge unregulated, you must run three reverse checks — any one that hits moves the case up one level:

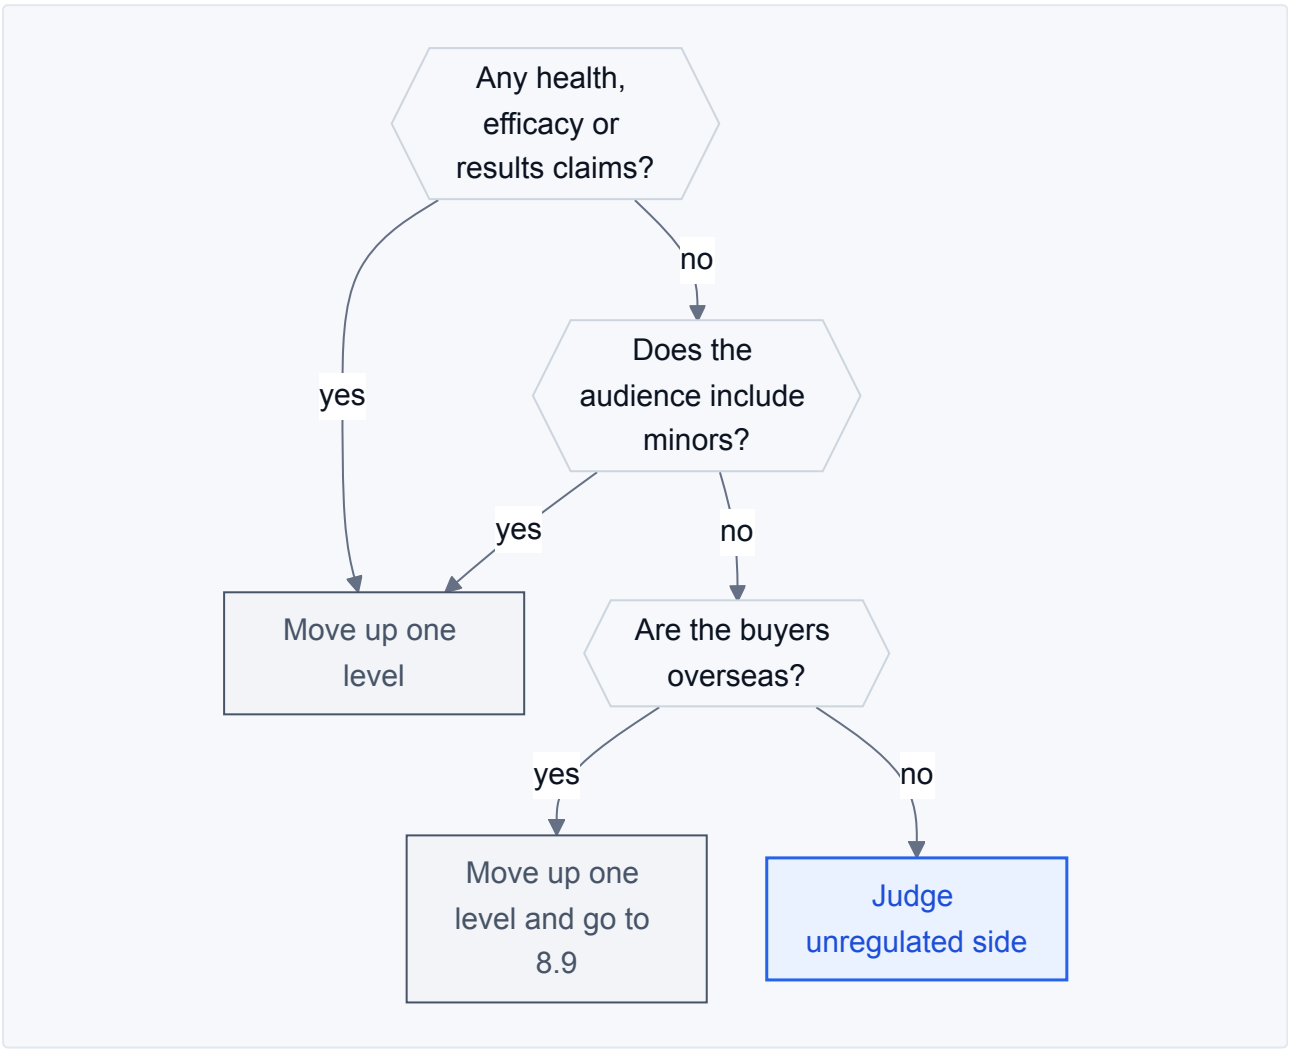


Figure: three reverse checks before you judge unregulated

The claim categories in the first question, in full, are six: health, efficacy, earnings, results, employment and salary; any one of these appearing in the copy counts as a hit.

Three criteria decide switch A (agency liability) and the cost of judging wrong

Run the documents found in step 3 through three independent criteria. First fix what counts as a "hit" for each one, then look at the combination of hits — don't judge by "all hit / not all hit".

Criterion	What counts as a hit	How to check
① Written advertising rules	There's a dedicated advertising regulation, or the ethical code / industry code has a section that actually binds outward-facing promotion	Search the document for advertisement / advertising / publicity / promotion, or their Chinese equivalents; if you can't find a binding clause, it doesn't hit, however thick the document is — a document existing isn't the same as a hit
② Penalties	The rule text states a fine amount or a prison term	Look for the penalties section or a "penalty" clause

Criterion	What counts as a hit	How to check
③ Agency liability	The rule text names the publishing party, an agency, an affiliate or the publisher among the liable parties	Look for wording like issue or publish / cause to publish / authorised person, or the Chinese equivalent of "cause to publish"

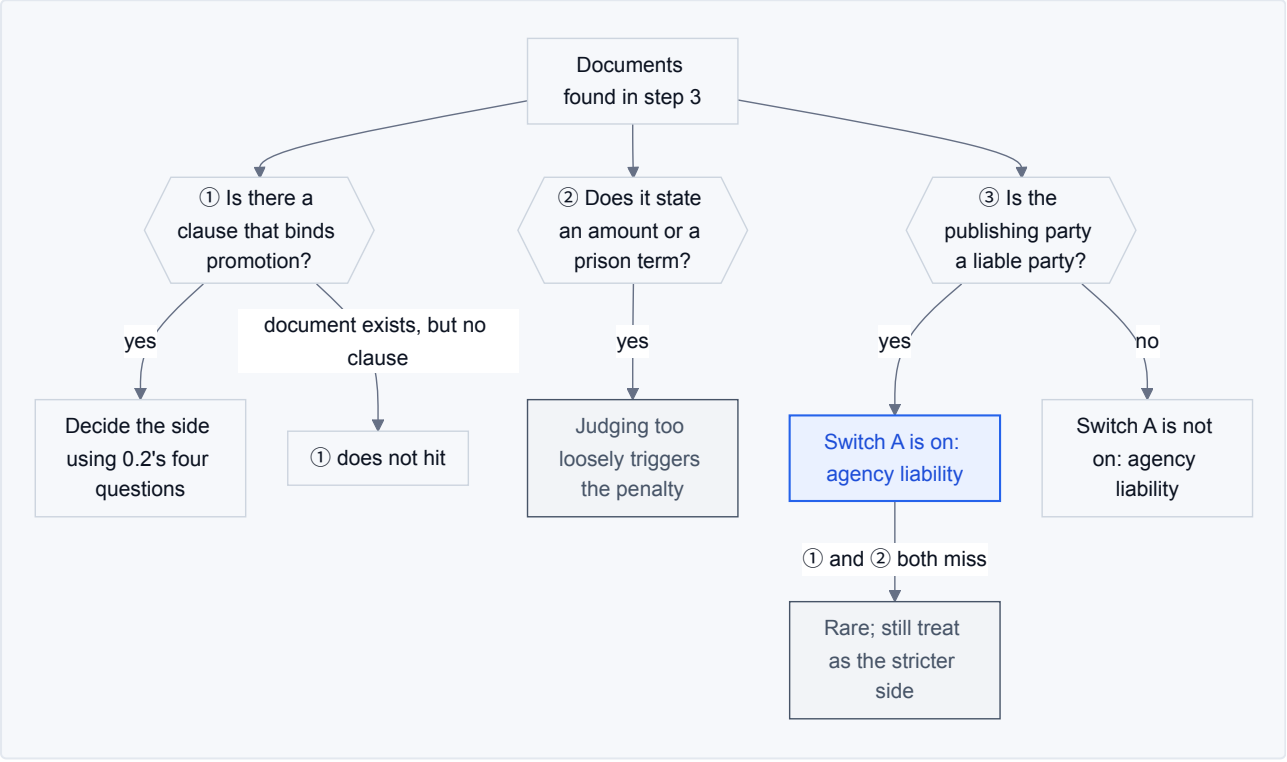


Figure: each criterion decides one thing: ① gives the four questions something to answer against, ② sets the cost of judging too loosely, ③ alone decides switch A

Hitting only ③ (agency liability exists but there's no advertising regulation) is a rare combination, but always treat it as the stricter side: ③ decides whether this approach can be carried out at all, and the cost of judging it wrong is the most asymmetric of all the combinations. For which form of agency publishing applies once switch A (agency liability) is on, and what hard consequences each one carries, see 8.6 item 7. Check switch B (legally required fields) and switch F (referral commissions banned) off the same batch of documents while you're reading them; the conditions for hitting them are in the switch table in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers.

The cost of judging wrong is asymmetric, and that's the only reason behind the rule "when unsure, always treat it as the stricter side".

Judged too loosely (should be strict, judged loose)	Judged too strictly (should be loose, judged strict)
May trigger penalties and joint liability	You only give up a few tactics

Figure: the cost of judging wrong is asymmetric; when unsure, always treat it as the stricter side

On the unregulated side, check these three places instead

On the unregulated side there's no regulator to check, so check three other places instead; what you're guarding against also changes to three "pass-through red lines" — the only pocket of strict regulation inside an otherwise unregulated industry.

On regulated sides, check these three kinds of document	On the unregulated side, check these three places instead
Advertising regulation or ethical code	Your target review platforms' review-solicitation policy
The regulator's FAQ	Your target app and integration marketplaces' listing policy
Penalty and joint-liability clauses	Your jurisdiction's consumer-protection and false-comparative-advertising clauses

Figure: with no regulator to check on the unregulated side, the three kinds of document become three places to check instead

Pass-through red line	When it triggers	What to do
A certification word is a legal claim, not an adjective	The page carries a security, privacy or payment compliance-certification word	Next to every such word, spell out the form of proof (type, how long the report stays valid, who issued it, whether it can be obtained under an NDA) and the last check date; writing "compliant" when the reality is only "aligned" is actionable
A buyer's own regulation passes through onto your claims	The product is sold to buyers in a regulated industry or to government buyers	Any wording that touches the buyer's industry regulation must be checked by legal, and written to the strict side's standard
Platform terms bind harder than the law in practice	Active on review sites, app stores, integration marketplaces or package registries	Read each platform's review-solicitation policy / listing policy clause by clause before you act; delisting or profile freezes arrive faster than regulatory risk does

Three edge cases (fixed)

- **One company spans two sides**: treat the whole site by the strictest side; never split strict and loose by subdomain — AI treats one domain as one entity, and splitting it apart just breaks up the entity's weight.
- **A group with multiple legal entities**: decide the side by whichever entity is actually named, and treat the rest of the group's entities by the strictest one.
- **An agency or distributor publishing on the brand's behalf**: decide who's the "authorised person" first; where it's unclear, treat it as the brand's own side — unclear liability = cannot publish = this approach does not exist.

Side decisions expire: recheck every quarter

Red line **In an industry the regulator is actively reviewing, the side may move up within the current year.** Recheck the regulator's website for the three kinds of document every quarter — don't treat a single side decision as a permanent verdict.

Acceptance: the output is one line of the *Side receipt* plus a *To-verify list* (each item records three things: what's to be verified, what it's stuck on — the register page was redesigned / the rule text is only a PDF / the agency hasn't published it — and the next recheck date), plus one more check of the register screenshots; the first time you actually use a register lookup entry marked "to be verified", it must be verified on the spot and filled back in.

8.3 The ten axes (1): page types, containers, anchors, regulation, question shapes

What you'll do in this section: use seven questions to label the baseline Top 20 by page type and work out the dominant page type (axis 1); use two questions to decide the control of citation containers for each URL (axis 2); read real answers to fix the entity anchor layer (axis 3); the regulation axis uses 8.2's conclusion directly (axis 4); and pick your battle tactics by the four question shapes (axis 5).

Axis 1 · Shelf page-type mix

Open every item in the baseline Top 20 (each intent-level top 10 counted once) and hand-label the page type — one label per item, never guess, judging only by that page's layout, never by the domain's brand. Work through the seven questions in order; stop at the first hit.

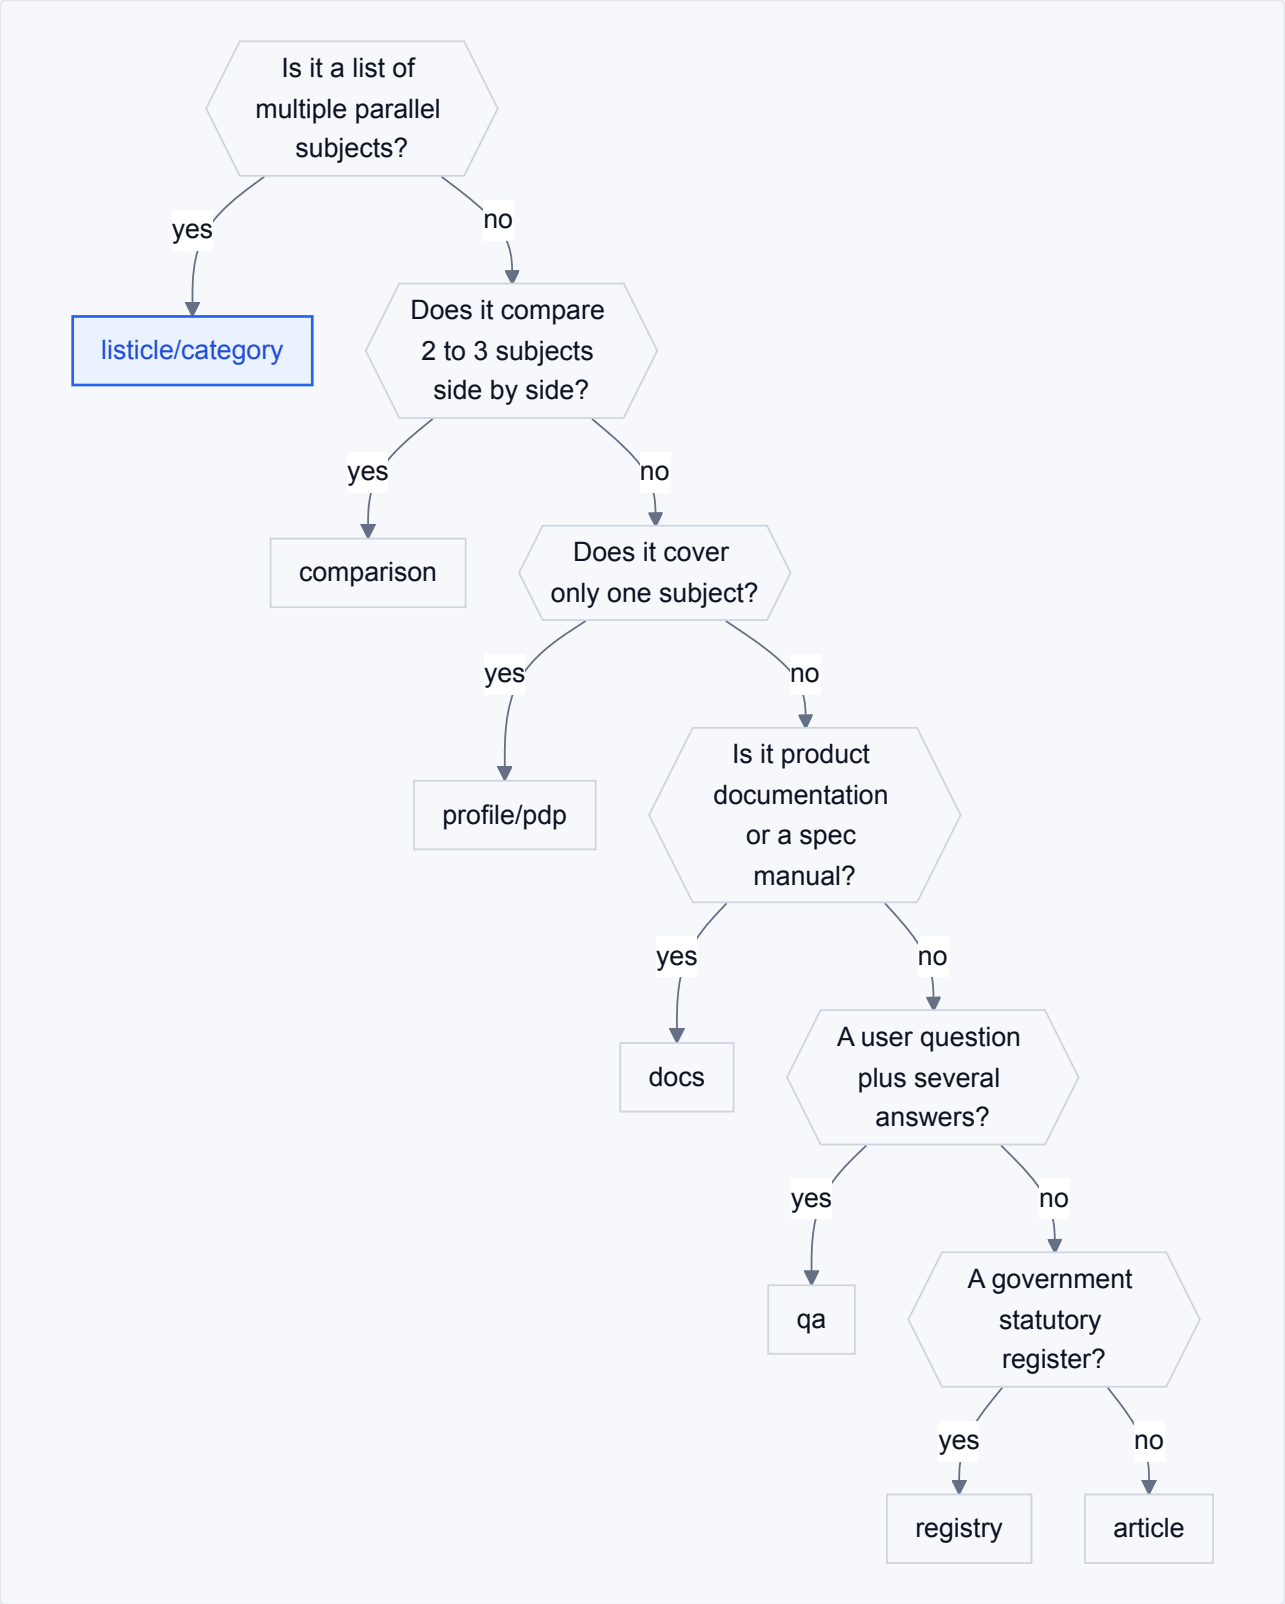


Figure: label the page type with seven questions in order; stop at the first hit

Three things the figure can't show: ① after hitting "listicle/category", ask one more question — "does it use ranking or better/worse wording?" — yes → listicle, no → category; ② after hitting "profile/pdp", ask "can you place an order?" — yes → pdp, otherwise → profile; ③ two people each label independently; where they disagree, don't negotiate a compromise — take whichever label comes from the lower-numbered question (the order of the seven questions is itself the priority), and anything still unresolved

is logged as unclassified, listed in its own column and not spread into any label's share; if the two people's disagreement rate is >20% for that round, the whole round is void and must be relabelled — never use it directly to decide where effort goes.

Dominant page type (share threshold)	Where effort goes
article ≥40%	Push earned-media bylined contributions
listicle ≥30%	First judge whether you can touch the list (axis 4); if you can, fight for an editorial slot; if you can't and you are listed, request removal
category + profile ≥30%	Directories move from a logging action to the main effort
comparison ≥15%	Build your own comparison page / price guide page, eating the same density
registry ≥30%	The action is checking your records are correct, not submitting
qa ≥20%	Go to axis 2 and judge whether the community is platform-owned content or reachable
pdp ≥30%	Go to axis 2 and judge C or D; answer the "words" question and the "door" question separately
docs ≥20%	Build your own documentation site + integration and open-source directories

Figure: the dominant page type decides where effort goes

Write the shares as `article x% / listicle y% / comparison z% / …`; you can hit several thresholds at once, and order effort by share, highest first. comparison has the highest per-page citation density of any type in the table (1.87 citations per search), and it's counted by page type, not by who owns the domain — a comparison page you build yourself eats that same density. pdp and docs are also judged by layout: a media outlet's long buying-guide article is an article, and a brand's own product page on its own site is a pdp, regardless of whose domain it sits under.

Page-type distribution is an industry variable, not a general rule: you must label the page types yourself; never extrapolate from another industry.

> **Example** (dental and aesthetics): in the medical sample, article accounts for 54% of citations and listicle 0% — lists simply aren't the shelf in that industry; directories and lists are only a logging action. In a different industry these two numbers could well be reversed (for the sample, see → General Edition A.1 Evidence for mechanics, the door and identity).

Axis 2 · Control of citation containers

Judge who owns the container URL by URL (not domain by domain) in the intent-level top 10. Ask two questions in order: are the words on this page mine, or can I submit them (does the page have a claim, submission, correction, nomination or seller back-office entry point)? And who owns the container? The basis for the verdict must be one of: a screenshot of the entry-point URL, an email receipt, or the location of an explicit statement on the page.

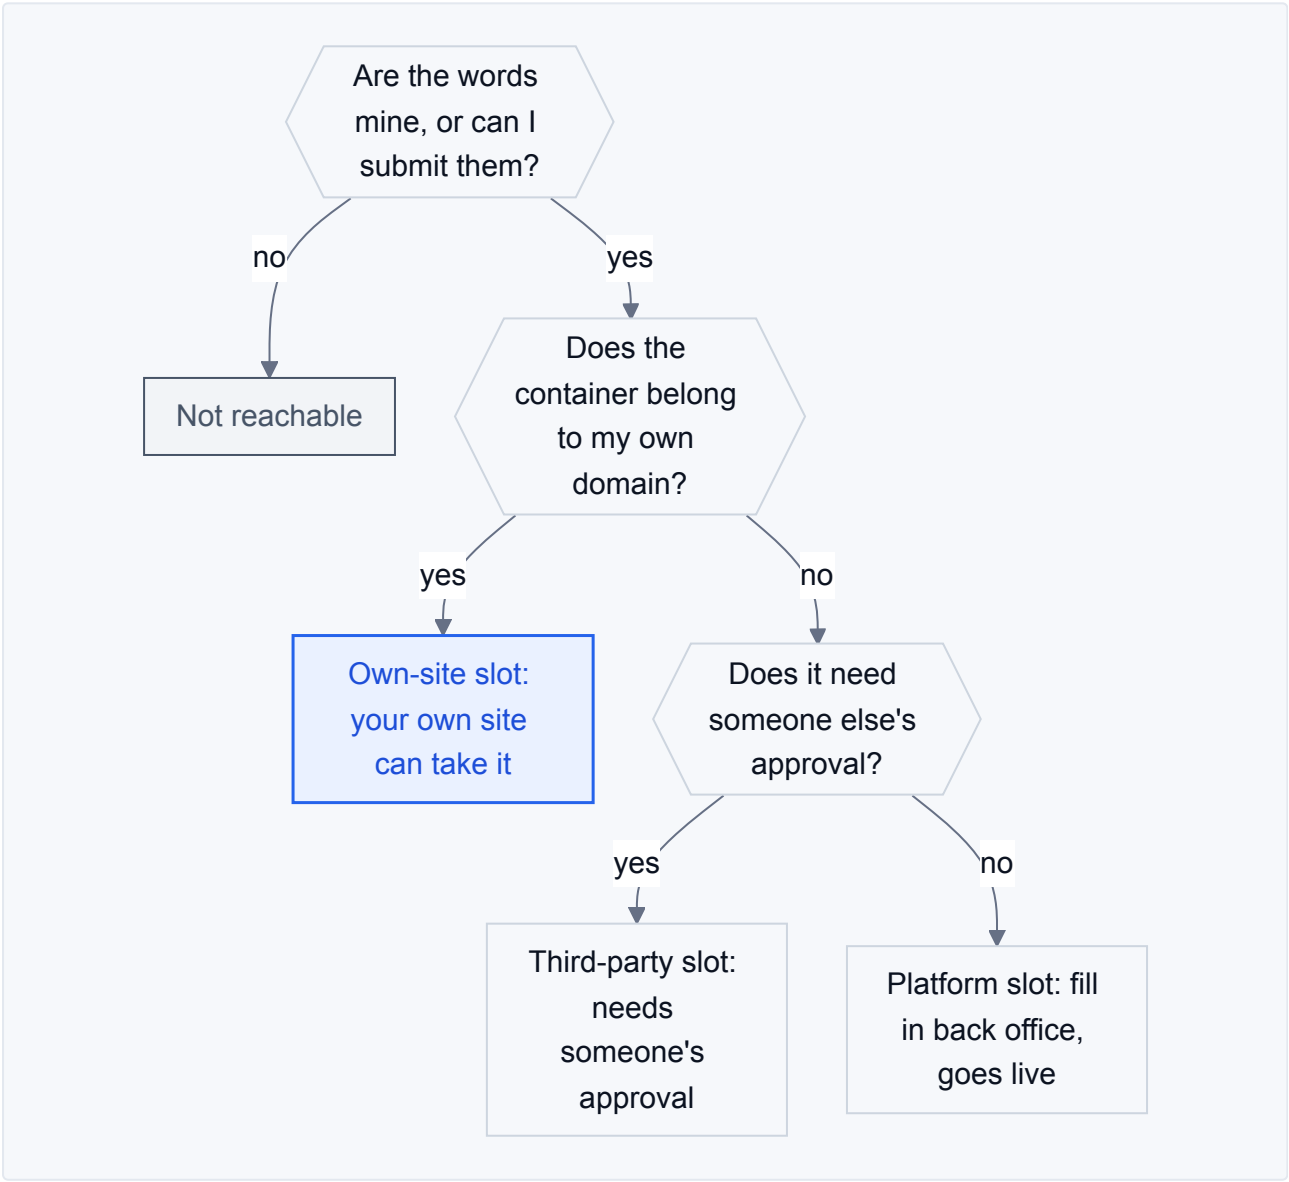


Figure: two questions to decide who owns the container

The own-site slot has one more hard test: this page type is either already present, or can be built this quarter, on my own domain — if you can't build it, it doesn't count as an own-site slot. The line between third-party slot and platform slot is whether there's human editorial review: editorial slots, submissions, listings and nominations need someone else's approval, so they are third-party slots; if what you fill in the back office goes live straight away, it is a platform slot.

The four containers here (own-site slot / third-party slot / platform slot / not reachable) and the four states in → General Edition 4.6 Four states, two denominators and the abstain line are two different things: the four states judge "can I get in", the four containers judge "who owns the container". How they map: own-site slot, third-party slot and platform slot all fall under the four states' "already present" or "reachable"; "not reachable" is simply the four states' "not reachable"; "to ask" isn't a fifth value — it's the state of a third-party slot where asking has not yet produced an answer. The count of "reachable enough" items is counted the way 4.6 defines it; neither not-reachable nor to-ask counts, and this is the number J1 uses.

Industry-level aggregate threshold (this is exactly what the decision table uses — never use an adjective instead): if a value appears in ≥5 items of the intent-level top 10, log it as "many / high"; ≥7 items, log "very high / saturated"; ≤2 items, log "few".

Which cell has more	Where it points
More own-site slots	The main battlefield is your own-built pages
More third-party slots	Put people on bylined contributions and listings
More platform slots	The cell most often missed entirely: you are the owner of that page's content — changing the words costs almost nothing, goes live almost instantly, and there is no 2–6 week outreach cycle
More not-reachable	Count the reachable-enough items; if that hits J1, rule it out

Three fixed rules: ① The definition of "not reachable" is "you cannot get your own words onto that page", not "this domain isn't yours" — a large share of "platform pages" ruled dead are actually C, not D; getting this wrong once means throwing away the whole main battlefield. ② The reverse also holds: it's only D when the platform itself has shut the door, and this must be judged retrieval-leg by retrieval-leg — never judge one overall state for a platform. ③ Measure platform by platform: of two equally large platforms, one may be half-blocked and the other fully open; you must `curl -L` it yourself, log it crawler-identity by crawler-identity, log it by date, and rerun it every quarter — what gets logged is "which leg is C, which leg is D", not "is this platform open"; a "to ask" that gets no reply within 7 days converts to D — never leave it hanging as a middle state.

> Red line **Fixed**: if the government writes about you, it's an official register — you can only check it for accuracy; if you write about yourself, and the government or the platform merely supplies the container, that's a C-type platform — judging it as D throws away the cheapest, fastest-acting cell entirely.

Axis 3 · Entity anchor layer

Read 5 real answers and look at what noun is the subject of the naming sentence; then open three third-party pages from the baseline Top 20 and see which layer their entry titles use.

Value	Where it points
Organisation	The anchor is the organisation facts page `/facts` (see → General Edition 3.2 The fact table and the organisation facts page)
Natural person	One page per person, `/team/<name>`, with every thick page, video and off-site material's byline linking back
Product / SKU	The anchor is the product page, not the company page; `sameAs` strings together product-level identifiers (version number, package name, GTIN, bundle id); third-party profiles are built by product, not by company
Parent brand	The sub-brand and the parent brand interlink and match word for word in the visible HTML
Course / cohort	The anchor is the course; the instructor is the second anchor

You can have dual anchors (the parent brand feeding the answer shelf + the SKU feeding the shopping shelf), or split by sub-trade. **The real mechanism behind entity strength isn't what the schema says — it's the same set of facts appearing word for word across several independent third-party pages that can be crawled.** Pick the wrong anchor layer and you've aligned the wrong object word for word — wasted work. Fill in the word-for-word alignment table for the six public traces as usual; only the six places change to this industry's own six.

The actions in the "identity" layer are the same across all three sides, with not one step dropped: unique spelling → `/facts` page → six-trace word-for-word alignment → send a correction the same day for any mismatch → write `sameAs` while you're at it. Only one thing changes — where the strongest anchor comes from, which depends on whether there's an official register.

Regulated sides: free for the taking	Unregulated side: you build it yourself
Official registration number + link to the official register lookup page	A general business-registration number (not an industry licence)
Write it in the first line of `/facts`	Leaving the registration-number line blank is a mistake — swap in an equivalent unique string instead
At least two of the six traces are official pages	A unique package name in a package registry or app store + your vendor-profile ID on a review platform
The anchor exists from day one	The five business profiles must match word for word; fixing the entity often comes before adding pages

Figure: whether there's an official register decides whether the strongest anchor is free for the taking, or something you have to build

Axis 4 · Side and switches

This axis uses 8.2's side decision directly: the side, plus whichever switches hit — the ones that most often change the actions are switch A (form of agency liability, see 8.6 item 7), switch C (regulated products), switch D (spans two sides: run the whole site to the strictest side) and switch E (peer comparison banned). This axis produces no rankings itself; it decides which of the ranking-producing actions you're not allowed to do. The "door" layer doesn't change by a single word: the robots.txt master template, the three `nosnippet` controls, matching the four identities, the CSR check — these are identical, word for word, on the strict side and the unregulated side; for the "identity" layer, see axis 3 above. **Regulation takes tactics away; it doesn't change the building site, but it does change what goes into the "identity" cell.**

Axis 5 · Question-shape mix

Run each of the 15 frozen questions on both engines, read the shape of each answer and count it, and write it as `recommendation x / spec y / comparison z / approach w`.

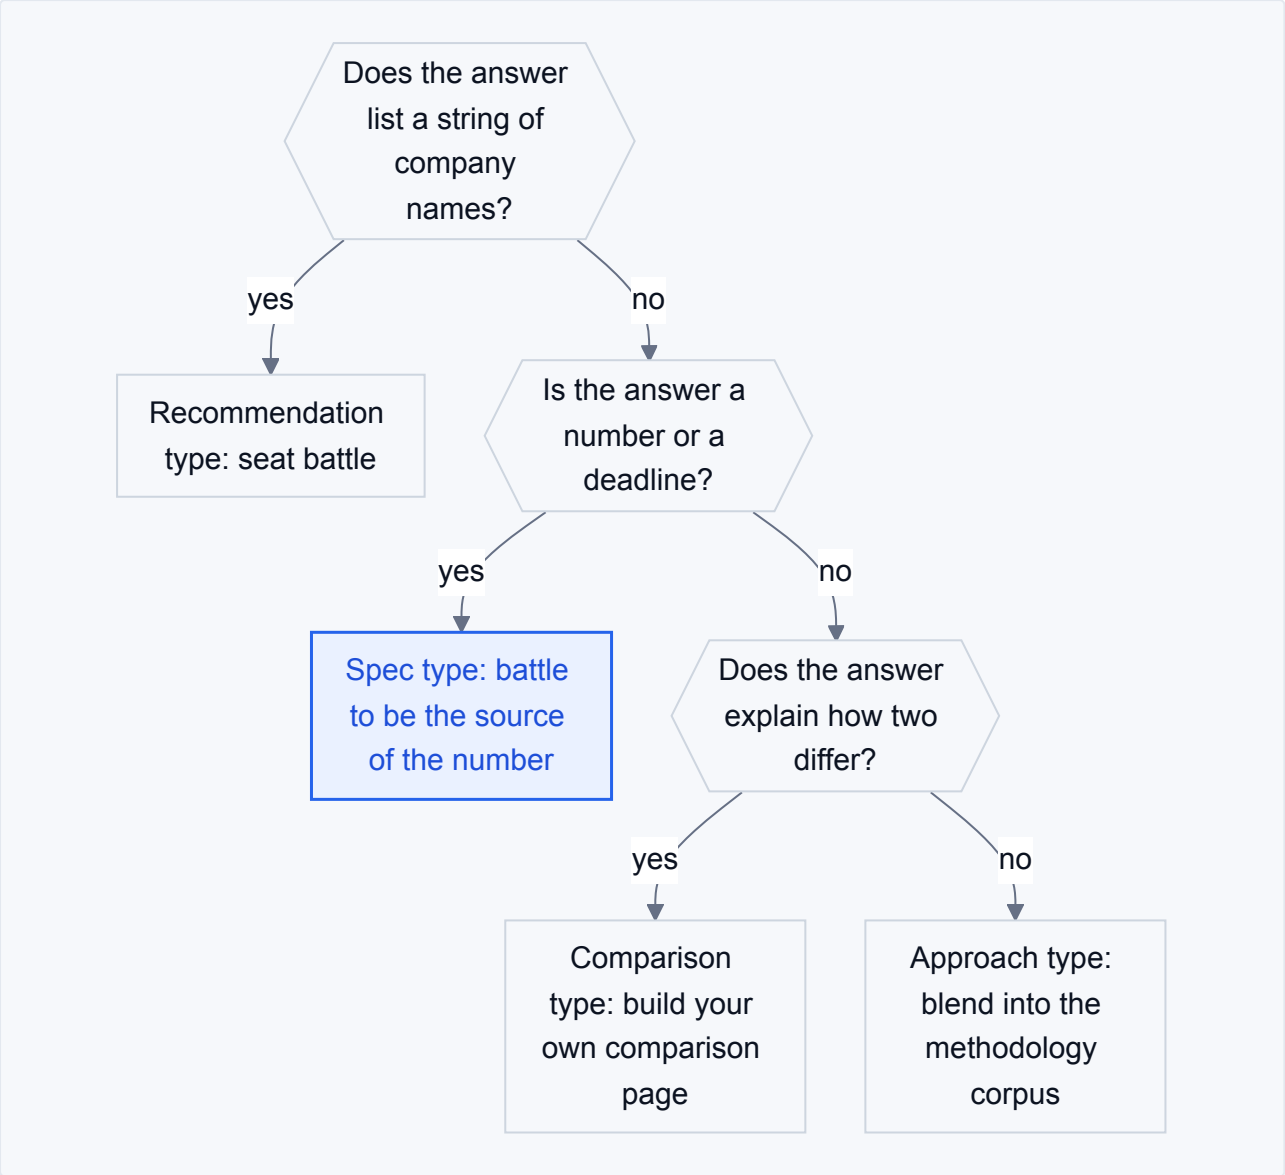


Figure: four question shapes, four battle tactics; use the tactics for the wrong shape and even a well-built page won't get in

The recommendation type fights over the 5–8 names in the answer; the main effort is off-site — planting your words into pages already cited, plus building your own list-style landing page (unregulated side only; the strict side does not self-build list pages, and neither does the light side — Conservative line (not statute text); on the light side the client's written reasons can relax it; see → General Edition 0.3 Which page types are limited for you, and what to build instead and 8.6). The spec type's target isn't getting listed — it's becoming the source of that number: a visible table + a stated basis and check date + ≥3 outbound links to checkable original sources (on the strict side, outbound links from your own website never go to commercial companies; linking only to government bodies and regulators is a Conservative line (not statute text), see → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions). **Price and cost queries trigger AI Overviews >80% of the time, and mixed-intent queries like "roughly how much does XX cost" 97% of the time** — the richest cell, shared by every industry.

The approach type is the hardest, and the one most easily given up on: your company name only appears as an example, and the only lever you have is getting your name into the body text of methodology someone else has written (bylined contributions, being cited on documentation sites, open-source and integration directories, long community answers). The basis: among the citation sources of approach-type answers, article has the highest share, and 84% of those article citations are earned media — label it "not measured" and do it anyway. Apply recommendation-type tactics to an approach-type category, and however well the page is built, it still won't get into the candidate set.

8.4 The ten axes (2): decision-makers, capacity, attribution, half-life, language

What you'll do in this section: take the last three deals and pin down who made first contact, who made the call, how many people signed off, and how many days to payment (axis 6); pin down the monthly capacity cap and pick the right annual-opportunity-value formula (axis 7); check whether the source of a sale can even be seen (axis 8); see how many times a year the price list changes and set the half-life level (axis 9); and compare the local-language named set to decide whether to build a local-language page (axis 10).

Axis 6 · Number of decision-makers and decision cycle

Ask about the **last three deals**: who made first contact, who made the call, how many people signed off, and how many days from first contact to payment — get three specific deals, not an average. Values: single, on the spot / single, over weeks / two people / two people + a third-party user / committee (3+ people across functions).

Single → one intent cluster, one page, no sections inside the page; **committee → the definition of a page unit must change**: a page is still one intent cluster, but is split into H2 sections by role inside the page, each section carrying its own checkable facts and sources; two people + a third-party user → a fixed two sections inside the page (a checkable-facts section for the buyer + one section for the user); high-ticket durable goods with no functional division of labour → split into H2 sections by "who suits which option". **AI picks up sections, not pages**: committee-purchase questions are already asked role by role, so multiple sections on the same page means the same URL has copyable body text against several intents at once (not measured).

Axis 7 · Shape of the capacity cap

Ask one question: "How many orders can this line take in a month at most, and what's the bottleneck?" Being unable to answer means there's no physical hard cap — switch the numerator. Values: headcount or machine count / professional's working hours / capacity × how often a cohort runs / inventory and supply volume / no hard cap.

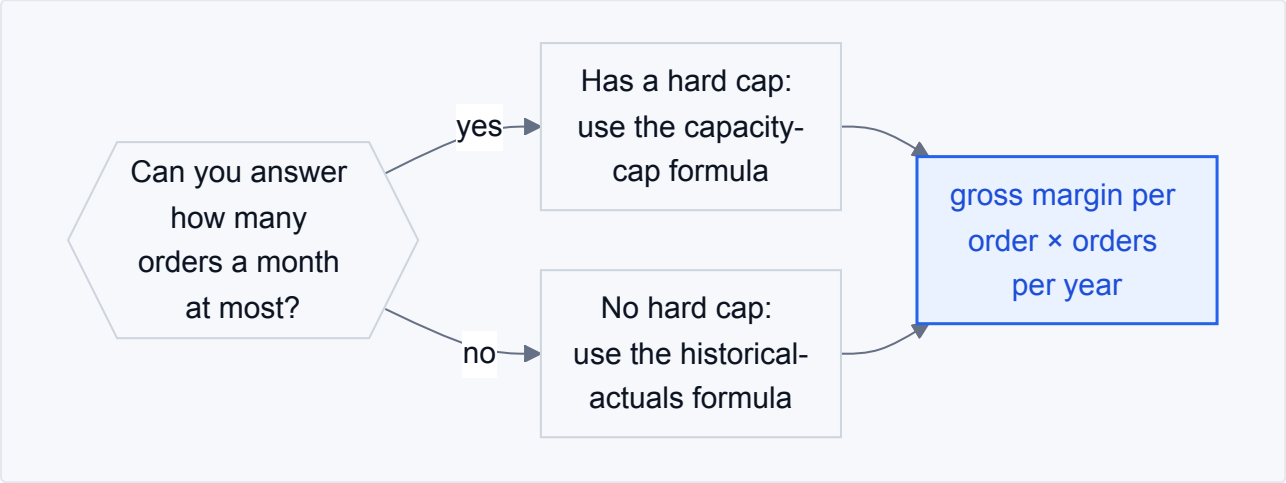


Figure: whether the question can be answered decides which annual-opportunity-value formula to use; the two formulas share the same shape

Both formulas must share the same shape: the annual opportunity value of one work unit = gross margin per order × orders per year; "orders per year" must be something you can look up from the back office, a rota, a course schedule, a stock table or a sales ledger — never estimated on the spot. With a hard cap: `price per order × gross margin × monthly capacity cap × 12`, where the monthly capacity cap comes from the rota or course schedule; for goods sold, write it as `gross margin per unit × monthly supply volume × 12`, where monthly supply volume comes from purchasing or restocking records — that number is itself the physical hard cap for this cluster of product categories. With no hard cap: don't ask "how many orders a month at most" again — switch to historical actuals instead: `first-year contract value × gross margin × expected retention years × this work unit's actual number of closed deals in the last 12 months`, with no "cap" of any kind appearing in the numerator; a new work unit with no history takes the actuals of the closest matching work unit at the same business, logged in the ledger as "proxy". **Two bans**: never mix "cap" and "actuals" into the same formula; and never assign a default value just because one factor can't be found — if you can't find it, mark the whole thing "cannot be calculated" and leave it out of the ranking. **Annual opportunity value is only for ranking — not one word of it may be called a revenue forecast, and it must never be multiplied or divided by seat counts, or shown side by side with them**; changing the numerator must not change the ranking rule (see → General Edition 4.5 Buyer types, annual opportunity value and the three checks).

The work unit varies by industry: one type of customer / one product or course line / one use case × one role / one category cluster. The ledger, landing pages and acceptance are all counted by number of work units, never by number of pages.

Axis 8 · Attribution channels

Check three things in the back office: is there a source field? Do sessions coming from AI carry a referrer? Can you insert a required self-reported-source field? Values: form / phone / WhatsApp / in-store visit / free-trial sign-up / add-to-cart order / app install.

Counts towards acceptance (the pass test)	Reference column only, never the pass test
Seat count, only counts as a rise once it clears the noise band	Attribution sources such as form / phone / WhatsApp / in-store visit
On the list and cited X/20, denominator fixed for the whole quarter	Clicks coming from AI naturally lack a referrer; any measure of this will undercount
Factual errors (count), read by hand from the brand six questions	Using it for acceptance digs your own hole

Figure: only three rulers count towards acceptance; attribution channels are for reference only

How to capture that reference-column number depends on the channel:

Channel	How to capture it
Form	Add a required single-choice field, "How did you find us?" (including "AI assistant recommended it")
Phone	A dedicated extension number for the landing page
WhatsApp	A link with pre-filled text, the text carrying the page's code
In-store visit	A fixed front-desk question, logged in the daily report
Free-trial sign-up	A required self-reported source field + the distribution of sign-up email domains
Add-to-cart order	A discount code unique to the landing page + the brand-term curve + the product page's direct-visit curve; the platform side has no referrer, so write clearly that it's reference, not attribution
App install	The app store's back-office search-term report

Axis 9 · Content half-life

Ask: "When did this price list / feature table / course schedule last change, and how many times a year does it change?" Look at historical versions — don't take an estimate. You can set the level field by field.

Very short (within 30 days)	Short (quarterly)	Medium (half-yearly)	Long (a year or more)
Don't write actual figures	**Quarterly recheck**	**Half-yearly recheck**	**Event-driven**
Write the structure, the range basis, the price variables and the last check date	Recheck every field when it comes due	Or recheck earlier if an event triggers it	Change it only when the regulation or the practice changes

Figure: four half-life levels; the recheck rhythm follows the field

The reason the very-short level doesn't write actual figures: if AI copies an out-of-date number, what it produces is one factual error, and factual errors (count) is one of your own rulers — that's marking yourself down. A decaying page both falls out of the candidate set and pours wrong words into the "get identified correctly" path; conversely, a visible time anchor gives whoever is fetching the page the confidence to copy it (78% of citations of videos with timestamps are repeat citations, and the same mechanism holds for pages).

One rule shared by every level: the update date must be a line of text in the visible HTML. None of the five major systems' live fetches read JSON-LD — a `dateModified` written into schema doesn't exist, as far as live fetch is concerned. For industries that need field-by-field levels, the specific field list is left to the industry edition.

Axis 10 · Language fork

Run the same batch of intents once in the local language, compare the named sets, and check whether the businesses in the difference set all have a local-language page.

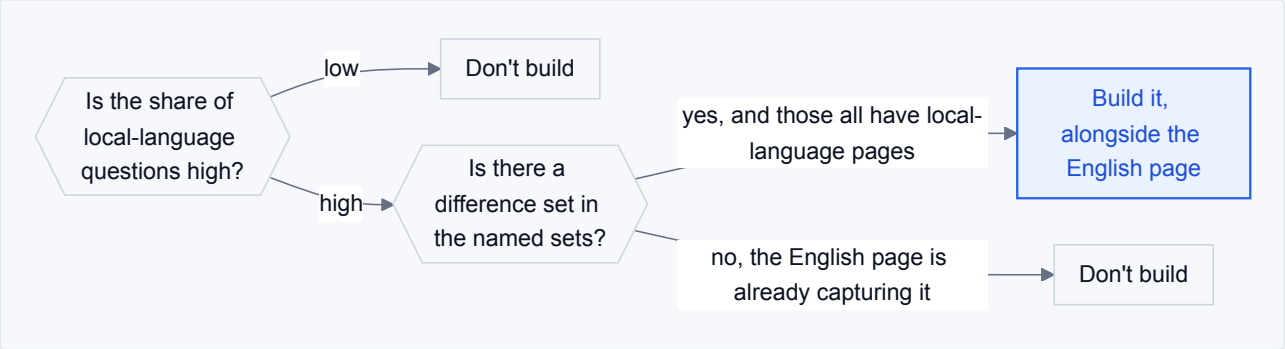


Figure: whether to build a local-language page

The language difference set is the cleanest causal evidence in the whole book: the difference set is itself the evidence — no inference needed. The businesses named in local-language answers are exactly the businesses that have a local-language page. When a product is bound by territorial regulation (pay, tax, invoicing, data residency, industry licensing), a variant appears: the language doesn't fork, but a separate page is built for each jurisdiction, carrying that jurisdiction's regulatory terms in full — this is a "jurisdiction fork", another expression of the same mechanism as a language fork.

8.5 The decision table and rulings

What you'll do in this section: use the four dimensions — control of citation containers × entity anchor × side and switches × question-shape mix — to find the closest row in the decision table; use the four-level ruling to handle conflicts where "more than one row fits"; and when nothing fits any row, use the off-table fallback and add a row marked "first case, not measured".

Once all ten axes are filled in, only four dimensions actually decide the main battlefield; the other six axes govern the details: **control of citation containers (axis 2) × entity anchor (axis 3) × side and switches (axis 4) × question-shape mix (axis 5)**.

#	Combination	Main battlefield	First actions	Notes
Example	Many platform slots + own-site slots high + product anchor + unregulated side + comparison type very high	Own-site pages (comparison / alternatives / category guide) + urgent pricing-page fixes + edit text in platform-slot profiles	The five door gates (robots per subdomain + docs site and trust page CSR) → product-level `/facts` + namesake-conflict check → complete the pricing page + strip `nosnippet` → mirror the trust page's visible HTML on the main domain → `X vs Y` / `X alternatives` / `best X for <scenario>` → docs site SSR → integration tutorial pages	For the values of the other industries already covered, see → General Edition C.1 Ten-axis values for the five industries covered, and three things for week one
Abstain gate	J1 does not hold (fewer than 4 reachable-enough URLs in the intent-level top 10; neither D nor "to ask" counts)	**Decide whether to do it at all first** — this takes priority over every row	Count the number of reachable-enough URLs; fewer than 4 → do not do it, and you don't need to look at the methodology row or the red-ocean row	Fewer than 2, with work already under way → wrong target, switch the work unit
Methodology row	Approach type ≥50% (the company is only mentioned in passing)	Enter the methodology corpus	The five door gates → make the docs site and long-form articles crawlable → bylined contributions (the byline is a real person) → long community answers and open-source traces; **do not build a list-style landing page** (that is the recommendation type's tactic — it doesn't enter the candidate set here)	Not measured

#	Combination	Main battlefield	First actions	Notes
Red-ocean row	Unregulated side + competitors' tactics all in use	Win on structure	On top of the row you've selected, add three things: fill out a full matrix of comparison pages, fill review and comparison-platform profiles to full completeness, and keep visible-update-date discipline (high-frequency updates are themselves the difference from competitors)	Not measured; this is an add-on, not a stand-alone row

How to use it: for a new industry, first decide the side (8.2), fill in the ten axes, then find the closest row in the table.

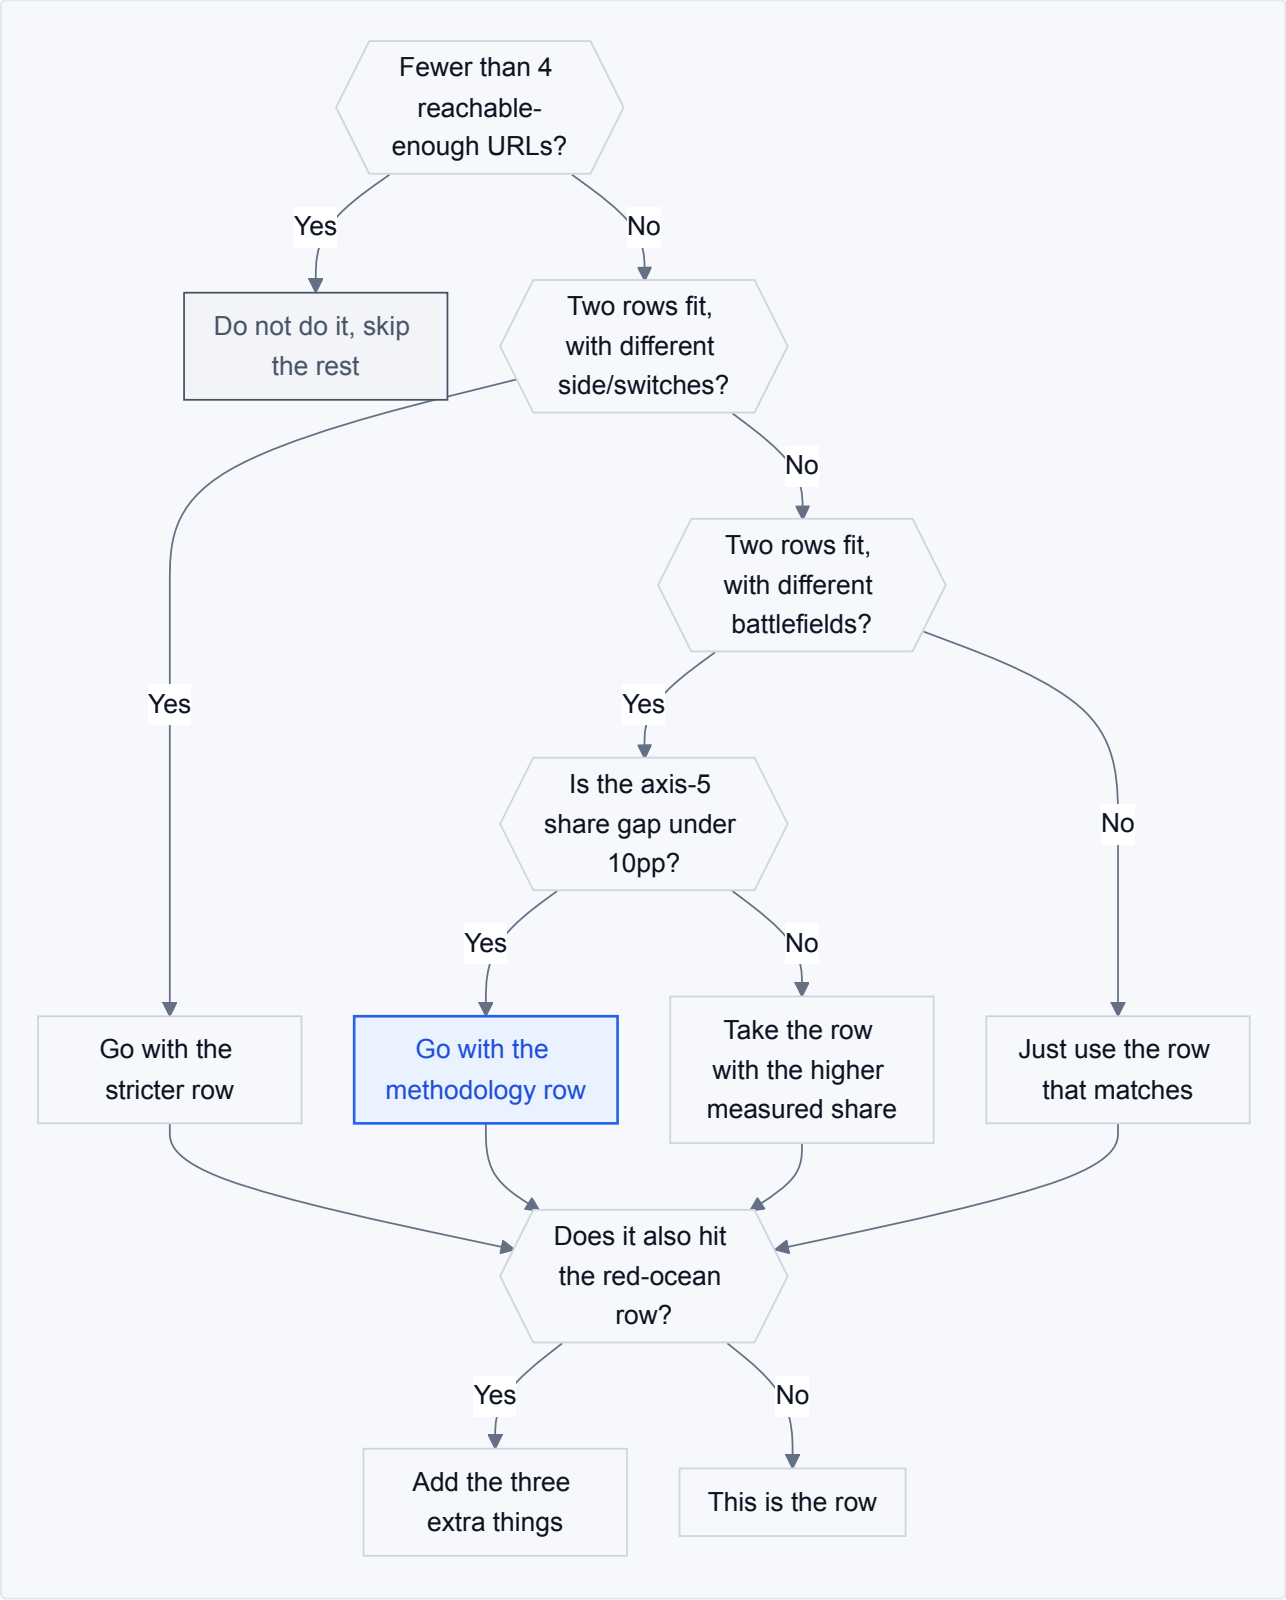


Figure: The four-level ruling, working down level by level; stop once a level matches; the abstain gate is always judged first

Two reasons the figure cannot show: at level 3, "a share gap under 10pp → go with the methodology row" is because the costs are asymmetric — using the approach-type tactic on a recommendation-type category is merely slow, while using the recommendation-type tactic on an approach-type category

means the whole page never enters the candidate set; level 4 is not a stand-alone row, it is an add-on — a hit simply adds three more things on top of the row you already selected.

Not fitting into any row is normal: this table only lists out the five verticals that have been run (only one row is kept here as an example) plus three general-purpose rows — the table isn't broken.

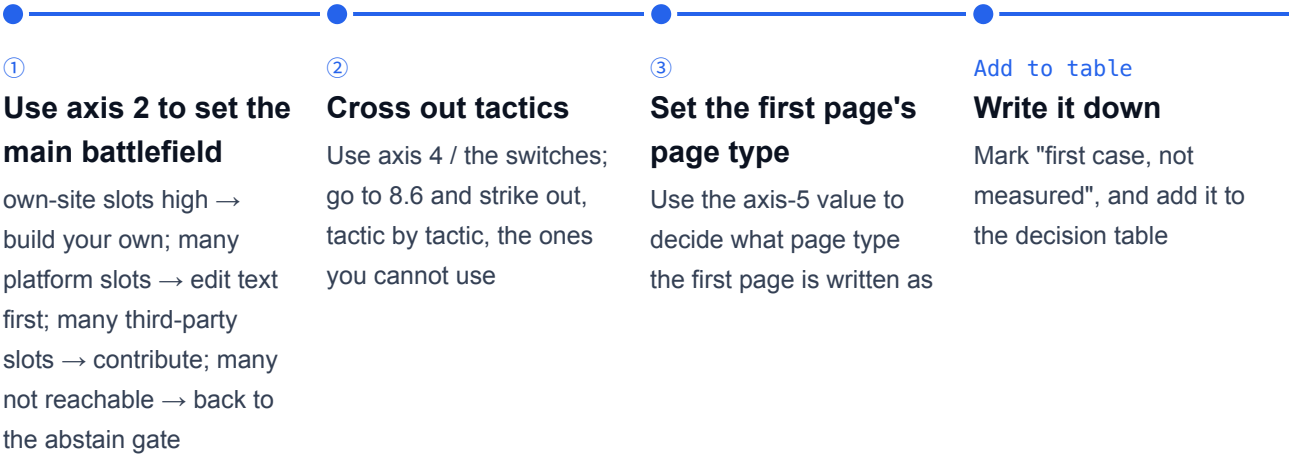


Figure: The order to fall back on when nothing fits any row

A known common combination outside the table (not measured): third-party slot (B) reachable + organisation anchor + unregulated side + approach type or spec type high → the main battlefield is bylined contributions + industry directory and register listings + turning case studies into checkable-number pages, **do not build a list-style landing page**, and the entity anchor for `/facts` has to be built yourself (see axis 3 in 8.3).

Example (dental and aesthetics vs K12 tutoring): on the same axis (shelf page-type mix), opposite values send the tactics in opposite directions.

Dental and aesthetics	K12 tutoring
listicle 0% of citations	Actual search top 9: 5 editorial lists + 3 platform directories
Best/Top pages: removal requests only	Can compete head-on for an editorial slot
Not worth the spend, both because of the 0% and because the regulations ban self-praise and buying list spots	No professional advertising regulations, and lists are the dominant page type

Figure: Same axis, opposite values, opposite tactics (example)

But on one point both sides reach the same conclusion: go for editorial slots only, never buy a paid slot. The reason isn't compliance — buying list spots is legal in K12 tutoring — **the reason is that paid placements and advertorials account for only 0.3% of AI citations: this is an efficiency judgement, not a compliance judgement.** "Don't buy a slot on a page that isn't already cited" governs whether you're wasting money; "if you buy it, disclose it" governs a separate line — the two criteria must not be mixed.

8.6 Seven tactics: how regulated and unregulated sides differ, and why each moves rankings

What you'll do in this section: go item by item through the seven tactics (reviews, lists, prices, testimonials and results, person entities, content boundaries, agency liability), checking whether each can be used on each of the three sides and why it moves rankings. The detail of exactly what you can write on each side is not spelled out here — see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers and → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions; the four figures below only show the strict side's "what gets banned → what fills the same slot".

1. Reviews

In recommendation-type questions, AI copies review text straight into the answer body (70% confidence); the rating and review count are also the sort field for the vast majority of listicle / comparison / directory pages, and listicle pages account for 19.6% of citations.

Cannot do on the strict side		Replace with this
Actively request reviews	→	Passive observation only; never a task metric or an acceptance measure
Reproduce review text, screenshots, star ratings or aggregate scores	→	Checkable-fact sentences: licence number, product model, process duration
Confirm in a public reply that this person is a customer	→	Never confirm anyone's identity

Figure: The reviews tactic: what is banned on the strict side is replaced with checkable-fact sentences

The light side may request reviews (conditions in 5.3); the unregulated side can and must, but there the gatekeeper is the platform's terms instead.

Red line A technical precondition that applies on all three sides and is the easiest one to miss: **AI reads visible HTML; none of the five major systems reads JSON-LD when fetching live — burying only an aggregate-rating structured-data block counts, to AI, as not having done it at all**. A review has to appear on the page as visible text; this is the mistake the unregulated side makes most often. The unregulated side isn't free to request reviews however it likes, either — review platforms have explicit incentive rules: any incentive must be token and applied equally to every invited customer; you must not invite only satisfied customers, and you must not reward only positive reviews; breaking these gets reviews purged and vendor pages actioned. The review gatekeeper has to be measured industry by industry — never inferred straight from the side.

2. Lists and directories

Listicle pages account for 19.6% of citations, and comparison pages get 1.87 citations per search (highest in the whole table) — the two richest shelves. Getting onto a list AI already cites and ranking first on it raises visibility and moves your position earlier in the answer; for the figures and their qualifying sentence see → General Edition D.1 Number discipline: how to label numbers, and what stays internal; for the sample and source see → General Edition A.1 Evidence for mechanics, the door and identity.

Cannot do on the strict side		Replace with this
Pay to get on a Best/Top page	→	Pure directory / price-comparison pages, with no self-praise in the title or entries
Build your own list and rank yourself on it	→	A "how to choose" criteria-based selection guide: list the criteria, don't rank
Claim to be number one	→	Send a correction letter for wrong entries in pure directories

Figure: The lists tactic: what is banned on the strict side, and what replaces it

On the strict side, even labelling it "paid listing" does not release you from liability.

Red line **One rule that cuts the other way (the most expensive mistake on the unregulated side)**: 84% of AI citations come from earned media; paid placements and advertorials account for only 0.3% (see → General Edition 1.1 How buyers ask, and what AI reads). Buying a list spot only works if that page is already cited by AI — buying a spot on a new, uncited list drops you straight into that 0.3% bucket. Before any paid slot, use the *Citation-slot table* to confirm the page is already cited within the question pool; if it isn't cited, don't buy it. Whether a list is worth the investment is an industry variable — every vertical has to measure its own page-type distribution and must not extrapolate the previous industry's conclusion.

3. How prices are written

Price-type queries trigger AI Overviews >80% of the time, and 97% for mixed intent — the highest-hit-rate question type there is. **A fixed price actually comes out ahead on the "extraction" front**: a precise, adjective-free, checkable number is easier to extract into an answer sentence than a "from X" price — a "from" price carries no information, so even if AI copies it, it still can't answer the user's question. The strict side is forced to write a fixed price, so it loses nothing on extraction; what it loses is on "placement" (it can't get onto a price-comparison list).

Cannot do on the strict side		Replace with this
"From", "as low as", original price, discounts, instalments	→	A tiered, itemised fixed final price
Bundles, time-limited offers, "free"	→	Write "included in the quoted price" into the tiered price
Comparing prices with peers	→	See 5.2 for replacement wording

Figure: The price tactic: what is banned on the strict side, and what replaces it

Prices on the unregulated side aren't a free-for-all either: mandatory fees must be included in the displayed price, a struck-through price needs a real baseline, "free" must come with no conditions attached, and fields with a very short half-life don't get written as hard numbers (8.4, axis 9). "Unregulated = nobody polices the price copy" is the most easily misread cell of all, and since the price page is the highest-hit-rate cell, misreading it costs the most. The real trap is here: drip pricing and fake struck-through prices are already enforcement targets — once a platform or regulator names you, what you lose is the credibility of the whole domain, not just one piece of copy.

4. Testimonials, case studies, result and efficacy claims

Short sentences with specific numbers are the material type AI extracts into answers most easily; article pages account for 23.7% of citations, and their citability rests almost entirely on the handful of checkable numbers inside them.

Cannot do on the strict side		Replace with this
Customer testimonials	→	Process steps + duration + number of sessions
Before-and-after photos, even with a disclaimer	→	Brand and model of the material or product + scope of use
Success rate, efficacy or result claims	→	Clearly defined, adjective-free, evidence-backed operating figures

Figure: Testimonials and results: what is banned on the strict side is replaced with evidence-backed operating figures

The strict side isn't out of cards to play — the cards just changed: the figures on the right are numbers too, just as extractable, and just as easy to back up with evidence — the regulations already require every statement to be provable, so a number you can prove is actually the safe zone. The unregulated side may make result claims, but must give the basis and sample, or the claim gets judged unverifiable and is, if anything, less likely to be copied; any "save X% / Y times faster" claim needs a clickable methodology page (sample, conditions, date) — delete any claim with no methodology page.

Red line **A conclusion shared across every side**: a number with a denominator, a year and a stated basis is both what the strict side's rules require and the form AI is most willing to copy — so **the line "we had to weaken the content for compliance" is not allowed**; what gets cut is always the self-praising sentence that was never going to be copied anyway.

5. Person entities

An official registration number plus a link to the official register lookup page is the strongest, cheapest disambiguation anchor in the whole playbook — it's a single-point confirmation of you from a third-party authoritative page; the regulated sides get it for free, the unregulated side has to build one itself (see axis 3 in 8.3). This net advantage only shows the strict side's full hand of cards when read together with two others: a fixed price comes out ahead on extraction (item 3), and evidence-backed numbers are the

safe zone (item 4) — grabbing only the registration-number point is playing with two fewer cards than you have.

6. Content form: where the limits of long educational articles lie

Comparison pages have the highest density (1.87), article pages come second (23.7%). The strict side may write long educational articles, keeping to four boundaries: ① must not lead into efficacy / results / success rate; ② must not contain anything on the ban list; ③ must be signed by the organisation or the licence holder; ④ must not recommend your own business over peers within the article.

The regulated sides are shut out of "naming and comparing competitors", and the legal way out is the "criteria-based selection guide" — list the criteria, don't rank, don't name peers. In page-type terms it still counts as listicle / article, so it still gets a share of that citation pool, without triggering either the self-praise ban or the peer-comparison ban; this is currently the only legal route for the regulated sides into that 19.6% listicle share (not measured, 65% confidence — recommend running one per vertical as a control).

> **Example** (financial advisory): if the content amounts to a product recommendation for an individual's specific situation, it may itself require a licence — adding a disclaimer does not remove the licensing requirement.

7. Legal liability for advertising published on someone else's behalf

This item doesn't move rankings directly, but it decides whether the tactic can be carried out at all. A tactic that can't be executed doesn't exist for this industry, so it has to stay in the playbook table — it must not be moved off into a legal appendix.

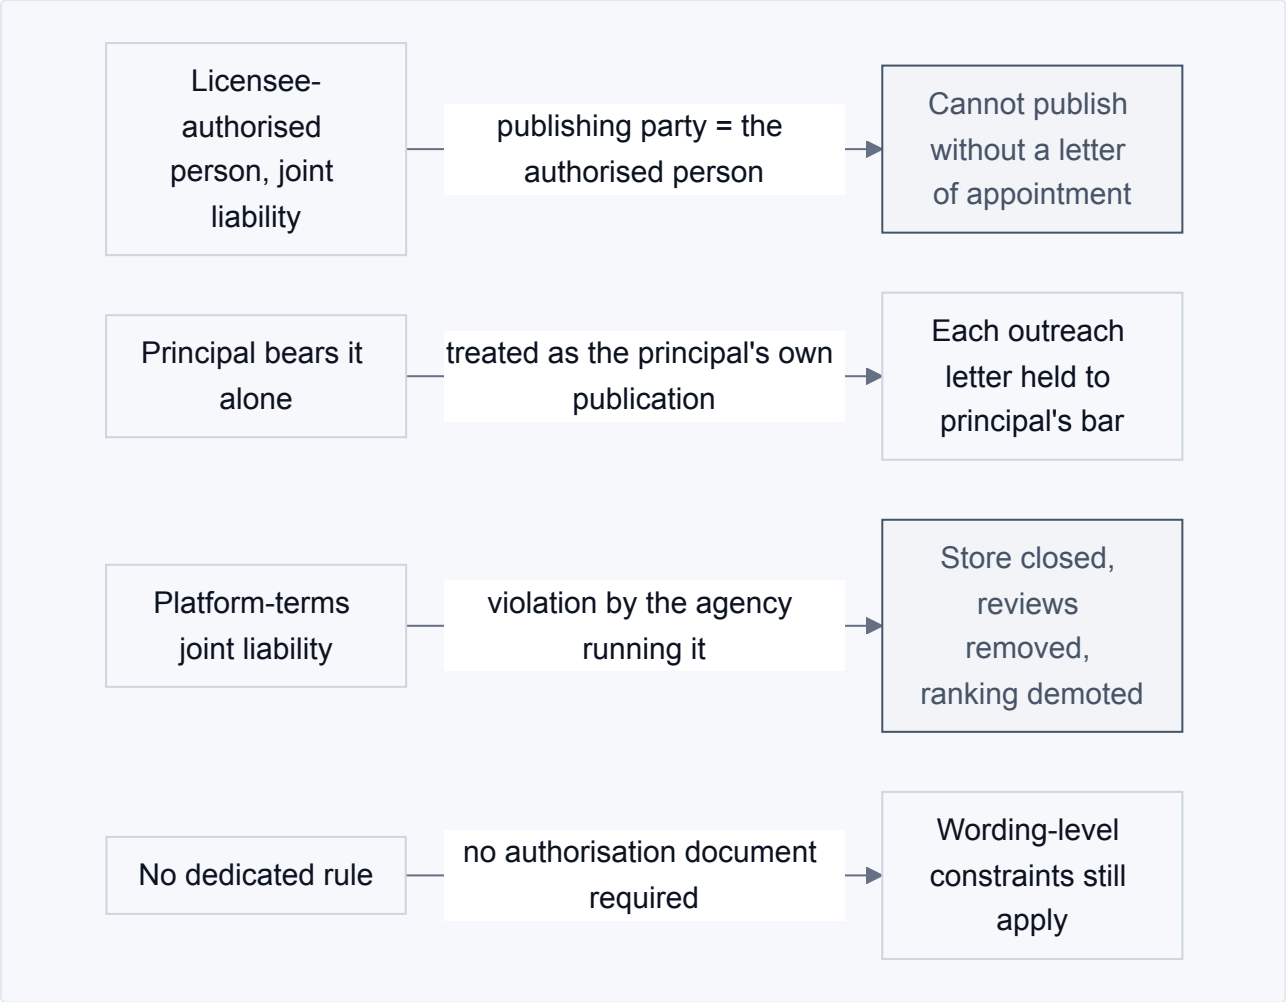


Figure: Four forms of publishing-on-someone's-behalf liability, each with one hard consequence, deciding whether off-site work can be handed to someone else

There's no downgrade path out of the "licensee-authorised person, joint liability" tier: the moment the party publishing on the client's behalf publishes, it becomes the authorised person and shares liability with the licensee, whose obligations cannot be transferred away; without a *Letter of appointment for advertising* ready to sign, nothing may be published in any capacity — no signature means a hard refusal. When the rule writes the publishing party directly into the offence as a liable party, and the business has even one such business line under its name, the whole site is run on the side of its strictest business line (switch D (spans two sides)), and the authorisation and sign-off documents come before work starts (switch A (agency liability)); if switch B (legally required fields) is on, any material published on the client's behalf that misses a legally required field is itself non-compliant material. "Principal bears it alone" also means the principal must not get a third party to publish, on its behalf, promotion that would be a violation if the principal published it itself. On the side with "no dedicated rule", the business bears primary liability, but the party publishing on its behalf may be pursued separately if it makes a misleading statement on its own initiative (inferred, 60% confidence; no direct case found). The real hard constraint comes from platform terms, which are enforced far faster than the law. If you cannot establish which form applies, judge it "cannot publish on the client's behalf" under 8.1 Q6b.

8.7 Cheap tactics on the unregulated side, and substitutes on the regulated sides

What you'll do in this section: match the eight cheap tactics on the unregulated side against their legal substitutes on the regulated sides; use the five substitutes to replace the tactics the strict side loses; and for UGC calls with a reward, once the reward is removed there are still three usable approaches. Fill in the specific form each one takes in your own industry.

Unregulated side	Legal substitute on the regulated sides
Actively request reviews at scale	Strict side: checkable-fact sentences take the extraction slot
Buy paid Best/Top list spots	Pure directory + price-comparison pages, plus a correction letter for wrong entries in pure directories; if listed on a Best / Top page, request removal
Build your own industry list and put yourself on it	A "how to choose" criteria-based selection guide (not measured, 65%)
Comparison pages naming competitors	Compare approaches or methods only, never companies
Before-and-after / success-rate / client-results numbers	Clearly defined, adjective-free operating figures + process and duration
Discounts, original price, time-limited offers	A tiered, itemised fixed final price
Affiliate marketing / ghost-written distribution	Bylined contributions: signed, unpaid, giving checkable facts
Rewarded referrals / UGC calls	Three approaches once the reward is removed (see below)

Figure: The eight cheap tactics on the unregulated side, and the legal substitutes in the same slot on the regulated sides

Item 1 doesn't need a full substitute on the light side: you may request reviews, but without offering any benefit, without filtering, and without ghost-writing them. Item 6, on the side where switch B (legally required fields) is on, becomes a complete price list with every legally required field filled in. Affiliate marketing and ghost-written distribution move a lot of volume but at a very low citation rate (they land in that same 0.3% bucket, see → General Edition 1.1 How buyers ask, and what AI reads); rate it low-efficiency and don't make it the main push (inferred, 70% confidence).

When the unregulated side uses these eight tactics itself, there are four traps: a self-built list that lists only you reads as obvious self-assessment, and both AI and editors will demote it — it only has citation value if it genuinely lists every peer; competitor information goes stale, and a stale comparison page bites back at the identity gate; results figures need a stated basis and date, or AI will keep repeating an old number as if it were the current state; and the constraints on discounts and struck-through prices are in 8.6, item 3.

Five substitutes for the tactics the strict side loses

These five are the strict side's entire outlet for persuasive power; every industry has to fill in the specific form each one takes.

Lost	Replaced with
Reviews + result claims	Third-party credential tier (on the strict side only technical-standard certification and registration status remain)
Self-described credentials	Official registration number + official register lookup link
Price comparison + promotions	A tiered itemised fixed-price page
Efficacy	Being listed on a manufacturer locator (they link to you — inbound)
Testimonials	A signed long-form video and bylined contributions

Figure: The tactics the strict side loses, and the five substitutes that replace them

The wording and the reasons the figure cannot show:

- **Third-party credential tier**: write only the tier's name and who grants it; never "highest tier / the only one / only X on the whole island"; never "most cases". The basis: in an actual search, the one recommendation-type answer that didn't copy a review copied the manufacturer's official tier plus its official registration status instead — when a verifiable third-party tier is available to copy, AI copies the tier, not the review (sample and source in → General Edition A.1 Evidence for mechanics, the door and identity). The strict side does not write manufacturer-granted tiers (Conservative line (not statute text); see the industry edition); "AI copies the tier" is an observation from a single actual search and has not been retested locally.
- **Official registration number**: this is also the hardest anchor on the identity path — a registration number is a unique string across the whole web, so it has the lowest disambiguation cost.
- **Fixed-price page**: captures the >80% trigger rate on price-type queries, and a self-built page gets the same density as the comparison page type.

- **Manufacturer locator**: the strict side's own page writes only the real model number and registration facts of the device, material or system, with no outbound link to the manufacturer; the payoff is on the inbound side — one application in exchange for one third-party inbound link and one brand endorsement.
- **Signed long-form video and bylined contributions**: you don't need a viral hit, just a reference that can be cited, see → General Edition 6.7 Long videos narrated by the named expert.

UGC calls with a reward, once the reward is removed

What the regulated sides ban is "giving something for content", not "obtaining third-party material" — the mechanism behind this tactic (creating citable Q&A material on a third-party domain) mostly survives once the reward is removed.

Once the reward is removed	The specific approach
8a No-reward call for questions	Open calls via contributions, association explainers and community long-answer columns
8b Swap the reward for a byline and exposure	Speaker pages, conference Q&A transcripts, association round-table transcripts
8c Use operating figures to carry the sense of volume	Operating figures with a stated basis, denominator and year

Figure: UGC calls with a reward: three approaches remain once the reward is removed, in order of effectiveness

8a has the licence holder sign and answer, with only checkable facts in the answer — the output is a qa page on a third-party domain, competing for the same shelf slot as UGC. 8b gives the person on camera a byline, not a benefit; the regulations govern benefits, not bylines. 8c replaces the role UGC plays in an answer — the sense that "many people have said so".

Red line **Red line**: 8a and 8b must never involve any form of reward — including prize draws, points, discounts, priority scheduling or free gifts. The moment any reward is involved, it reverts to tactic 8 itself and the whole approach is off-limits. Platform terms often ban rewarded calls outright too, and review sites' incentive rules are almost as strict as the regulated sides' — read the target platform's policy clause by clause before starting.

8.8 Unregulated ≠ easy: admission criteria J0–J7

What you'll do in this section: on the unregulated side, competitors can use every cheap tactic too, so first use the eight testable admission criteria J0–J7 to judge whether it's worth doing at all — if J0 does not hold, don't do it; J1 must hold; at least two of J2 / J3 / J4 must hold; and if competitors have already maxed out the cheap tactics and there is no gap at either the door or the entity, don't do it.

The intuition is "unregulated = easy to win"; the reality is the opposite: competitors have all six tactics open too, and in a category where every tactic is fully open, recommendation-type answers default to copying whatever social proof already exists — content quality alone can rarely jump the queue, and you have to win on structure + scale + update frequency instead. Compliance substitutes can't open up a gap on this side, so differentiation comes down to just three things: a cleanly fixed door, a well-aligned entity, and thoroughly done off-site work.

The method has three layers. **L-a compliance substitutes** (getting persuasive power back under a ban) are worth the most on the regulated sides; on the unregulated side they don't drop to zero, there's just no single set of regulations you can look up — the constraints come from the four base layers in 8.2 step 5, and each layer happens to press exactly on the richest shelf slot: price transparency presses on the price-and-fees page, the ban on misleading statements presses on the results-figures page and review display, the Do Not Call Registry presses on outreach pacing, and platform terms press on review collection and filling in platform-slot (C) profiles. So on the unregulated side, L-a is "scattered everywhere — you have to collect it yourself": make a two-page sheet with links to the original text. **L-b technical door** (robots / WAF / nosnippet / CSR / multiple subdomains / platform doors) is often worth more on the unregulated side — these industries run newer tech stacks, with a higher share of SPA and headless setups. **L-c ruler and target picking** (frozen question pool, noise band, four-state list, ranking by annual opportunity value) is worth the same on both sides. **The criterion isn't "regulated or not" — it's how big the gap is in L-b and L-c.**

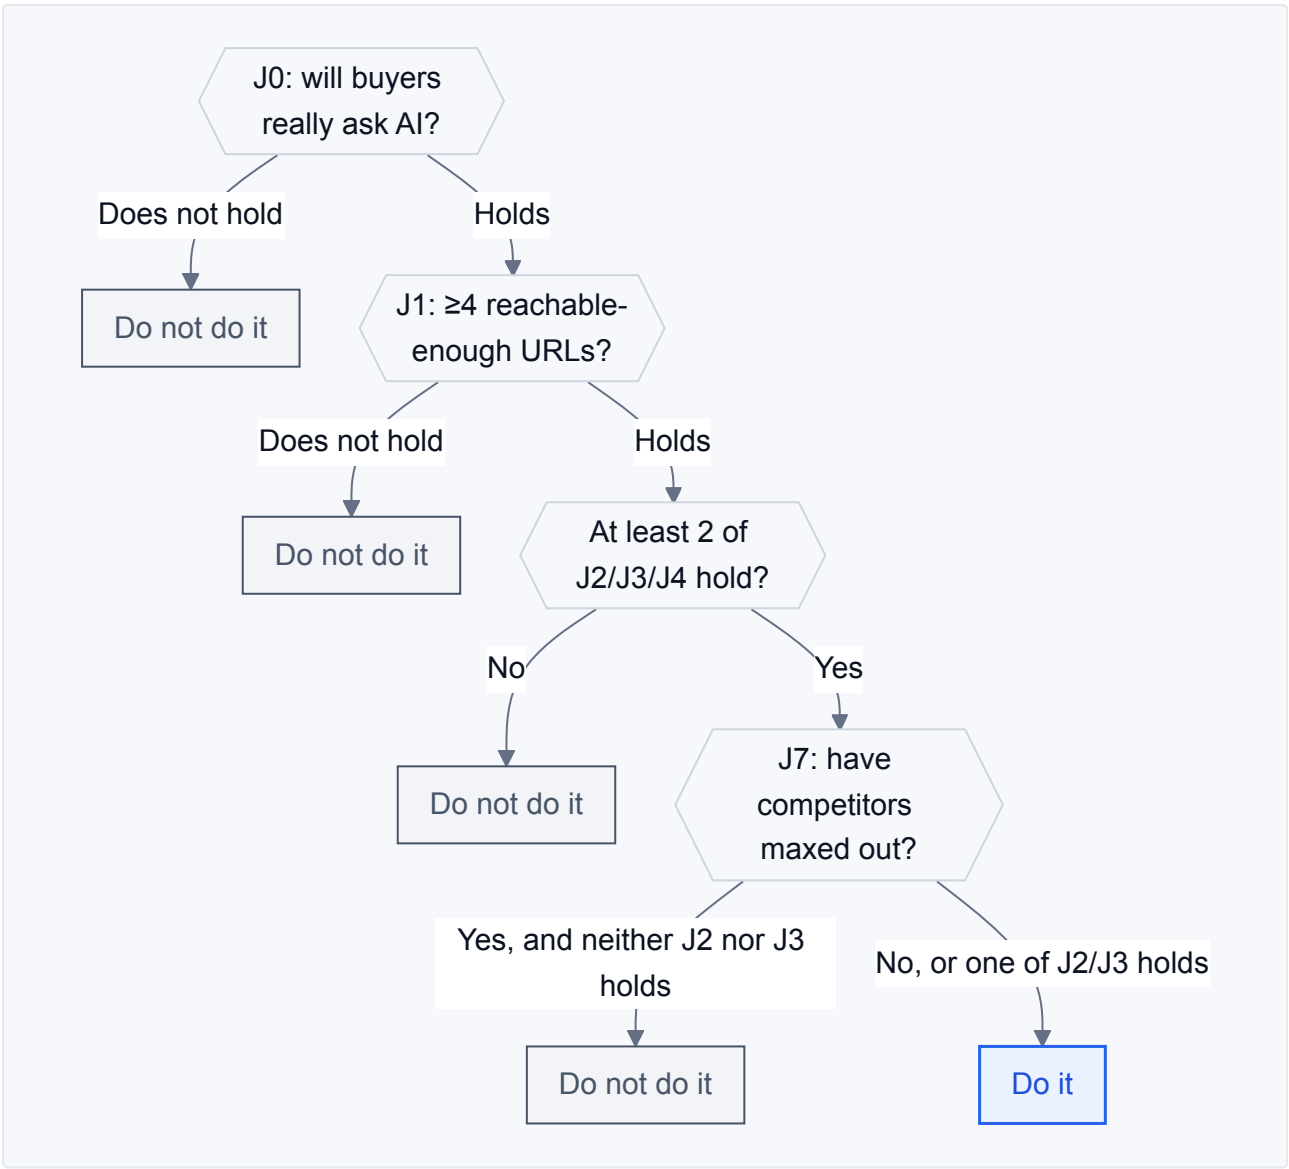


Figure: J0–J7 decide whether it's worth doing; J1 is the book's only abstain line

All eight criteria are testable — never score by impression; the figure only shows the ones that decide whether to do it:

Criterion	What it means	How to test it
J0	Whether the AI entry point holds: buyers really do ask AI (holds) / AI is only used for an initial screen (half-holds) / it runs through a panel, pitch or personal network (doesn't hold)	Run 15 frozen questions through both engines; count whether any company at all gets named
J1	Reachable-enough count ≥4 (fewer than 4 reachable enough in the top 10 for a given question → don't do it)	Tag the four states for every URL in the intent-level top 10
J2	The door is generally closed, and competitors haven't opened theirs either	Turn off JS and view source + the crawler hit table, sampling 5 competitors

Criterion	What it means	How to test it
J3	Entity disambiguation is hard (namesakes, abbreviations, multiple product lines, no official registration number)	Ask each of the five major systems "what company is X", count the misidentifications
J4	The key citation slots are gatekept by editors or an independent third party, and cannot be bought	Check those pages' submission-policy pages: is there a rate card, is anything marked sponsored
J5	The buyer shape still fits "buys infrequently, can't judge quality"	Admission questionnaire
J6	Annual opportunity value is high	Verbal + to be verified
J7	Whether competitors have maxed out the cheap tactics	Sample 5 competitors, count: are they actively collecting reviews, have they bought list spots, do they have vs pages

Why J2 / J3 / J4 specifically: they correspond, respectively, to the parts of the door, identity and shelf-placement gates that competitors cannot buy with money. Reviews can be bought, list spots can be bought, advertorials can be bought — but a door that won't open, an entity that can't be identified, and an editorial slot that can't be bought can only be earned by doing the work. The rankings you can actually win sit in those three cells.

> **Example** (weddings, not one of the industries covered so far): the buyer shape fits perfectly (a once-in-a-lifetime purchase, high price per order, buyers can't judge quality), and it still gets ruled dead — not because the target was wrong, but because this category has no regulatory red line, and a review-management agency can get most of the same effect for less money: that is, J7 hits while J2 and J3 do not hold.

Red line **The most common mistake**: assuming that unregulated or lightly regulated means you can skip the door and identity layers and go straight to writing content — **regulation only takes tactics away, it does not change the building site** — skip them once, and everything after it gets multiplied by 0.

Strict side	Unregulated side
Main effort: finding legal substitutes	Main effort: volume and structure
Technical-standard certification, registration number, fixed price, listing on manufacturer locators	Comparison-page matrix, alternatives pages, review-platform profile completeness
Main off-site push: government and association registers + bylined contributions	Main off-site push: review and comparison-platform vendor profiles + long community answers + integration marketplaces
Fastest payoff: the five business profiles + checking the official register	Fastest payoff: completing platform-slot (C) profiles + comparison pages
Most often missed: the door (all the attention goes to compliance)	Most often missed: the door ("nobody's watching anyway" so it never gets checked)

Figure: The two sides' tactics have opposite centres of gravity, but the layer most often missed is the same one

Government and association registers are free, lose almost nothing between effort and listing, and their share is rising; platform-slot (C) profiles and comparison pages pay off fastest because the container is already inside the cited set, so a text edit takes effect immediately. What both sides miss is the same layer — the door.

Three rules that hold across every side: ① the door and identity layers are exactly the same on all three sides — the robots.txt master template, the three `nosnippet` controls, matching all four identities, the `/facts` page, six-trace alignment word for word — not one item may be dropped; ② the compliance verdict outranks priority ranking: however rich an off-site slot is, if the compliance verdict is "correction only / request removal / do not do", you only send a correction letter or a removal request, you don't schedule it or give it priority; ③ **the side where the only constraint comes from consumer-protection law and platform policy is the easiest one to crash on**, precisely because there's no single set of regulations to check — for these categories, before starting work you must make a separate two-page sheet on "what this platform does not allow for collecting reviews / setting prices", with links to the original text.

8.9 Cross-border: whose law applies

What you'll do in this section: first answer three questions (where the service is delivered, where the offer and acceptance are made, and whether the page targets a given country); a country counts as targeted only when at least 2 of the six targeting signals are ticked. Then apply the five rules in order to take the intersection and write one version to the strictest standard — don't split domains or brands by country.

Three questions to decide

Answer three questions before asking which law applies: ① where is the service / goods delivered (check the place of contract performance, the distribution of shipping addresses, the store or delivery location — use back-end data where you have it, don't go on what people say); ② where is the offer made and where is it accepted (check the merchant entity named on the checkout page, where the party receiving payment is registered, the jurisdiction clause in the terms); ③ does this page target a given country — this is the targeting test countries actually use for territorial jurisdiction.

For question three, tick each of six signals: ① a dedicated page in that country's language, or hreflang pointing to that country; ② prices quoted in that country's currency, or that country's local payment methods offered; ③ the page states delivery to that country / available in that country; ④ it uses that country's ccTLD or a subdirectory for that country; ⑤ it is listed or advertised on that country's platforms, app stores or ad networks; ⑥ the page names that country's local regulatory terms, local credentials or local customer-service channels.

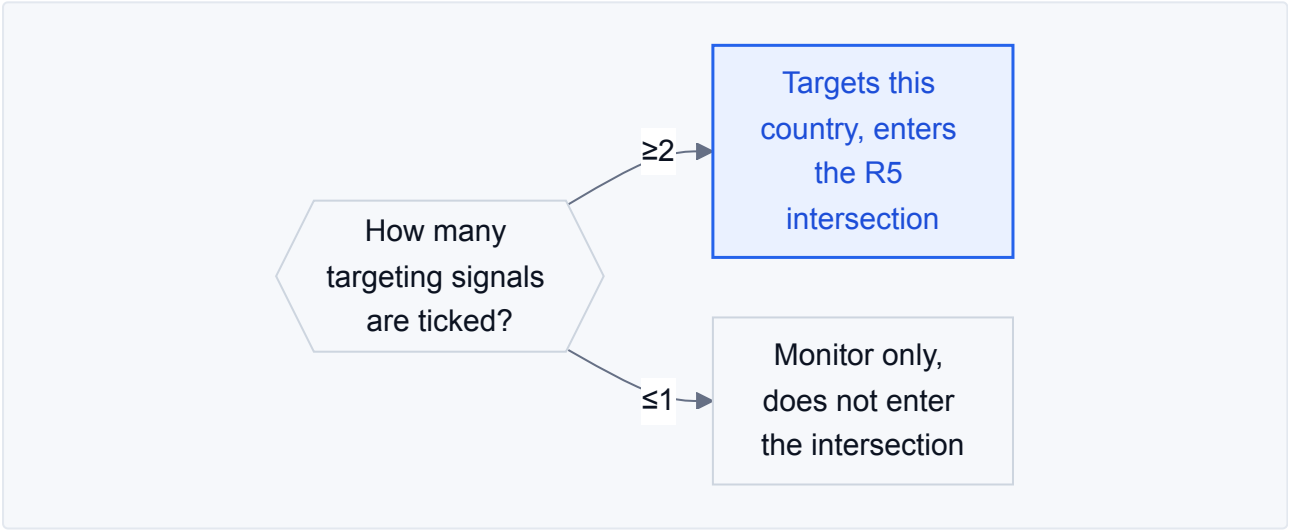


Figure: Six targeting signals; a country counts as targeted, and enters the intersection, only when ≥2 are ticked

With ≥2 ticked, treat the page as "targeting that country": that country's advertising and consumer law applies. The test has to be reproducible: all six signals are facts visible on the page, so two people opening the same URL should tick the same ones; if they disagree, go with whichever side ticked more.

> Note **A common misunderstanding**: "where the ad can be seen" is not a criterion. A web page is naturally visible worldwide, so this question gives the same answer for every page — it tells you nothing — and treating it as a criterion pushes "take the intersection and write to the strictest standard" into "write to the strictest standard worldwide", which is a standard nobody can actually follow. The only thing you can judge is the six targeting signals in question three above.

Five rules, applied in order

R1

Territoriality doesn't loosen

If the licensee is local and the service is delivered locally, the full set of bans applies exactly as before

R2

The base layers still apply

Either the supplier or the consumer is local, or the offer and acceptance are made locally

R3

Two routes can reach you

① platform delisting, store closure, review removal ② payment and logistics — they don't need to wait for a regulator to come knocking

R4

Stack the stricter one

If the destination country has its own banned-word list, stack it with the local one and write to whichever is stricter

R5

Fix the intersection

Don't write versions with different degrees of leniency for different countries; split by language using hreflang

Figure: The five cross-border rules, applied in order

R1 is inferred, 85% confidence: a page aimed at foreign buyers is still bound by the full set of bans; you may not wave through any item in the local ban list on the grounds that "this page is for foreigners". The reason: the regulations govern the licensee's advertising conduct and the licensable services delivered locally, no provision has been found that exempts by audience location, and under the regulator's own position, all of the internet (social media and websites included) counts as a regulated medium. R2 has a source (the territorial provision). R3 is inferred, 75% confidence: platforms enforce faster than the other country's regulator. R4's destination-country statute has a source; "how it actually applies to an overseas operator" is inferred, 70% confidence.

> **Example** (a Chinese-language page aimed at consumers in China): China's *Advertising Law* bans absolute terms (most / first / only / national-grade / highest-grade / best), with penalties of RMB 200,000–1,000,000; there is also the SAMR *Enforcement Guidelines on Absolute Terms in Advertising*, and, as paraphrased second-hand, not every phrase containing "最" (most) or "第一" (first) counts as a violation — wording that points at the operator itself rather than the product, that expresses an order in time or space, or that is limited to a specific scope the business can substantiate itself, may leave room for not being treated as a violation (Original text not obtained; not used as a basis for leniency). For overseas operators targeting consumers in China, enforcement in practice runs through platforms (inferred, Conservative line (not statute text)).

The strict side's self-praise ban list		The destination country's absolute-terms list
Best, optimal, first, only, leading	→	Mostly the same set of words
Top-tier, authoritative, national-grade, highest-grade	→	Has exceptions (not treated as violations)
No exception clause	→	Its exceptions never clear a word the strict side bans
Write to this list and both sides are satisfied at once	→	Banned outright on the strict side does not mean penalised outright here

Figure: The two banned-word lists mostly overlap; writing to the stricter one satisfies both, but not the other way round

So when the two lists mostly overlap, writing the destination-country language page to the strict side's list satisfies both sides at once; the local-language page tactic (axis 10 in 8.4) can be carried straight over into both the regulated sides and cross-border cases, with no need to write a separate set of copy rules (inferred, 80% confidence).

R5 explained: why you must not split the entity by country

Splitting domains / splitting brands shatters the evidence for the identity gate. AI confirms an entity from consistent signals across sources (same name, same address, same registration number, same account) — split into three brands and none of the three is strong enough, so none of them gets named. The correct approach: **one domain + hreflang + each language page linking to the others + identical entity information across the whole site.**

On the data side: forms, tracking scripts and remarketing pixels are configured to the strictest law among the buyers' countries — take the strictest, don't tier by country; the specific differences between each country's rules have not been measured, so this only gives the basis, not the detail.

Appendix A

The evidence: the numbers and sources behind every conclusion

Conclusions, sample sizes, and what you can say publicly

A.1 Evidence for mechanics, the door and identity

A.2 Evidence for picking targets, writing pages, off-site and retests

A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types

A.4 Page-type measurements (2): tier B and C samples, and the AI Mode retest list

A.5 List of sources

Every sentence in the main text that says "see Appendix A for evidence" ends up here: a one-line conclusion / a number / the sample and timing / the source / the confidence level and whether it can be said publicly, one row per item. This appendix does not repeat the main text's rules; it only gives the row of data that sits behind each rule.

A.1 Evidence for mechanics, the door and identity

What you'll do in this section: check the original basis for any number in the mechanics, door or identity layers — sample size, source, confidence, and whether it can be said publicly.



Figure: Out of 8 test prices, how many each system retrieved — all from visible HTML

searchVIU measured the live fetches of five systems — ChatGPT, Claude, Perplexity, Gemini and AI Mode: not one of them reads JSON-LD, and hidden Microdata or RDFa are ignored the same way, even when that piece of information exists nowhere else on the page. Measured, may be said publicly; it can only be used to say "facts must be written out, one field per line, in visible HTML" — it does not follow that "schema is useless".

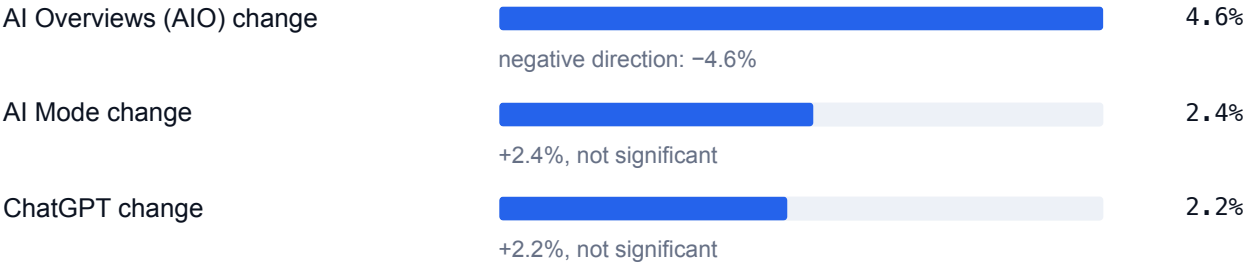


Figure: After adding JSON-LD, only AI Overviews showed a significant change, and it was negative

Ahrefs tracked 1,885 pages that added JSON-LD against 4,000 control pages (2025-08 to 2026-03). After adding JSON-LD, AI Overviews (AIO) citations changed −4.6% — the only significant result, and negative; the other two platforms showed no significant effect. The only conclusion is that "schema needs one half-day job to seal it once, and it does not go on every page's checklist"; it does not follow that "schema is useless": it works through Google's Knowledge Graph, an indirect path.

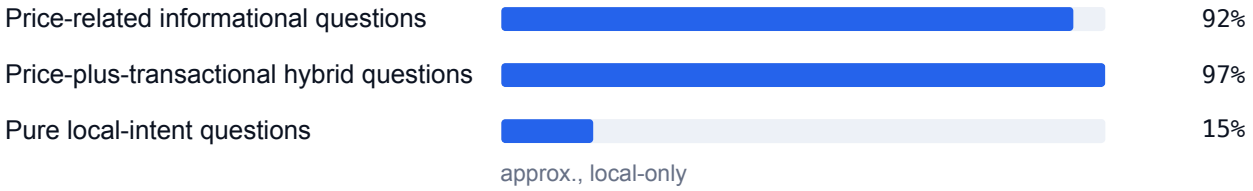


Figure: Cost questions have the highest AI Overviews coverage of any question type

Whitespark, 540 queries / 3 US cities / 6 industries (including dentists, medical clinics, personal injury lawyers). Can be used to say "cost questions are the question type with the highest AI coverage, not just one of them"; US sample, must not be written as a Singapore number.



Figure: Most AI Mode local citations point back to the business profile

Steady Demand, 1,487 queries / 50 metros / 14,472 citations. May be said publicly, but the study states outright that it did not separately quantify the citation contribution of business hours, the service list, photos, Q&A or posts — the profile is an entry ticket, not a lever, and must not be used in place of other actions.



Figure: Even without a published price, AI quotes one for you anyway



Figure: A category page full of prices, without a single judgement sentence (example: an e-commerce category page)

A Singapore brand's D2C category page, raw-fetched with a retrieval-crawler UA. Two more numbers from the same page: 460,655 bytes of HTML, only 2,182 characters of visible body text left after stripping tags (the original test record gives body text as five parts in ten thousand; bytes and characters are different units, not recalculated here). Measured, 1 case on a single page — directional only.

Conclusion	Number	Sample and time	Source	Confidence · can it be said publicly
Mainstream AI retrieval crawlers execute zero JS; the only exceptions are Gemini and Applebot	500M+ fetches sampled, no evidence of JS execution; none of `GPTBot` / `OAI-SearchBot` / `ChatGPT-User` / `ClaudeBot` / `Claude-User` / `PerplexityBot` / `Bytespider` / `Meta-ExternalAgent` execute JS	2024-12	Vercel × MERJ	Measured, may be said publicly; "69% of AI crawlers don't execute JS" is a mis-cited second-hand figure — do not use it
Taking `llms.txt` out of the prediction model made citation-frequency prediction more accurate, not less	No specific improvement figure given	Sample size not stated	SE Ranking (XGBoost model); Google's 2026-05-15 AI optimisation guidance states outright that it isn't needed	Backed by both an official document and a third-party study, may be said publicly; "this is the biggest fake GEO lever of 2026" is this book's own judgement, not the study's wording
Video view count / likes / subscriber count are almost uncorrelated with citation frequency	Correlation coefficient about −0.03	1.7M data points	Otterly	Measured, may be said publicly (the full set of video evidence is in A.2)
GPTBot is a training crawler, not a retrieval crawler; blocking it does not affect whether you get cited	88.2% of sites that block GPTBot are cited anyway	—	BuzzStream (other sources in A.5)	Measured, may be said publicly
Applebot-Extended is only a training opt-out switch and does not affect the Applebot retrieval leg; all three Anthropic bots obey robots.txt	—	—	Apple's *About Applebot*; Anthropic's official documentation	Official documentation, may be said publicly

Conclusion	Number	Sample and time	Source	Confidence · can it be said publicly
At least eight retrieval sources; local answers mainly go through Foursquare + Yelp + Labrador-local, medical questions through Labrador-medical	OpenAI's own index, Labrador (split into vertical sub-indexes such as general / local / medical / legal / news / shopping), Bright Data, Oxylabs, a third SERP-fetching path, Yelp, TripAdvisor, Microsoft Web IQ, and two internal pipelines; OpenAI and Foursquare publicly announced a partnership in 2024-12, 100M+ POIs / 200+ countries	2026-05-21 to 07-21, captured through ChatGPT's own SSE `result_source` field	Peec	Measured, may be said publicly; IndexNow only feeds Bing and Yandex, it is not a switch for the ChatGPT leg
OpenAI has signed a deal with Yelp, and Apple Business has broad coverage — claiming the five business profiles adds two more trustworthy sources	OpenAI signed with Yelp on 2026-07-23; Apple Business covers 200+ countries including SG	—	Compiled internally (no URL given in the original)	70% confidence, no URL given in the original; may be said publicly, but only in the "free, done in passing" wording, never promising a result
Buyers rarely click the links in an AI summary; once they have a name, they search for it separately	About 1% of visits click through	2025, 900 US users	Pew	US sample, may be said publicly but must carry a qualifier

Conclusion	Number	Sample and time	Source	Confidence · can it be said publicly
Most AI citations come from earned media; news coverage alone accounts for 27%; paid placements and advertorials are only 0.3%	84% / 27% / 0.3%	25M links, 17 industries	Muck Rack	Cross-industry aggregate, may be said publicly with a qualifier (not separately validated for the three Singapore industries); the correct reading is "84% is other people writing about you on their own initiative, not you buying a placement" — it must not be read as "84% comes from third-party sites" and then used to put effort into directories, paid listings and listicle submissions, which is exactly that 0.3% bucket
Even when the buyer searches a brand name, owned content still accounts for a low share	23%	Sample basis not detailed in the original	Omniscient	May be said publicly; the denominator differs from the 84% in the row above — the two numbers must not be subtracted or compared side by side
Page-type distribution is an industry variable, not a general rule (example: healthcare)	Across all industries: articles 23.7%, listicles 19.6%, comparison pages 1.87 citations per search (top of the whole table), category 5.2%, profile 3.8%; healthcare/medical breakdown: articles 54%, listicles 0%, category 5.2%	25,337 citations	deltaV	Measured, may be said publicly; the healthcare breakdown is an example only and must not be applied to other industries

Conclusion	Number	Sample and time	Source	Confidence · can it be said publicly
Scouting numbers can only order the work; they must not be said publicly as the current state	Recommendation-type / non-recommendation-type review-copying frequency 67%/0% (n=15); second-opinion type: 0 businesses named at answer level, 6 in citation slots; the per-tier competitive conclusions for Chinese and Indonesian; the item on Invisalign's official tiers replacing reviews; the off-site page-type share group	Generic search API, region not locked, single engine, one-off	Internal measurement, 2026-09-21 etc.	Lifting this restriction requires all three conditions at once: ① already retested on your own frozen question pool using the OpenAI direct leg + Gemini model leg, region locked to SG ② the retest date and engine are stated ③ what you cite publicly is the retested number; any number derived from these is equally not for external use (general rule in → General Edition D.1 Number discipline: how to label numbers, and what stays internal)
Google's own wording confirms the AI Overviews gate is `Googlebot` + `nosnippet`, not `Google-Extended`	"Google-Extended does not impact a site's inclusion in Google Search nor is it used as a ranking signal in Google Search." "a page must be indexed and eligible to be shown in Google Search with a snippet" "robots.txt directives for Googlebot is the control"	Read on 2026-09-21	Google's crawler documentation + *AI features and your website*	Official documentation, verbatim; may be said publicly

Conclusion	Number	Sample and time	Source	Confidence · can it be said publicly
Google's official documentation lists Merchant Center and Business Profile in its AI Overviews / AI Mode best-practices checklist, and states outright that no extra schema is needed	"Checking that your Merchant Center and Business Profile information is up-to-date" is on the checklist; "You don't need to create new machine readable files, AI text files, or markup to appear in these features. There's also no special schema.org structured data that you need to add."	Read on 2026-09-21	Google, *AI features and your website*	Verified (official wording), may be said publicly; the line "AI Mode / AI Overviews / Gemini's product answers all draw from the same Shopping Graph" is an internal-pipeline detail that has not been verified — use it only to order the work, never as Google's official position
Wikipedia is the single largest node in ChatGPT's citation set, and Wikidata is its structured anchor	Top-10 source share about 47.9%; 7.8% of total citations out of 680M citations; Ahrefs 2026-07 tracked 8.9%	Profound (680M citations); Ahrefs 2026-07	Profound; Ahrefs	Measured, may be said publicly; Wikidata's admission criteria require one of three conditions, and the relevant one requires that the subject "can be described using serious and publicly available references"; entry is not automatic
Pure directory listing is the page type with the lowest citation density in the measurements; the single biggest effect is the entity-strength multiplier at the person layer	category 5.2% + profile 3.8% (in healthcare: listicles 0%, articles 54%); the single biggest effect is an 8x entity-strength multiplier (Wikidata record, business-media coverage, structured disclosures); independent businesses with only Yelp reviews and a Google profile appear only unreliably, and 78% of independent local businesses have an AI citation share of roughly zero	deltaV, 25,337 citations; 5WPR, 320+ prompts / 10 metros / 5 engines	deltaV; 5WPR	Measured, may be said publicly; a single third-party study, US sample, not separately validated for the three Singapore industries — when cited, the qualifier must be written in the same sentence, and it must not be used to promise any result or amount

Conclusion	Number	Sample and time	Source	Confidence · can it be said publicly
On the regulated sides, almost all persuasion is banned at the organisation layer, but not at the person layer (example: healthcare)	The regulator explicitly allows practitioners to publish their qualifications, areas of practice, practice arrangements and contact details	—	The industry regulator's own rule text (rule text and numbering in the industry edition's appendix)	The original record says "explicitly allowed." Both internally and externally, be clear that "allowed ≠ unregulated": the industry's nine advertising standards (factual / accurate / verifiable / not extravagant / not misleading / not sensational / not persuasive / not comparative / not disparaging) still apply to personal profiles; laudatory terms have their own separate ban list from the regulator
Identity fixes are the only lever where "what AI says often changes within 7–14 days of the fix"; writing pages, off-site and video all pay back on a scale of months	7–14 days	Internal tabletop observation / internal review	Internal	70% confidence, not measured, **never state the scale publicly**
The one recommendation-type question that didn't copy from reviews had its answer body copy from the manufacturer's official tier + the regulator's specialist register instead (example: dental)	4 of 6 recommendation-type questions copied reviews and star ratings; the one that didn't copy reviews copied the manufacturer's official tier (Black Diamond / Platinum Elite) + the regulator's specialist register	n=15, single engine, 2026-09-21	Internal measurement	85% confidence; before a region-locked retest, use for internal ordering only, not in external messaging; this is a mechanism observation, not compliance clearance; on the strictly regulated side, do not put manufacturer tiers on your own pages (the conservative line)

Industry rule-text sources: where the main text marks something "Statute text" and the rule belongs to a particular industry's regulator (for example, the rule that disclosing payment does not exempt you

from liability, for list pages whose title is itself laudatory, or the four items explicitly allowed in a practitioner's personal profile), the rule text, its number and its URL are kept in that industry edition's appendix, not in the General Edition; the main text always writes "rule text in the industry edition's appendix", never pointing to this appendix.

A.2 Evidence for picking targets, writing pages, off-site and retests

What you'll do in this section: check the original basis for the ratios used in picking targets, the web control leg, the basis for annual opportunity value, where off-site effort goes, the video correlation coefficient, the noise band and the confidence list.



Figure: Bare words can produce a fake zero: the same query, differing only by a location qualifier

PRWeb, 1,120 queries / 7 industries / 13 cities. Question-pool wording can produce a fake zero visibility; the question-pool ratios are therefore fixed: with a location ≥60%, price words ≥30%, bare category words ≤10%. Measured, may be said publicly.



Figure: For Chinese-language questions, the named set is mostly empty for price and scenario types

Searched on 2026-09-21, region not locked to SG, single engine, one-off, 80–85% confidence. This is for internal ordering; before it goes into external material, retest it region-locked to SG on your own frozen question pool; before launch, spot-check 3 questions from your own pool. The bar is not whether you're on a list — it's whether you have a Chinese-language landing page.

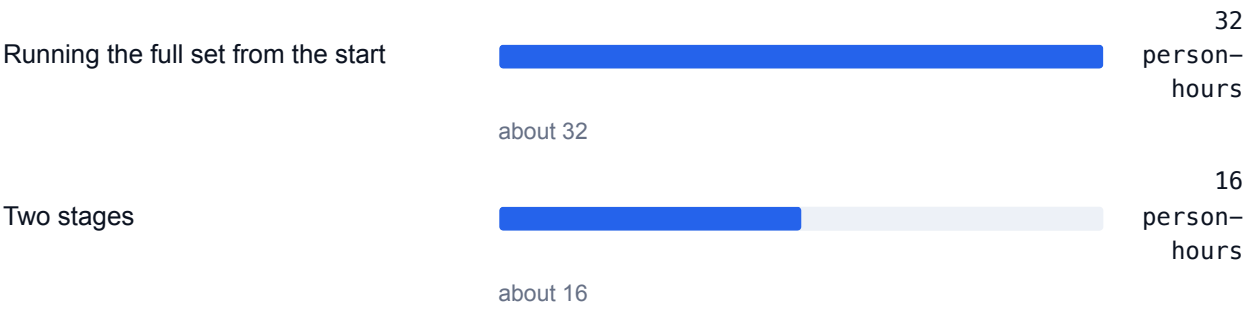


Figure: A two-stage probe saves half the person-hours

Internal figure, 80% confidence. Probe count: running the full set from the start is 30 × 5 × 2 = 300; two stages is 60 + (12–18) × 5 × 2 = 180–240, plus one round of same-week retesting. The cost: dead-pile questions only ever get an n=2 conclusion, which can never be reported externally; a question wrongly knocked out gets no second chance for the whole quarter. The dead-pile criterion is an **operating convention, not a measured conclusion** — what backs it is the direction seen across our own 200 tests (educational-type questions almost never name a business), 70% confidence; the red line for external wording is that you may only write "no business named in the coarse screen", never "zero visibility".

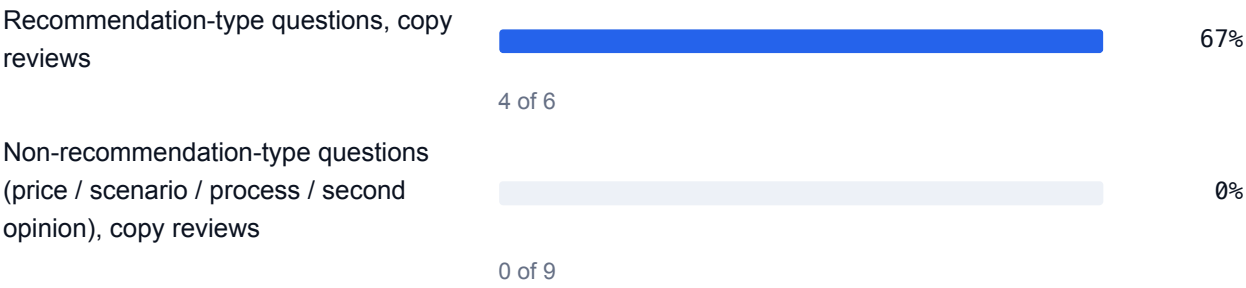


Figure: Only recommendation-type answers copy reviews word for word

Searched on 2026-09-21, n=15, single engine, 85% confidence. Before a region-locked SG retest, use this only for internal ordering and to decide the direction of the work — not as a promise and not in external messaging; how this prints in the monthly report is in the confidence list at the end of this section.

D1–D30: why each step moves rankings

Time	Why it moves rankings (with sources)
D1 morning · the door's five technical gates	A closed door = everything after it is multiplied by zero. A ×0/1 switch, cause and effect certain, zero cost; gate 0 comes first because it's the only direct evidence, and because access permissions take the longest to wait for
D1 afternoon · the work order	The three inputs — side, acceptance tier, and annual-opportunity-value ranking — are already in hand; this is filling in a form, not research
D2–D3 · the shelf map and comparison table	Turns probe data you've already paid for into tasks, without spending on new probes
Week 1 · entity and person pages	When machines identify the wrong business, all page work returns zero or even negative; the person page also doubles as the landing point for bylined off-site contributions and video credits
Weeks 1–2 (in parallel) · the four Group B steps	Institutional and government source share rose from about 1/6 to nearly 1/3 (Otterly / Axios, ChatGPT's 2026-08 source reshuffle), and effort put in converts to being on the list with almost no loss — it doesn't depend on a third-party editor's approval

Time	Why it moves rankings (with sources)
Weeks 1–2 (in parallel) · the first off-site round works only the "already present" class	Putting content into a container that's already being cited pays off on a scale of days to weeks (inferred, 65% confidence, not measured)
Weeks 2–4 · the thick-page main line	Price-related queries trigger AI Overviews >80% of the time, hybrid intent 97% (Whitespark); comparison pages get 1.87 citations per search, top of the whole table, and this is counted by page type — a self-built page earns it just the same (deltaV, 25,337 citations); Chinese price-type / scenario-type / second-opinion-type questions are 80% empty, the only open field within 90 days where you can pick up seats starting from zero
Within 14 days of each page launching · a long video narrated by the named expert	Correlation coefficient between YouTube mentions and AI Overviews visibility: **0.737** (backlinks are only 0.218, Ahrefs 2026); YouTube accounts for **23%** of citations in Google AI answers (5W Citation Share); **94%** of YouTube AI citations go to long-form videos, and **78%** of videos with timestamps get cited repeatedly (Otterly, 1.7M data points)
Weeks 2–3 · bylined contributions to earned media	**84%** of AI citations come from earned media (journalism alone is 27%), while paid placements and advertorials are only **0.3%** (Muck Rack, 25M links / 17 industries)
Every Monday, weeks 2–4 · the weekly competitor seat report	You can see week by week which cells competitors are fighting for; it also doubles as next week's task list — whoever pushed you out, go and take the source they're standing on
Weeks 3–4 · cut the second off-site round to a minimum	Every hour freed up converts to pages; Group A is a structural gap in Singapore: staking a whole line of work on third-party inventory that cannot be created comes up empty every quarter, automatically
One day at the start of every month · the web control leg	The citation distribution between the two API legs and the web version is a systematic skew, not noise: overlap with authoritative-media rankings is 45.5% on the web vs 27.3% on the API; public-broadcaster sources are 34.6% on the web vs 12.2% on the API; a single domain can shift from position 1 on the web to position 61 on the API (welt.de); running 5 rounds shows the same size of skew as running 50 (University of Hamburg + Leibniz Institute, 2025-11, 24,000+ answers / five weeks)
D30 · the monthly review checks the door and identity unconditionally, once each	These two layers fail silently — a page that's never even being read, or AI identifying the wrong business, can easily coexist with a seat count that's rising for some other reason

The full basis for the video leg

Beyond "video view count / likes / subscriber count are correlated with citation frequency at about −0.03" (Otterly, 1.7M data points), already given in A.1: the correlation coefficient between YouTube mentions (title / transcript / description) and AI Overviews visibility is **0.737** (Ahrefs 2026; backlinks are only 0.218 — the strongest single correlation of any factor measured, more than three times that of

backlinks); YouTube accounts for **23%** of citations in Google AI answers (5W Citation Share); **94%** of YouTube AI citations go to long-form videos, **78%** of videos with timestamps get cited repeatedly, and description length has r=0.31 (Otterly, the same batch of 1.7M data points); by platform, Perplexity 38.7% + Google AIO 36.6% — the two together account for three-quarters of YouTube citations, and the gain doesn't land in ChatGPT's seat count. 0.737 is a correlation coefficient, not causation: the mechanism hypothesis is that the YouTube page is already in the cited set (the container is already inside the candidate set, so words written on it have a chance of being copied) — causation has not been measured.

Two numbers that must not be used side by side

Number	What it measures	The only place it may be used	Where it must never be used
Ahrefs 43.8% (ChatGPT citations from best-of blog posts, 750 queries / 26,283 URLs)	Share — how much of all citations the listicle class accounts for	Deciding whether to work a category at all (whether listicles belong in this quarter's scope)	Setting single-page priority, deciding how many letters to send, or putting it in the same table as page-type density
deltaV page-type density (comparison 1.87 citations/search; category 5.2%, profile 3.8%; in healthcare: listicle 0%, articles 54%; 25,337 citations)	Density and within-industry distribution — within a category you've already decided to work, which single-page type is more likely to be cited	Ordering single pages within a category you've already decided to work	Ruling out an entire category; applying it across industries

If the two appear side by side in the same table, that is a mistake: send it back for rewriting.

Reddit and off-site source shares

Reddit's share of ChatGPT answers fell ≥73% in 2026-08 (daily average 497 → 132), while ChatGPT's total citations actually rose 3.5% over the same period; but a crash of the same size in 2025-09 later fully recovered, and ChatGPT is still reading Reddit (about a quarter of pages fetched during the study period were Reddit threads) — what stopped is putting it into the answer, not reading it (public tracking by Otterly / Axios + the same batch's fetch-side records). Criterion: only put it on the task list when this month's share is > 2%; at ≤ 2%, just log it; only write it into the regular cadence once it's been > 2% for two months running. In the same batch, institutional and government source share rose from about 1/6 to nearly 1/3 (same Otterly / Axios tracking, ChatGPT's 2026-08 source reshuffle).

Off-site timing and hit rates

- Putting content into a container that's already being cited pays off on a scale of days to weeks; a new self-built page, 2–8 weeks (inferred, 65% confidence, not measured).

- Paid listings and directory submissions: a free editorial slot, 2–4 weeks; a paid slot, 4–8 weeks (both inferred, not measured).
- Bulk outreach reply rate 15–25%; from a placement going live to a change showing up in AI citations, 2–6 weeks (inferred, 55% confidence, not measured). These two numbers are used only to schedule internal work hours.
- The time for identity fixes to pay off (7–14 days, 70% confidence, not measured) is in A.1.

Page length and differences by side

- Cited pages average 2,290 words, 78% over 1,000 words (Trakkr, 1,465 cited pages). For the length distribution broken down by page type, see rule 7 of the seven rules in A.3.
- In AI Mode, professional services win with educational content, consumer services win with review volume (PRWeb, 1,120 queries / 7 industries / 13 cities, does not include Singapore, 75% confidence).

> **Example** (home repair): among plumbers and electricians, the winners had 337% more reviews: 546 vs 125.

This winning mechanism cannot be turned into "so the strict side should go and ask for reviews too": on the strict side, reviews from consumer-type buyers can only be observed passively.

The statistical basis for retests

If you run 20 samples for one intent, a Wilson interval only lets you judge "it moved" once it is named **10/20** times; yet a baseline of 0/10 already has an interval upper bound of **27.8%**. Run it on this basis and week 13's default output is bound to be "not demonstrated", whether or not anything actually moved. The fix isn't to loosen the statistics — it's to change what counts as the acceptance measure: judge by firm, countable numbers (seat count, pages on the list, factual-error count) plus a measured noise band; the Wilson three-state test is internal reference only, used to trigger the not-moved triage, never a conclusion (rule in → General Edition 7.1 What a retest produces, and the rulers).

Question breakdown: the searches ChatGPT actually ran

- Sample: 225 stored ChatGPT answers (gpt-5.5 API with web search, Singapore location); no new queries were sent: 101 multi-industry (2026-09-14, 17 industries), 60 law firm (2026-09-29), 64 dental and aesthetics (2026-09-29). From each answer's web-search log we took the searches the model actually ran, then used Google's top 10 (Serper, gl=sg) as a proxy, looking up the original question and each search once.
- How many searches: for "who's best" questions, the median per answer was 16 multi-industry and 11 dental; for information questions, 8.5 multi-industry and 6 dental.
- Which search found the cited pages: 97 "who's best" answers cited 743 URLs; 49 (7%) were in the top 10 for the original question, 363 (49%) in the top 10 for one of AI's searches. Placebo control: swapping in the same number of searches from another question in the same industry hit only

46/736, against 314/736 for the question's own searches (the denominator 736 counts only answers that could be paired with another "who's best" question in the same industry; the question's own side is cut to the same number of searches, which is why it is 314, not the 363 above). A single search is no more accurate than the original question (on average 0.57 vs 0.51 cited URLs in each top 10); the searches win by being many and coming from many angles.

- Which words were added: of 995 multi-industry searches for "who's best" questions (one search can fall into several categories), 373 were about credentials and official sources, 238 price, 208 a year, 105 reputation, and 247 used `site:` to limit the search to one website (about 150 of them to one business's own website); 79 (8%) still carried best or top. In dental, AI swaps the patient's words for specialist terms (deep bite becomes orthodontist, periodontist).
- Dental and aesthetics, measured separately on 8 procedures × 4 phrasings (2026-09-29, each question asked of ChatGPT 2 times and AI Mode once): for a broad "best" question, 50/71 of the clinics named had their own website cited in the same answer, and 36/62 of the names overlapped with third-party lists for the same procedure (rough control: of the clinics in Google's top 10 for the same procedure's price and detail questions, 27/83 are also on these lists, and only 9/40 once the clinics AI named are removed); for a narrow "best" question (best plus a specific condition or group of people), ChatGPT 16/16 and AI Mode 8/8 gave a list, 76% of the clinic pages cited were deep pages such as doctor pages and procedure pages, and asking the same question twice gave a list overlap of 0.46, against 0.28 for broad "best".
- Limits: Google is only a proxy, and ChatGPT's search back end is not Google; only the ChatGPT API version was measured, so the searches AI Mode breaks a question into can't be seen; Singapore and English only, and each question was sampled on one day only; all of it is correlation, and there is no intervention experiment on whether adding pages gets you onto the list. Use it to set the categories and the build order of the question-breakdown table; never promise externally that you will get onto the list (main text in → General Edition 4.7 Actions by state, the ten question types and this quarter's slots).

Confidence list: what to print and what not to

Statement	Confidence	How to mark it
The noise band for seat count is measured every month (3 same-week retests); it is no longer unknown. Before the first measured value is in hand, seat counts are only listed side by side, never judged	Measured	**Print**, the noise band figure must be on page 1 of the monthly report
"Completing the five profiles → ChatGPT's local answers benefit" (from a public partnership)	70%	**Print**, on the "free, done in passing" basis; no result is promised

Statement	Confidence	How to mark it
"Improving your position on a list is cheaper than getting onto a new one"	Cross-industry aggregate basis; +16.5pp is from a B2B SaaS sample, not validated in the three local industries	**Print**, carry the qualifier unchanged
Splitting industries into professional and consumer services (each follows different fallback items) is based on PRWeb's 1,120-query sample, not including Singapore	About 75%	**Print**
Reviews get copied word for word into the answer body: recommendation type 67% (4 of 6), non-recommendation type (price / scenario / process / second opinion) 0% (0 of 9) (searched through a generic search API on 2026-09-21, n=15, SG not locked, single engine)	85%, very small sample	**Not printed externally.** Its use is to set the direction of the work: thick pages target price / scenario / process questions, and the review-copying frequency on that class of question is 0 — so thick pages are not paired with any review action, and the substitute for regulated industries is registration status, technical-standard certification and checkable-fact sentences. What the monthly report prints is your own measured values for your own M questions, split into two lines, recommendation type / non-recommendation type. The two lines must be printed side by side — printing only the recommendation-type line would lead people towards asking for reviews, something regulated industries simply cannot do
"60–70% of the 90-day increase lands on brand questions"	70%, tabletop projection, not measured	**Not printed**, internal reference
For one intent, "0 businesses named at answer level; several local pages are already fighting for citation slots" (searched through a generic search API on 2026-09-21, SG not locked, single engine, one-off)	85%, one-off sample	**Do not print the specific figure.** For internal ordering only; what the report prints is the named-business count and citation-slot ownership from your own baseline numbers, with the retest date and engine stated. No material may say "no one is doing this" or "competition is near zero"

The gain from getting onto a list: the one wording for the whole book

This data point may only be written this one way anywhere in the book; quote it word for word:

> Getting onto a list that AI already cites and ranking first on it = visibility **+16.5pp** and a position **0.8–1.8 places** earlier in the answer (Peec, 5.7M data points / nearly 200,000 answers / 8 engines; **+16.5pp is from a B2B SaaS sample; for emerging categories it is +13.4pp; not separately validated in Singapore dental, aesthetics or legal. Use it to order the work, never to promise any result**).

A.3 Page-type measurements (1): two legs, seven rules, the nine tier-A types

What you'll do in this section: check the basis for the two-engine round (526 citations, 92 pages), how often the seven rules hold, how the count for the 46 types adds up and the disputed calls, and the cited samples for each tier-A type.

What the two probes are called: "the two-engine round" = the 2026-09-23 five-industry probe, 60 questions, the ChatGPT and AI Mode legs, 526 citations, 92 pages taken apart one by one; "the ChatGPT-only round" = 5 industries × 10 questions, ChatGPT leg only, 241 citations.

How labelling works: measured n/N (N is how many pages of this type were taken apart, n is how many of them match), measured, 1 case (backed by only one cited page — follow it, but know it's a single case), not measured (inferred or an external rule, with no cited page to check against). **Two limits on the measured basis**: the seed keywords all come from Google Autocomplete (gl=sg); only the first 600 characters of each answer were stored, so for the second half of an answer there's no way to match each citation to its sentence.

Denominator: of the 526 citations, only **75** could be traced to the exact sentence copied on the original page (the rest either failed to fetch, or the copied sentence was a multi-paragraph rewrite with no single source). Below, any figure like "73%" or "about 55/75" has this 75 as its denominator, not the full 526; each item's position (e.g. "about 1% in" or "about 53% in") is a **visual estimate** based on how far down the page it sits, not character-level coding — treat it as a strong trend, not a precise measurement to put in a client report.

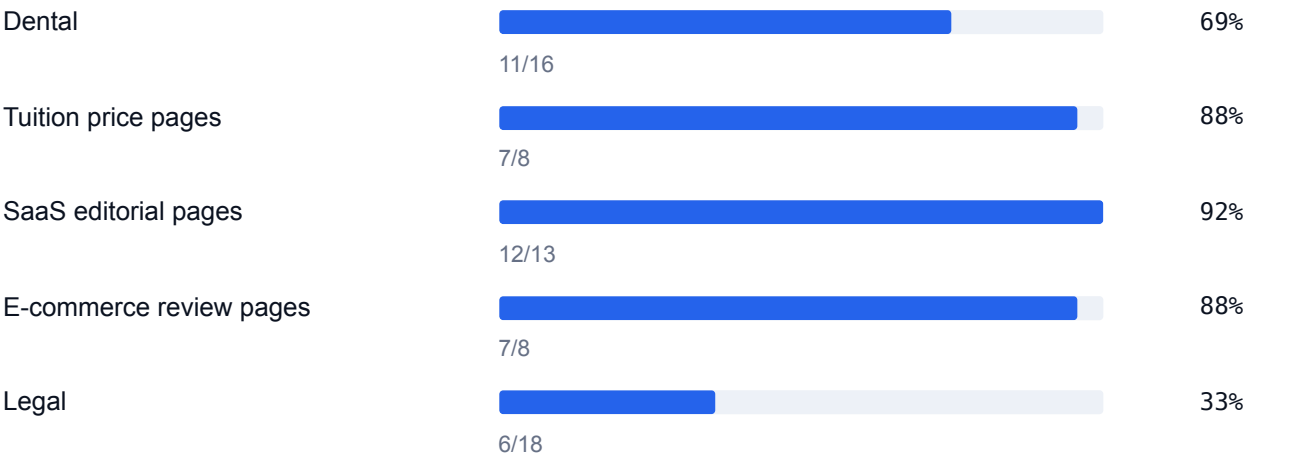


Figure: Share of cited pages that have a table (example: one line per industry, five industries, converted to percentages)

Tables are the carrier copied from most often; AI copies even the column headers straight across (examples of copied headers: "Part-time | Full-time | MOE", "Best for | Finish | Price"); on 6 of the 16 dental pages cited, what got copied was a table row. A table placed later on the page still gets copied (a government fee table at the 45–51% mark).

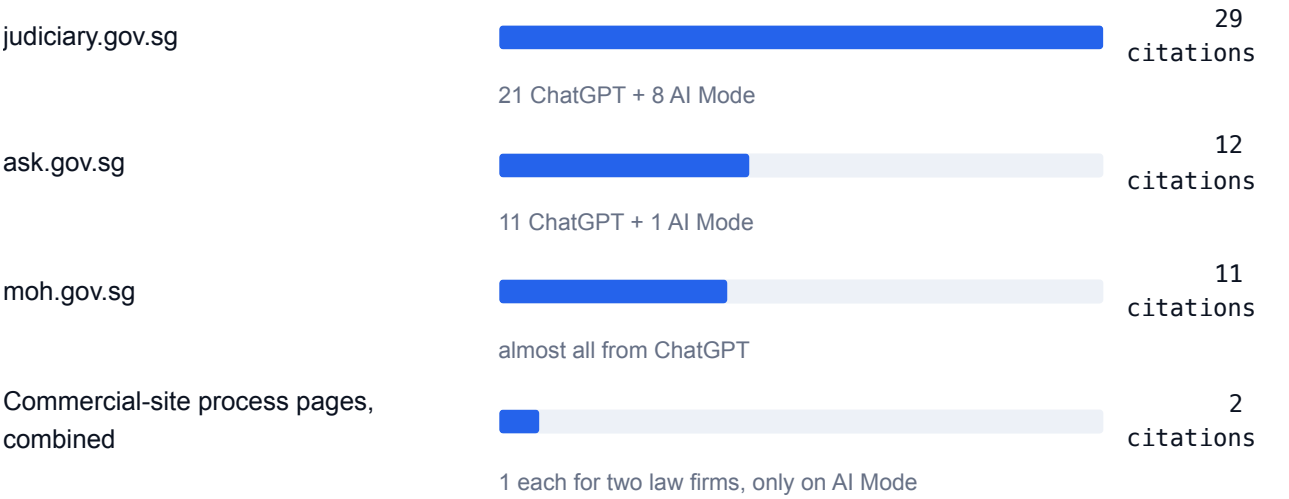


Figure: The most-cited Eligibility and process pages are official pages

The two legs differ: ChatGPT looks for the source, AI Mode for second-hand summaries (measured in five industries)

Industry	What ChatGPT mainly cites	What AI Mode mainly cites
Dental · aesthetics	Government and public institutions 76% (38/50: MOH fee benchmarks, CPF, SingHealth, SMC)	Commercial sites such as clinics 75% (39/52); for aesthetics "who's best" questions, goes through Google cards and booking product pages
Family law	.gov.sg 77% (46/60; judiciary 21 times, ask.gov 11 times)	Commercial sites such as law firms 69% (36/52)

Industry	What ChatGPT mainly cites	What AI Mode mainly cites
Tuition · education	Price questions cite peers' price-list pages; "who's best" cites organisations' own websites (about 35 of 44 URLs); judgement-type questions cite only government and media, 9/9	Price-list pages, lists, Reddit
B2B SaaS	Official pricing pages about 53%, plus about 7 citations of the help centre, then writes its own conclusion	YouTube, Reddit and LinkedIn together 40%, the rest split between peer comparison pages and lists
E-commerce · consumer goods	Product pages (PDP) 54%, category pages 21%, government pages 13% (used as the basis for a judgement)	Google Shopping cards 44%, YouTube 15%. URL overlap between the two legs is 0

The denominator for each percentage is that engine's citation count for that industry's questions; site type is classified by a domain-matching rule (recomputable from the public dataset, stats.md 3.1).

Inference 2: "who's best" intents are not won by writing pages, measured in five industries

When buyers ask "who's best / best clinic / best lawyer / best tuition centre", neither leg cites a business's own "best of" article. For where to put the effort instead, see → General Edition 5.6 Two legs, two kinds of sentence.

Industry	Where ChatGPT goes	Where AI Mode goes
Aesthetics, dental (n=1 question)	Public-hospital specialist pages, clinic procedure pages	2 Google searchviewer cards + 4 booking product pages
Family law (n=1)	Doyle's ×2, Legal500, Chambers, ask.gov, judiciary	Law firms' own websites, SassyMama, the Law Society directory, a map card
Tuition (5 questions)	Organisations' own websites, about 35/44; zero citations of lists	Mainly lists (singaporetuitionteachers, sethlui, tutorcity)
E-commerce	PDP 54%	Shopping cards 44%; for "vs" questions, YouTube 3/3
SaaS ("best CRM")	Official pricing pages 4/4	Reddit, YouTube ×2, business.com, the Slack blog

How often the seven rules hold

1. **What gets copied is almost always the first sentence of a paragraph**: legal 15/18, and the other four industries can all be matched to some paragraph's first sentence. Counter-example: on a law firm's prenup-agreement page, the H2's first sentence "is a complicated matter" wasn't copied —

what got copied on the same page was the FAQ entry that gives the conclusion directly. The strongest rule in the whole sample.

2. **The core fact falls in the first 30% of the page**: about 55 of the 75 locatable citations (73%); by industry: dental 12/16, legal 11/18, tuition 7/7, SaaS about 15/19, e-commerce about 10/15. Between the H1 and the first H2: 32/75 (43%); the first or second H2's opening sentence, or the first table: 23/75 (31%); the middle or later section: 17/75 (23%); in the FAQ alone, rarely. Position is a visual estimate, as noted at the start of this section.
3. **Whatever can be made into a table, make it a table**: frequency in the first figure at the start of this section.
4. **The title carries a year**: dental price blog posts 8/8, tuition price pages 7/8, SaaS editorial pages 12/13, e-commerce lists 5/6; exception: official pricing pages 0/7 carry a year; pages of the One question, one page type add one only when the answer changes over time (not measured).
5. **A statement H1**: cited clinic pages 10/11, tuition price pages 7/8, SaaS official pages 8/8; pages with ≥3 question-form H2s: only legal 3/18, SaaS 4/13, government process pages 0/4.
6. **The FAQ is rarely the source of the copied sentence**: pages with an FAQ — dental 11/16, SaaS 17/21, tuition price pages 7/8, legal 8/18; only about 6 cases had the copied sentence come directly from the FAQ; in most, either the body didn't have it or the FAQ just restated the conclusion.
7. **A precise number matters more than a clean layout**: mindstretcher's structure is messy, but its $272.50 is precise to the cent; a clinic's price page has 40 H2s, but the first one is a price list; singaporelegaladvice hasn't been updated in 4 years — all of these pages still get cited.

Three more rules that show up by page type:

- **Word count is not a criterion**: cited pages range from as short as 40 words (ask.gov), 62 words (a law firm's price list), 260 words (Talenox), up to as long as 9,000 words (MissLobang); the median is about 1,300–1,500 words; pages in the 1,800–3,000 word range: legal 2/18, tuition 2/8, e-commerce 3/11.
- **Price, results and specs must be in the first-screen HTML**: hard numbers on tuition-organisation pages, 6/6 readable straight from plain HTML. Counter-examples: axiom's S$420/month cannot be found in the homepage HTML, so ChatGPT had to look elsewhere; zoho switches currency by IP; mi.com's air-purifier-4 static HTML has no price.
- **Outbound links and a byline are not conditions for being cited**: dental pages with an authoritative outbound link, 1/16; legal, 2/18; a named individual author: legal 2/18, tuition price pages 1/8, SaaS official pages 1/8, PDP 0/5. Exception: medical pages — a named doctor in the page header, 6/12; a byline is worth doing there.

46 types: how the count adds up, and the disputed calls

```
The two-engine round (two legs, buying-decision questions)                    9 ty
New in the ChatGPT-only round (ChatGPT leg only, non-decision questions)     +33 ty
Absorbed from the external teardown                                           +5 ty
```

```
Merged as a duplicate (the teardown's first-party data report page folded
  into the Statistics source page as its self-built variant; not a new type)   -1 t
─────────────────────────────────────────────────────────────────────────────────────
Final                                                                          46 ty
```

46 is a count that comes from the data: the two-engine round's 60 seed questions were all piled onto buying-decision questions (how much / who's best / A vs B / can it be done); the ChatGPT-only round swapped the phrasing for "what is X", "how to … step by step", "what percentage", "what to bring", "is it legal", "near <MRT station>", "is X good" and "complete guide" — and the page-type count immediately jumped from 9 to 42. The number of page types isn't an industry trait; it's a function of how much question-phrasing ground you cover.

Two pages count as the same type only when all three merge rules hold: ① the triggering questions belong to the same family ② the position and form of the copied sentence on the page belong to the same family (a single first-screen sentence / a table cell / an H3 paragraph's first sentence / a chart data label / a footer NAP string / a numbered rule) ③ both can be written from the same block blueprint. Three disputed calls:

- **Type 13, Step-by-step procedure page ≠ type 7, Eligibility and process page**: type 7 answers whether you meet the conditions, how many times and how long, and what gets copied is the condition sentence; type 13 answers what to do at each step, which form to submit and how much the official fee is, and what gets copied is a cell in the `Step | Result` table.
- **Type 17, Subsidy and limit rules page ≠ type 7**: what gets copied is the amount cell in the limits table, not a condition sentence.
- **Type 22, Collected FAQ page ≠ type 6, One question, one page**: type 6 is one question per URL, with the first sentence giving a yes/no; type 22 packs 4–20 questions onto one page, and in the measurements the same page was copied from three different Q&A blocks, separately.

The ruler bias of the ChatGPT-only round: because it only ran the ChatGPT leg, the 33 new types it added skew towards the Primary sources family, and page types that lean on second-hand summaries can't be measured — this is a bias from the ruler, not the full picture (the reason and the retest plan are in A.4). Tallied by page type (formal run and rerun combined; not a citation count): 228 for the new page types, 63 for the nine known types.

The nine tier-A types: cited samples

n is written as "two-engine round → ChatGPT-only round".

Page type · n	Sample URL	Where it was copied from / which leg
① Single-service price page · 6 → 2	A dental clinic's Invisalign cost page	The first sentence after the H1 (about 1% in) + the first sentence of H2 "Invisalign or braces?" (about 53% in), cited by both legs; another example: a law firm's price list, 4 lines in all, cited by ChatGPT for a Chinese-language question. For this type, 5/6 were cited by ChatGPT only
② Price guide page · 16 → folded into 21 full-guide citations	A dental clinic's crown-cost blog post	Right under the first H2, a `Crown type | Cost (before GST) | Suitable for` table, cited by both legs, more by AI Mode; another example, tutorbee, is in the flagged samples below
③ Official pricing page · 31 → 14	https://monday.com/pricing	Copied from three places — the price card, the footnote, the FAQ; ChatGPT 27 : AI Mode 4, and AI Mode only cited it for brand-price questions (HubSpot price questions 3/3)
④ Comparison page · 20 → 1	https://ventureharbour.com/hubspot-vs-salesforce/	The "The short version" and "So which should you choose?" paragraphs, cited by AI Mode; AI Mode 17 : ChatGPT 3, and 2 of those 3 ChatGPT citations were third-party lab pages (RTINGS, TechRadar)
⑤ List page · about 29 → 2	https://www.misslobang.com/article/best-robot-vacuums-singapore-2026	"S$999.90 as of mid-September", cited by ChatGPT. AI Mode about 23 : ChatGPT about 6, and all 6 of those ChatGPT citations were third-party (G2, TechRadar, MissLobang, Doyle's ×2, Legal500); these 4 list domains, 4/4, state their methodology or a check date in the title or first screen

Page type · n	Sample URL	Where it was copied from / which leg
⑥ One question, one page · 23 → 6	https://ask.gov.sg/sgcourts/questions/clywqg7te006bq97hezb0sc5l	The whole page is 40 words in two sentences; the first sentence, "There is no legal requirement…", cited by ChatGPT. Government-page version: ChatGPT 20 : AI Mode 3; the 4 commercial-site single-question-page URLs (one law firm, one divorce-lawyer keyword site, two dental clinics) were all cited by AI Mode (n=4)
⑦ Eligibility and process page · 29 → 7	https://www.judiciary.gov.sg/family/understand-requirements-getting-divorce	The table of statutory facts × explanation × when you can apply, paraphrased row by row by ChatGPT; the same page was cited under 4 different questions; citation counts by domain are in the second figure in this section
⑧ Remedy and second-opinion page · 2 → —	A dental clinic's failed-implant page	The TL;DR box before the first H2 (the first 144 words), AI Mode's first citation; another example, an aesthetic clinic's page, was used by ChatGPT as its safety criterion. 2 of 2 answers discussed safety first; the legal-industry version was not measured
⑨ Entity anchor page · multiple pages → 10	https://axiomeducation.com/	The results card (at the 10–11% mark) cited number by number by ChatGPT; fees are not in the homepage HTML. PDPs were only cited by ChatGPT; the person profile page itself was not measured — what was measured was the byline in the page header (medical 6/12, legal 2/18, tuition price pages 1/8, PDP 0/5)

Three flagged samples:

- **A range sentence on a single-service price page** (a clinic's Invisalign page): "Invisalign at [the clinic] costs S$3,900 to S$7,900, depending on how complex your case is" — **a range: this breaks the rules on the strict side**. Do not copy this wording.
- **A law-firm page that writes its own price and a market range in the same H2, cited by both legs** (a law firm, measured, 1 case): "At [the firm], our fixed fee for an uncontested divorce ranges from $1,690 to $2,890." and "An uncontested divorce in Singapore generally costs between $1,500 and $3,500 (Source: Singapore Legal Advice)." — the rules do not say directly whether a law firm may put its own fees and a market range in the same place; the closest is the rule against comparing fees with other lawyers, so on the stricter reading we advise against copying it (Conservative line (not statute text)). This case only proves that both legs will cite this sentence shape (the copied sentences on three other cited pages have the same shape) — it is not a wording you may copy.
- **A price-guide page's by-level price table** (tutorbee): the card at the top says Primary $35–$55, the table says $25–$35 — AI skipped the card and used the table.

A.4 Page-type measurements (2): tier B and C samples, and the AI Mode retest list

What you'll do in this section: look up, type by type, the cited samples and n for each of the 33 tier-B types; tier C's statement that nothing has been measured; the priority order for retesting on the AI Mode leg; and the three things you must record alongside every run.

Statement of basis: tier A's data comes from the two-engine round (526 citations, 92 pages). Tier B comes from the ChatGPT-only round, run on the same day (2026-09-23): 5 industries × 10 questions, ChatGPT leg only (gpt-5.5 + web_search, Singapore location), 241 citations. Each type's n is a page-by-page count for that page type and combines the formal run with one extra run, so it is not the same as the number of citations. Tier C has no measurement of our own: its specs are taken, unmodified, from three files in an external teardown (the teardown target is a GEO agency's public blog). Every n = 1 finding is flagged, and it is directional only, never a settled conclusion.



Figure: The top 12 tier-B types by times cited

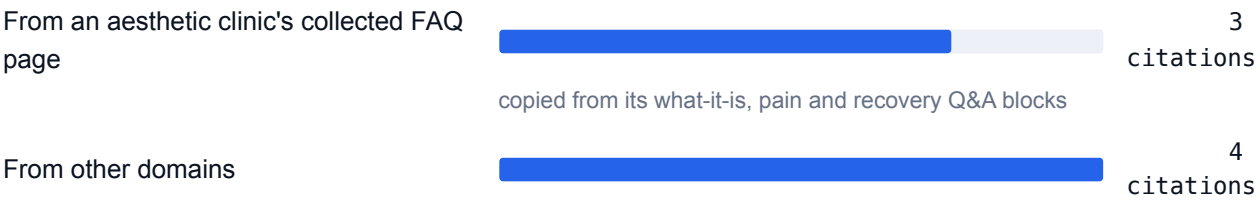


Figure: Of the 7 citations for one definition question, 3 come from three Q&A blocks on the same collected FAQ page

The question was in Chinese, asking what HIFU is and what its side effects are. This is the only self-built commercial page that ChatGPT cited repeatedly in the ChatGPT-only round (every other commercial citation was a price guide): one page supplied several answer fragments (3 of the 7 citations came from this one page; we have no comparison data for the hit rate of single-question pages, so no multiple is given).

The 33 tier-B types: n and key evidence

Page type	n	Key evidence
10 Regulatory obligations and penalties page	38	IRAS AIS's obligation sentence on the first screen was copied: "AIS employers are required to submit … by 1 Mar each year. Late submission may lead to a fine of up to $5,000."; for B2B software, the ten questions were run twice (the formal run and one extra run), and 59 of the 113 citations combined (52%) landed on .gov.sg; figure in → General Edition 5.23 Tier-B questions and rules types (1): regulatory obligations, definitions
11 Store / branch page	30	A dental clinic branch page's "Sun: 9.30am – 1pm" was copied from the end of the page; brizosystem's address line was copied from 90% of the way through the body text
12 Definition page	19	medical 10, legal 5, training 2, SaaS 2; all 10 citations of a US hospital's health library landed on the first sentence under an H3 question, and all were conditional sentences; in e-commerce, "HEPA vs true HEPA" got 0 citations, because ChatGPT never called web_search
13 Step-by-step procedure page	14	judiciary's `Step | Result` table was reordered into the answer; the $56 official fee was copied exactly
14 Preparation and bring-list page	12	3 of NDCS's items, grouped by patient category, were copied; the same question run twice gave 0 citations the first time and 8 the second, so stability is poor
15 Legislation text page	12	all 8 citations of sso.agc.gov.sg carried a `ProvIds` provision-number anchor; sso.agc itself returned 403 and could not be retrieved
16 Schedule and deadline page	9	commonwealthsec's date range on the first screen was copied; in one run, a single question cited 9 URLs of this type to cross-check them
17 Subsidy and limit rules page	8	MOH CHAS's limits (up to $830 a day, $240–$5,290 for surgical items; paraphrased) came from cells in the itemised limits table
18 Regulator guidance PDF	8	in PDPC's 62-page PDF, the Chinese question hit the line in the question-style table of contents that was phrased the same way
19 Statistics source page	8	the parallel restatement sentence in PubMed's CONCLUSIONS section (77.6%) was copied, not the raw data in the RESULTS section. This type overturns "core facts go in the first 30%": for statistical questions, ChatGPT digs all the way to the end of the page (the SingStat sentence sits 85% of the way through, the IMDA chart at 53%)
20 Register / approved list page	7	line 79 of the IRAS AIS PDF (at 84%) and the disclaimer at 2% were stitched into a single answer sentence; in another case, a vendor's row sat at 94% of the way through a table (ch5:2260)
21 Vendor legal terms page	7	the list of country names in the definitions section of HubSpot's DPA (at 17%); the `jst-asia-pacific` commitment sentence (at 73%)

Page type	n	Key evidence
22 Collected FAQ page	6	see the bars figure above
23 Official replies and speech records	5	the whole sentence in mlaw's parliamentary reply giving 4,150 of 6,220 cases, or 66%, was copied
24 Third-party single-business review	4	tutorly's at-a-glance fee table and its Pros/Cons were copied; rtings' static HTML is only 181 words of navigation, so it was listed as a citation but had no sentence that could be copied
25 Sentiment, forum and news pages	4	of the 5 citations for the question asking whether a dental chain is any good, 2 came from news and Reddit; neither page could be retrieved (the Reddit and mothership pages were both blocking pages)
26 Policy hub page	4	IRAS's question-style anchor structure = the AI answer's section structure; IMDA's three section pages are JS-rendered and could not be retrieved
27 Parameter and rate basis page	3	CPF's "It was gradually raised and has reached $8,000 in 2026." was copied; the calculator tool page itself got 0 citations
28 Review aggregate page	3	birdeye's first-screen "4.2/113 reviews" was copied; reviews.io's first-screen 4.9 does not match its JSON-LD 4.85
29 Self-built reputation and credentials page	3	a law firm's "[The firm]'s own family-law page shows 4.9 stars, with 3,455 Google reviews" was accepted, with the answer noting that it came from the firm's own page
30 Verification and lookup page	3	SDC's statement that every dentist must register before practising was copied; mlaw's register lookup page had 0 words copied and was only attached at the end of the answer as an exit link
31 Help centre / how-to page	3	the list of objects in the first paragraph of HubSpot's KB, at 13–14%, was copied; the "quotes do not sync" negative list, at 25%, was copied
32 Device and product regulatory documents	3	NEA's line that side effects are mostly temporary and ease in about a week was copied almost verbatim; the clinic's own self-written side-effects paragraph got 0 citations
33 Third-party directory listing	2	the first run copied its Sunday hours as 9am–9pm; in the second run of the same question, ChatGPT instead flagged it as unreliable and said to call to confirm (n = 1, directional only)
34 Change notice / old-vs-new page	2	judiciary's "replaces" sentence was copied; talenox's "View-only mode" transition-period clause was copied
35 Integration / marketplace listing	2	HubSpot's listing page is JS-rendered with 0 words of body text; the object list that actually got copied sits in the first paragraph of the KB page on the same site

Page type	n	Key evidence
36 Case-law page	2	a 689-word summary page and the 16,152-word full text were cited side by side; the holding sentence sits directly under the H1
37 Category list page	2	mothercare's whole sentence (169 models, S$112 to S$3,045, 9 brands; paraphrased) was copied; another site's whole page is JS-rendered, with 0 words of static body text
38 Calculator page (n = 1)	1	sgschoolkaki's tiered conclusion on the first screen and its comparison table were copied; the calculator component itself got 0 citations
39 Trust centre / compliance proof page (n = 1)	1	trust.hubspot's "SOC 2 Type II … June 23, 2026" was copied; not one sentence of the same vendor's 3,657-word marketing version was copied
40 Misconception page (n = 1)	1	healthhub's "When you buy and make payment for a vape, it is considered a purchase, which is also prohibited." and "Effective 1 May 2026, anyone caught owning, using and/or buying is liable to a fine up to $10,000" were copied
41 Verdict-first page (n = 1)	1	shopback's prices (S$1,299 non-HEPA / S$1,349 HEPA) were copied, with the source credited as ShopBack, not Dyson's own site
42 Time-limited promotion page (n = 1)	1	a dental chain's nett-price sentence (scaling + polishing + intraoral 3D scan for S$109 nett at participating clinics; paraphrased) was copied; the page itself was blocked by Cloudflare and could not be retrieved

The six n = 1 items in the table (38–42, plus the single observation in type 33) are directional only, not settled conclusions.

The four tier-C types: nothing measured

Page type	n	Note
43 Buyer's selection framework	0	the teardown target's version of this archetype cannot be found anywhere in the Google US top 30; the ChatGPT-only round had no matching question
44 Concept pillar page	0	adjacent to the ④ Comparison page and type 12, the Definition page, but with a different intent; to be verified
45 News and policy explainer	0	for rule-type questions, almost 100% of citations land on the original .gov.sg page; explainers on commercial sites get 0 citations
46 Market observation page	0	the teardown target's own 5 Type=Market articles are the only internal-link dead ends on the whole site (0 links out), and they are not even in the `/blog` index

The teardown target's own pages of these types do not even rank organically. Its raw ranking data shows 8 of 31 non-brand terms in the top 10 and 5 at #1, all of them list pages; its tool-comparison pages and single-competitor review pages cannot be found in the Google US top 30. The teardown target measured organic rankings and we measured AI citations: numbers from the two rulers cannot be combined.

Three conclusions that recur across page types: ① AI does not click calculators. For "calculator" or "how do I calculate" questions, ChatGPT did not cite a single interactive calculator tool, and this repeated across three industries. It cites the static table next to the calculator, the tiered conclusion sentence on the first screen and the authoritative source for each parameter in the formula, then does the arithmetic itself (shared by types 17 / 27 / 38). ② ChatGPT looks for the source, AI Mode for second-hand summaries: full evidence in A.3. ③ The teardown target's market-observation articles are internal-link dead ends: see type 46 in the table above.

Waiting for the AI Mode leg

During the ChatGPT-only round, the AI Mode leg's data interface (SerpApi) had run out of quota, so the AI Mode leg was not run at all. All 33 tier-B types therefore have evidence from the ChatGPT leg only. "Only the ChatGPT leg has been measured" does not mean AI Mode does not cite them; it only means they have not been measured. The known direction of the bias: ChatGPT looks for source pages and AI Mode for second-hand summaries, and in e-commerce the two legs' URL overlap was 0. The 33 newly added types skew systematically towards the Primary sources family, and page types that lean towards second-hand summaries could not, by construction, show up in the ChatGPT-only round.

Priority	Page type	Expected difference	Why we expect this
High	24 Third-party single-business review	AI Mode should be markedly higher than ChatGPT (only 4 citations in the ChatGPT-only round)	In the two-engine round, the ④ Comparison page was AI Mode 17 : ChatGPT 3 and the ⑤ List page 23 : 6; reviews and lists are AI Mode's home ground
High	28 Review aggregate page, 29 Self-built reputation and credentials page	AI Mode may go to Google cards / map cards instead and not cite these two types	In the two-engine round, for "who's best" questions AI Mode went to searchviewer cards + booking product pages
High	25 Sentiment, forum and news pages	AI Mode should be higher (all 4 citations in the ChatGPT-only round depended on ChatGPT digging to page 3 / page 23 of Reddit)	In the two-engine round, AI Mode's Reddit + YouTube + LinkedIn made up 40% combined for SaaS

Priority	Page type	Expected difference	Why we expect this
High	12 Definition page	The cited pages may change wholesale: ChatGPT cites the source, while AI Mode may switch to peers' long educational articles	The cleanest comparison point for the source-vs-second-hand split
Medium	37 Category list page, 41 Verdict-first page	AI Mode may switch to citing Google Shopping cards and YouTube	In the two-engine round, AI Mode in e-commerce: Shopping cards 44%, YouTube 15%
Medium	14 Preparation and bring-list page	Stability itself needs retesting (the same question run twice gave 0 → 8 citations)	Whether checklist pages win citations swings wildly
Medium	38 Calculator page	Test whether "AI does not click calculators" holds on both legs	It repeated across three industries in the ChatGPT-only round (n = 3 questions), but all on the ChatGPT leg
Medium	19 Statistics source page, 23 Official replies and speech records	AI Mode may use media paraphrases instead of the primary source's own figures	If so, "build your own statistics page that states n = and its basis" needs a different way to win on the AI Mode leg
Low	10 Regulatory obligations and penalties page, 15 Legislation text page, 18 Regulator guidance PDF	Both legs are expected to cite them, with little difference	Legal facts have only one source; a second-hand summary cannot replace it
Must test	43 Buyer's selection framework, 46 Market observation page (tier C)	The ChatGPT-only round cannot measure them, by construction	Both are second-hand-summary / discursive forms, so only the AI Mode leg can give the first piece of evidence. If type 46 still gets 0 citations on the AI Mode leg, it can be closed out and not built

Whatever question you run, record these three things alongside it:

1. **The full answer text**: both probe rounds stored only the first 600 characters. Next round, save the whole answer; otherwise, in a long answer, "which sentence got copied" can only be guessed.
2. **Rerun the same question**: strong swings were observed for type 14 (checklists) and types 24/28 (reviews). Run each question ≥2 times and record the difference in citation counts between the runs.
3. **Fill in the pages that could not be retrieved**: the ChatGPT-only round flagged 7 places as "could not be retrieved" (sso.agc.gov.sg 403; a dental chain's page blocked by Cloudflare; ishopchangi.com / ecosystem.hubspot.com / IMDA section pages JS-rendered; PubMed's cookie wall; G2/Capterra

403; mothership/Reddit blocking pages). Next round, render them in a browser before breaking them down; this is exactly why type 42's format spec is currently empty.

A.5 List of sources

What you'll do in this section: look up, entry by entry, the original sources for the evidence and rules used across the whole book: full name, URL, date retrieved.

This section lists only general sources: consumer-protection and price-transparency guidelines, do-not-call guidelines, destination-country advertising law, platform documentation and data research. Each industry regulator's own advertising regulations, FAQs and register lookup pages belong in that industry edition's appendix, not here. Where the original gives no URL or date retrieved, write "not given" or "not stated" as it stands; do not fill it in.

Consumer protection and price-transparency guidelines

- CCCS · Guidelines on Price Transparency: `cccs.gov.sg/consumer-protection/legislation-and-guidelines/guidelines-on-price-transparency` (date retrieved not stated in the original)
- The CPFTA statute (Consumer Protection (Fair Trading) Act): `sso.agc.gov.sg/Act/CPFTA2003` (date retrieved not stated in the original)

Do-not-call guidelines

- PDPC · Do Not Call Registry and Your Business: `pdpc.gov.sg/overview-of-pdpa/do-not-call-registry/business-owner/do-not-call-registry-and-your-business` (date retrieved not stated in the original)

Destination-country advertising law

- China's *Advertising Law* (2021), English translation: `chinalawtranslate.com/advertising-law-2021/` (date retrieved not stated in the original)
- SAMR's *Enforcement Guidelines on Absolute Terms in Advertising* (no URL given in the original; date retrieved not stated in the original)

Platform documentation

- Gartner Peer Insights · Incentives FAQ: `gpivendorresources.gartner.com/en/articles/6812506-incentives-faqs` (date retrieved not stated in the original)
- Gartner Peer Insights · Review Sourcing FAQ: `gpivendorresources.gartner.com/en/articles/6812574-review-sourcing-faqs` (date retrieved not stated in the original)
- OpenAI merchant feed spec: `developers.openai.com/commerce/` (date retrieved not stated in the original)

- OpenAI merchant portal: `chatgpt.com/merchants` (date retrieved not stated in the original)
- Google, *AI features and your website*: eligibility for AI Overviews and AI Mode, Merchant Center and Business Profile listed as best practice, no extra schema needed; original wording in A.1 (no URL given in the original; read on 2026-09-21)
- Google's crawler documentation (what `Google-Extended` is for, and the original wording "does not impact a site's inclusion in Google Search"; see A.1); Google's 2026-05-15 AI-optimisation guide (`llms.txt` not needed) (no URL given for either; date retrieved not stated in the original)
- Apple, *About Applebot*: `Applebot` serves Spotlight / Siri / Safari retrieval, `Applebot-Extended` is only a training opt-out switch (no URL given in the original; date retrieved not stated in the original)
- Anthropic's official documentation: `ClaudeBot` / `Claude-SearchBot` / `Claude-User` all respect robots.txt (no URL given in the original; date retrieved not stated in the original)

Data research (sample sizes, time periods and qualifiers are all in the matching entries in A.1–A.4; this is only an index and does not repeat the numbers)

- Citation density and page types: deltaV, 25,337 citations (A.1)
- Earned media share: Muck Rack, 25M links / 17 industries (A.1)
- Zero JS execution: Vercel × MERJ, 500M+ fetch samples (A.1)
- Share of price queries that trigger AI Overviews: Whitespark, 540 queries (A.1)
- The gain from getting onto a list: Peec, 5.7M data points / nearly 200,000 answers / 8 engines (A.2)
- Skew between the citation distributions of the web leg and the API legs: University of Hamburg + Leibniz Institute, 2025-11, 24,000+ answers / five weeks (A.2)
- Video correlation coefficients: public correlation studies, namely Ahrefs 2026, 5W Citation Share and Otterly's 1.7M data points (A.2)
- The rest, which have only the researcher's name and no URL in the original: searchVIU, Ahrefs (JSON-LD tracking, Wikipedia share), Steady Demand, 5WPR, Pew, Omniscient, SE Ranking, BuzzStream, Profound, Peec (`result_source` capture) (all in A.1); PRWeb, Trakkr, Ahrefs (listicle share), Otterly / Axios (Reddit and source-reshuffle tracking) (all in A.2)

Appendix B

Templates: robots.txt, wording, outreach letters and tables

Copy-ready masters, what to tell clients, outreach and review replies, and work-order headers

B.1 Door templates

B.2 Approved wording and scripts

B.3 Outreach emails and review replies

B.4 Headers for work orders, ledgers and monthly reports

B.5 Blueprints for unverified page types (for industry editions to include)

B.1 Door templates

What you'll do in this section: Copy the robots.txt master template, the crawler hit table, the read-only six-line robots record and the platform-by-platform curl commands straight from here, replace the placeholders, then use them.

robots.txt master template (12 retrieval groups + 6 training groups)

This is the only master template in the whole book. The person who edits the site copies it straight from here. `<N lines of Disallow>` is a placeholder: **it must be replaced with the N original lines you copied in step 1 of → General Edition 2.6 Fixing the door: nosnippet, the four-step robots merge, the WAF allowlist, not one line short, in every single group**.

```
# —— Retrieval crawlers (12): allow them. They have to read you before you can even
# [Note] Every group below must carry the original wildcard group's N Disallow lines,
#    otherwise these crawlers fall completely outside the wildcard group and none o

User-agent: OAI-SearchBot      # The crawler ChatGPT uses to fetch search material.
Allow: /
<N lines of Disallow>

User-agent: ChatGPT-User       # Live fetch when a user opens your page inside Chat
Allow: /                       # [Note] This line is a declaration, not a switch. Ope
                               # "Because these actions are initiated by a user,
                               # robots.txt rules may not apply". The real switch i
                               # the IP ranges from openai.com/chatgpt-user.json to
<N lines of Disallow>

User-agent: OAI-AdsBot         # Only needed if you run ads; delete the whole group
Allow: /
<N lines of Disallow>

User-agent: Bingbot            # Bing is one of ChatGPT's eight retrieval sources (
Allow: /
<N lines of Disallow>

User-agent: Googlebot          # Google Search, AI Overviews and AI Mode share one
Allow: /
<N lines of Disallow>

User-agent: Applebot           # The Siri / Spotlight / Apple Maps / Apple Intellig
Allow: /                       # and one of the very few crawlers that execute JS.
<N lines of Disallow>          # [Note] Don't claim your Apple Business Connect profi
```

```
User-agent: Amazonbot          # Alexa / Amazon retrieval
Allow: /
<N lines of Disallow>

User-agent: PerplexityBot
Allow: /
<N lines of Disallow>

User-agent: ClaudeBot          # All three Anthropic bots honour robots.txt (unlike
Allow: /
<N lines of Disallow>

User-agent: Claude-SearchBot
Allow: /
<N lines of Disallow>

User-agent: Claude-User
Allow: /
<N lines of Disallow>

User-agent: DuckAssistBot
Allow: /
<N lines of Disallow>

# —— Training crawlers (6): nothing to do with "whether you can be cited" — decide
# All of them can be changed to Disallow: / without affecting any retrieval leg abo
# To allow one, also carry the same N lines of Disallow

User-agent: GPTBot             # OpenAI's training-corpus crawler, not used for ret
Allow: /
User-agent: Google-Extended    # Controls Gemini training and grounding inside Gemi
Allow: /
User-agent: Applebot-Extended  # Controls Apple training; doesn't affect the Appleb
Allow: /
User-agent: Meta-ExternalAgent
Allow: /
User-agent: CCBot              # Common Crawl
Allow: /
User-agent: Bytespider         # ByteDance
Allow: /

# —— The Sitemap line belongs to no group and must go at the end of the file ——
Sitemap: https://<your-domain>/sitemap.xml
```

`Perplexity-User` is not in the master template: like `ChatGPT-User`, it is triggered by a user action and robots.txt cannot control it. The only way to allow it is to add its official IP ranges in the WAF. For

the original wording and the link, see → General Edition 2.6 Fixing the door: nosnippet, the four-step robots merge, the WAF allowlist.

curl commands to check robots platform by platform

Put this in the *First work order*. When off-site citation slots are spread across several third-party platforms, run it once for each platform; if a platform's key pages sit on a different host (a separate subdomain, third-party hosting), run it again on its own for that host.

```
for d in <platform1-domain> <platform2-domain> <platform3-domain> <other-platform>;
  echo "=== $d ==="
  f="./robots-$d-$(date +%F).txt"
  curl -sL -m 20 -A "Mozilla/5.0" "https://$d/robots.txt" -o "$f" \
    || { echo "  not retrieved (connection failed) → log it as 'undetermined'; do n

  echo "--- ① Per-leg named group: judge listed legs by their own group; legs not
  grep -inE -A3 "^user-agent:[ [:space:]]*(OAI-SearchBot|ChatGPT-User|OAI-AdsBot|GP

  echo "--- ② What paths the wildcard group blocks: check whether your key page pat
  tr 'A-Z' 'a-z' < "$f" | awk '/^user-agent:[ \t]*\*[ \t]*$/{f=1;next} /^user-agent
done
```

The four columns of the crawler hit table

Column	What to fill in
Crawler name	Judge by IP-range ownership, not by the UA string
Hit count	30-day total
Top 20 hit URLs	Sorted by hit count, descending
Status-code breakdown	Count of each: 200 · 301 · 403 · 429 · 5xx

The read-only six-line robots record (blank headers)

For what each of the six lines checks and how, see → General Edition 2.2 Gate 1 read-only audit: what to check in nosnippet, robots and the WAF; this is only the skeleton of the table — fill it in as you read that section.

What to check	How to check it	What to record	How often

What to check	How to check it	What to record	How often

B.2 Approved wording and scripts

What you'll do in this section: For external use, for clients and internally, pick the line that matches the situation and say it exactly as written. Do not improvise. Where an entry also gives the sentence you must never say, attach that too.

Check the table below for your situation first, then jump to the matching entry:

When	Which one to use	Who can use it
Making a PDF, screenshot or slide for internal eyes only	Internal material	All sides
Telling someone outside how AI currently mentions you	Three approved lines for outside use ①②	All sides
Telling someone outside which items you skip because of compliance	Three approved lines for outside use ③	Strictly regulated side only
A client asks why the price page doesn't show a range	On pricing	① all sides, ② strict side only
A client asks whether to pile on more schema, or whether completing business profiles is worth it	On structured data, On the five business profiles	All sides
A third-party page has you wrong, and the client wants AI to say it differently	How to write a correction letter, Red lines in what you say	All sides
Signing the frozen question pool, the compliance officer reviewing a draft	When signing the frozen question pool, For the compliance officer	All sides
Writing the monthly report and the work order	On this quarter's outreach, Two off-site disciplines, On the retest noise band	All sides

Internal material

When to use: Any material for internal eyes only that names third-party organisations or contains competitor data (PDF, screenshot, slide).

Wording (burn it into the PDF or image itself, not the body of an email; add an `–INTERNAL` suffix to the file name):

> Internal reference material · Not to be published, posted, forwarded or used for any promotion. This page contains third-party organisation names and counts produced by our own measurement method, and is not an assessment of any organisation's merits.

Why this comes first: In industries that regulate peer comparison and advertising, material that compares you with peers runs straight into the rules on comparison and disparaging peers, and on third-party promotion, the moment it goes up at the front desk or into a public channel (rule text in each industry edition's appendix). And in Singapore, the bar for a named peer to file a complaint is low.

Never do this: The first time you show someone in person how AI currently talks about you, that screen often carries a recommendation-type answer that names peers. Only point at the screen and talk it through: no screenshot, no saved copy, no PDF, no copy kept.

Three approved lines for outside use

① On seats

> "I asked K times and AI named us X times. That count comes from this frozen set of questions, this engine and this number of passes. In the same week I re-ran it 3 times and it varied by ±N seats, so only a change larger than ±N counts as real movement."

Never say: "We rank number X on AI", "Our visibility score is XX", "Users will see us in ChatGPT" (banned whenever the only evidence you have is from the API leg).

② On the two legs

> "OpenAI API leg X seats, Gemini model leg Y seats, Google AI features Z impressions: three different systems, three lines side by side, never added together."

③ On the items you skip for compliance (strict side only)

> "These are the things I choose not to do (<name each one; common examples: not proactively requesting reviews, not buying a spot on a Best / Top N list, not writing a 'from' price>), because doing them would break the rules."

This is an explanation, not a release from liability: **responsibility was never transferable in the first place**. Inside the brackets, name only the items that Statute text on your side bans. An item labelled

Conservative line (not statute text) must never be described as "doing it would break the rules".

On pricing (two sentences, then stop)

When to use: A client asks "why doesn't the price page show a price / a range / the word 'from'".

> ① "Even if you don't publish a price, AI quotes one for you. We just asked it, and the figure it gave is not yours." (Say this pointing at the screen, using the number you measured and screenshotted yourself and the source URL it cited. **Do not cite any industry percentage.**)
> ② "The regulator only allows one fixed figure, so the price page will have no range and no 'from'." (**Strict side only**; look up the rule by industry — for the rule text and its number, see each industry edition's appendix.)

Stop there — never add a third sentence: "which happens to be exactly the price format machines read best", "you're already the only one whose whole line can be copied", "so AI will copy you". Delete the whole sentence; it must not appear in any variant. The public evidence for how price format affects extraction rate is inferred, not measured: the study itself says it did not measure differences in extraction by price format.

On structured data

When to use: A client asks "should we pile on more schema".

> Structured data is an industry baseline, not a lever. Ahrefs tracked 1,885 pages that added JSON-LD: after adding JSON-LD, AI Overviews (AIO) citations changed by −4.6%, the only significant result, and it points negative; the other two platforms were not significant (AI Mode +2.4%, ChatGPT +2.2%). Seal it in half a day and spend your time elsewhere.

For details, see → General Edition 2.7 Gates 3 and 4: indexing paths, and JSON-LD sealed once.

On the five business profiles

When to use: A client asks "is it worth completing these five business profiles".

> "Completing these five profiles is a basic step, done in passing."

Never say: "Complete them and ChatGPT will recommend you". The mechanism for the three (Foursquare, Apple and Yelp) is held at only **70% confidence**; Steady Demand's own researchers also say plainly that **there is no data isolating which profile field drives citations**.

How to write a correction letter

When to use: You find a fact wrong on a third-party page and need to send a correction letter.

> Point out only the error, plus a link to a checkable source (pointing to `/facts`) — no mention of ranking, no promotion tacked on.

Red lines in what you say

When to use: You're asked "can you make AI say this about us".

Never say: "I'll make AI describe us this way."

> "I'll write this sentence on our own page, in a form machines can read. Whether it gets copied is up to the machine; whether it's written correctly is up to me."

When signing the frozen question pool

When to use: Before a client signs the frozen question pool — read it to them sentence by sentence, change it on the spot, sign it on the spot.

> "These 30 questions are a ruler, and they're also an asset — the original answers, screenshots and ledger are yours from day one, and you can take them any time. Once signed, they don't change for the whole quarter, because changing a question is changing the ruler, and once the ruler changes, the before-and-after numbers can no longer be compared. That's the same for everyone."

For the compliance officer

When to use: Explaining to a client why some cells on the page don't exist, and who makes the call.

> "We didn't set these rules — the regulation did; we're only translating it into layout: which cell can exist, and what can go in it. Your compliance officer crosses out what shouldn't appear, and we do the other half: fill that cell back in with what Statute text allows. A page that only deletes and never fills in is safe, but it brings nothing. The final call is always your compliance officer's; all we guarantee is that the draft they receive isn't a blank page."

This sentence must always be followed by this bottom line: in three situations we stop and escalate to your compliance officer for a written sign-off: ① facts only you know; ② judgements the Playbook labels Conservative line (not statute text); ③ entries the Playbook labels Original text not obtained. Item ③ is not released even with a sign-off: write the Playbook's default wording; it is not used as a basis for judgement; assume neither that the rule exists nor that it does not. **A written sign-off resolves only these three situations and can loosen only ① and ②; it never turns something banned as Statute text into something publishable.** For how each of the three is handled, and which ones a

sign-off can loosen, see → General Edition 0.4 The three labels for compliance sentences, and the stop-and-escalate rule.

On this quarter's outreach

When to use: Telling someone outside how much outreach you did this quarter. Just this one sentence — don't extend it.

> "This quarter we sent to every reachable slot in Group A and Group B — how many letters, which day each was sent, whether there was a reply, whether it went live, all logged line by line in the *Outreach ledger* and reviewed once a month."

Never say: Don't give an expected reply count (e.g. "a batch gets about 2–5 replies") — that number comes from an inference at 55% confidence, and saying it out loud means making a promise with a number you yourself labelled "not measured". Also never dress up a low letter count as "precision targeting" or "quality over quantity". How many letters go out is an internal scheduling call; explain it internally and leave it there.

Two off-site disciplines

When to use: Writing the work order and the monthly report — both disciplines must be spelled out for the client.

> "Off-site asset maintenance: during the service period, recheck every off-site asset once a month (business profiles, directory listings, list spots, review count and star rating), and list each one in the monthly report as 'on the list / dropped out / information out of date'."
>
> "Paid-listing disclosure: any off-site listing obtained through a paid placement must be marked 'paid listing' line by line in the work order and the monthly report, with the cost shown separately. We do not buy reviews, organise fake recommendations, pass anything off as an independent review, or use a platform that charges per enquiry or on a commission-share basis."

On the retest noise band

When to use: Reporting a change in seat count in the monthly report — this sentence must be printed on the report exactly as written.

> "The same questions, the same engines and the same number of passes, run several times in one week, give different seat counts on their own. We measure that difference every month and print it below (this month's noise band ±N seats). A monthly change in seats is written as 'increased' only when it is larger than this band; otherwise we write 'within normal variation, cannot be determined'."

B.3 Outreach emails and review replies

What you'll do in this section: Copy the main letter (English + Chinese translation), the swap-in paragraphs (update letter / correction letter / removal request), the follow-ups at +7 and +21, and the review reply template; before sending, clear the send gate in → General Edition 6.2 The send gate, paid listings and the citation-slot table; on regulated sides, insert the optional paragraphs where marked.

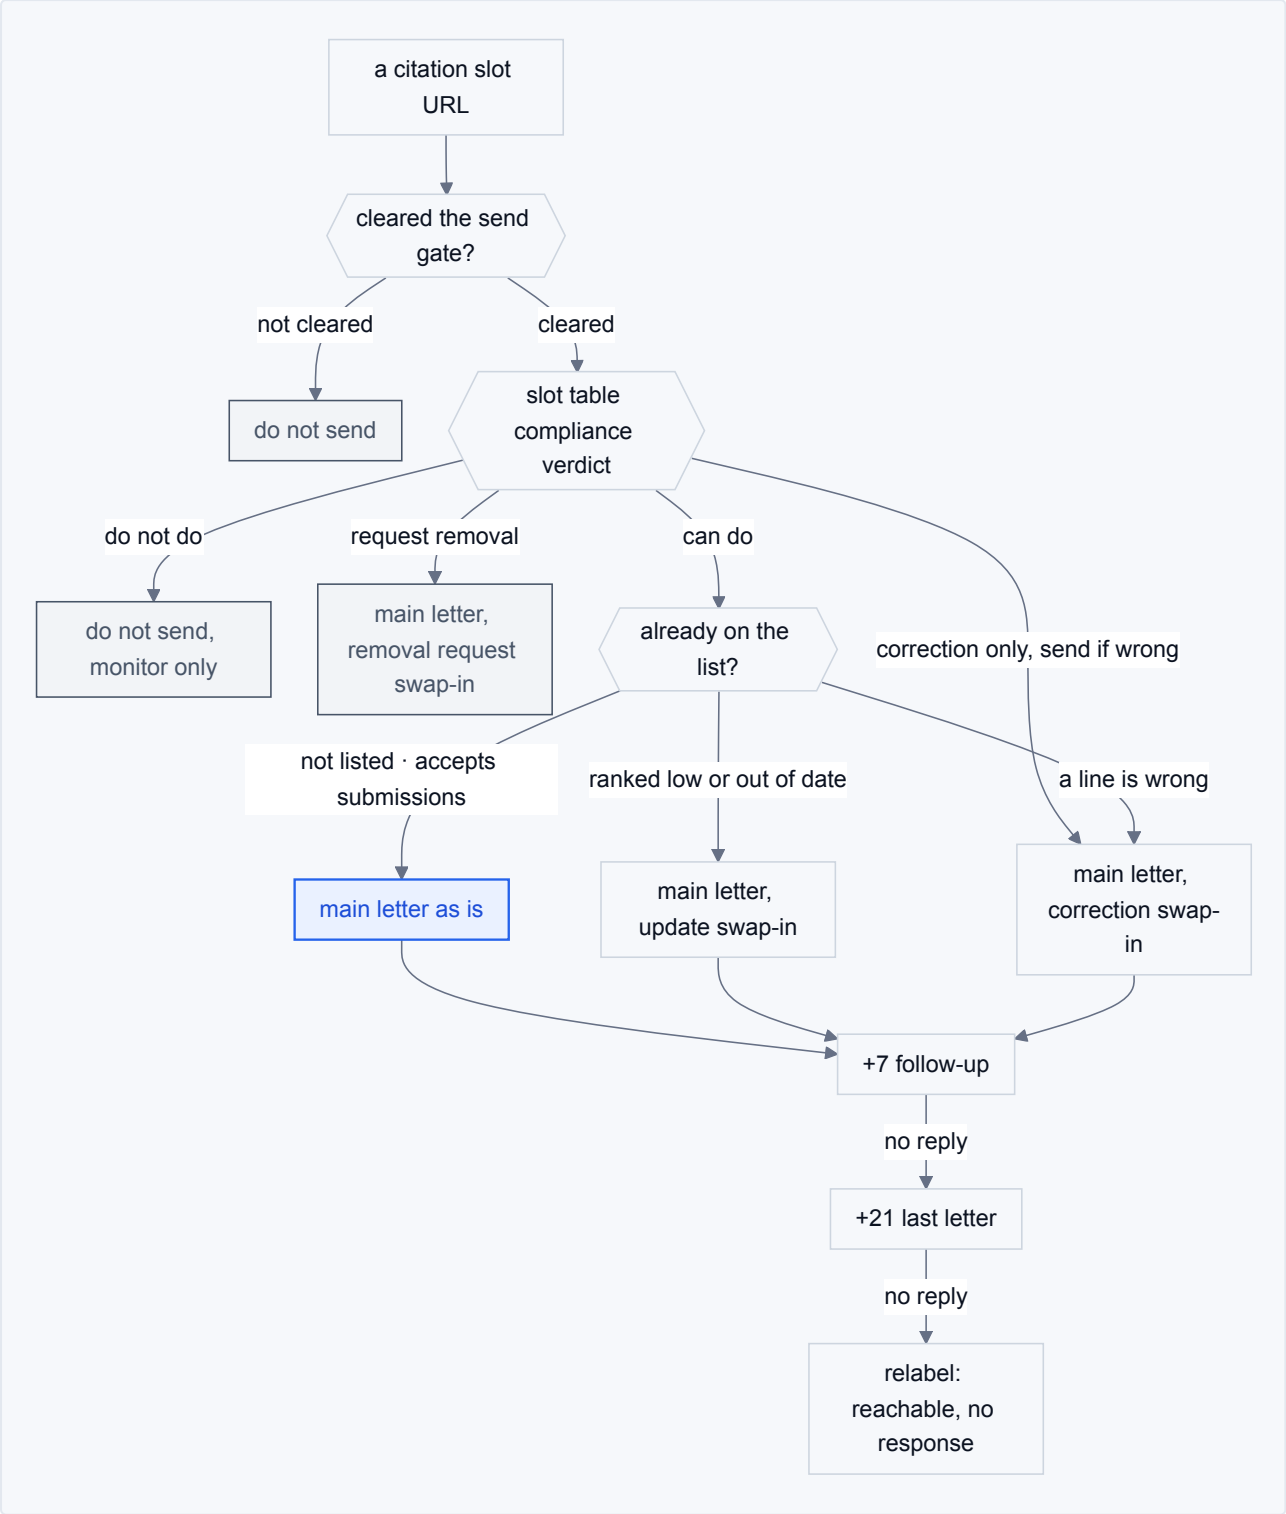


Figure: clear the send gate and the compliance verdict first, then pick a letter by whether it's already listed; all three letter types get only two follow-ups, at +7 and +21

For which groups to send to, and which "three items" Group A needs complete, see → General Edition 6.3 Six dispositions, bylined contributions and outreach letters; this section only gives the letters themselves.

Main letter (English, default)

Subject: Listing / editorial enquiry — [Business Name], Singapore ([category])

> Hi [Name / "Editorial team"],
>
> I'm [Your Name], [role] at / authorised representative for **[Legal Name]**.

Insert for regulated sides (optional, appended to the end of the previous sentence): `, a licensed [<licence category>] in Singapore ([<registration-number type and number>])`.

> We came across your page **[exact URL]** — it's one of the pages Singapore buyers and AI assistants surface most often for "[query]". Two questions:
>
> **1.** Is there a process to be considered for that page — editorial submission, a directory listing, or a paid placement? If there's a rate card or media kit, could you send it? **Could you also confirm how placements are charged — a fixed rate, or on a per-enquiry / commission basis?** (We can only work with fixed-rate placements.)
>
> **2.** What materials do you need? We can provide, on request: • registered entity name and UEN • addresses and opening hours • our published prices for [service], itemised and matching our own pricing page line for line • high-resolution photos we own the rights to • a factual 80-word description

Insert for regulated sides (optional, add to the list above): `• [role]'s name with their registration number`.

Insert for regulated sides (optional, whole paragraph):

> Two things we will not send, so you know up front: **no testimonials, no success rates, no before/after images, no comparisons against other practices, and no superlatives such as "best", "top" or "leading" — including in any headline or list title that carries our entry.** Singapore's advertising rules for our sector don't allow them, and we'd rather you hear that from us first.

> Could you let me know the process, the charging basis, any cost, and your usual lead time? Happy to work to your format.
>
> Best regards, [Name] · [email] · [phone]

Main letter (Chinese version, translated)

Subject: “收录 / 合作咨询 —— [机构名称]（新加坡 [行业]）”

> “您好 [称呼 / 编辑部]：
>
> 我是**[法定名称]**的 [职务] / 受其书面委任的对外发布负责人。”

Insert for regulated sides (optional, right after the previous sentence): `（新加坡持牌 [<持牌类别>]，[<注册号类型>]：[填]）`.

> “我们注意到贵站的这一页：**[完整 URL]**。它是新加坡买家和 AI 助手在搜「[查询词]」时最常出现的页面之一。想请教两件事：
>
> **一、** 本地机构如果想进入这一页，贵站的流程是什么——编辑投稿、目录收录，还是付费刊登位？如有刊例或媒体资料包，烦请一并发来。**另请确认计费方式是固定刊例价，还是按询盘 / 佣金分成**（我们只能采用固定刊例价）。
>
> **二、** 需要我们提供哪些材料？我们可以提供：注册实体名称与 UEN；地址与营业时间；[服务项目] 的**逐项公开价格**，与官网价格页逐行一致；我方拥有版权的高清照片；一段 80 词的事实性简介。”

Insert for regulated sides (optional, add to the list above): `从业人员姓名及其注册号`.

Insert for regulated sides (optional, whole paragraph):

> “有几类内容我们不会提供，先说清楚：**客户证言、成功率数字、前后对比图、与同行的比较，以及「最好 / 首选 / 领先」一类的褒扬用语（包括承载我方条目的标题或榜单名）**。新加坡本行业的广告规则不允许。”

> “烦请告知流程、计费方式、费用与一般上线周期。
>
> 此致 [姓名] · [邮箱] · [电话]”

Swap-in: update letter (already listed, but ranked low or described out of date)

When to use: The entry is already on the page, but the price or credential information is out of date.
Replace points 1 and 2 of the main letter; leave the rest unchanged:

> [Name] is already listed on that page. Two details have changed since it was published, with sources attached: • **Pricing:** [service] is now **S$X, inclusive of GST**, covering [items]; [excluded items] are priced separately — see [pricing page URL], updated [date].

Insert for regulated sides (optional): `• Credentials: [Name] is registered with [<registering body>], registration no. [XXXX] — verifiable at [official register URL].`

> I've also attached a refreshed 80-word description and two photos we own. **No change to the ranking or to the other entries is being requested** — just current facts.

Swap-in: correction letter (wrong information)

When to use: One fact on the entry is wrong; correct only that one line. **Replace points 1 and 2 of the main letter; leave the rest unchanged**:

> One line on that page is no longer accurate. Current text: "[quote the exact sentence]". The correct position is [one sentence], verifiable at [official source URL]. **I'm not asking for any change to the ranking, to the list itself, or to the other entries — only that one line.** Happy to send anything else you need to verify it.

On the strict side, for a Best / Top N page, don't use this paragraph; use the removal request below.

Swap-in: removal request (strictly regulated side, listed on a Best / Top N page)

When to use: on the strict side, your business is listed on a list page whose title or entries carry laudatory terms (Best / Top N, "best N businesses"). **Replace points 1 and 2 of the main letter; leave the rest unchanged**; send it in the name of the client themselves (the licensee), have the compliance officer look it over before it goes out, and attach no material of any kind:

> We have just seen that [business name] is included in "[exact page title]" at [URL]. We did not ask to be listed and have not authorised it. The advertising rules that apply to us do not allow our services to be promoted in rankings or award lists of this kind, so **please remove [business name] and any details about us from this page.** We are not asking for any other change to the article. Could you let us know once it has been removed? We will check the page again on [date, 14 days from today].

Check on the due date: on day 14, open the original page and save a screenshot to the archive; if you are still listed, follow up at +7 and +21 as usual; if it still hasn't been removed after both, move it to the monthly recheck column, and in the monthly report write it as it stands: "Removal requested; the publisher has not acted."

Follow-up at +7

English

> Hi [Name], following up on my note of [date] about **[exact URL]**. Just need one line from you: is inclusion on that page editorial, a paid placement, or closed — and if paid, is it a fixed rate? Either answer is useful — if it's closed I'll stop bothering you. Thanks. — [Name]

Chinese

> “您好 [称呼]，[日期] 关于 **[URL]** 的邮件跟进一下。只需要您一句话：这一页的收录是编辑决定、付费刊登（固定刊例还是按询盘），还是不开放？哪种答案都行——如果不开放，我就不再打扰。谢谢。—— [姓名]”

Follow-up at +21 (the last one)

English

> Hi [Name], last note from me on **[exact URL]**. I'll mark it as "not open" on our side unless I hear otherwise. If that changes — a [2027] refresh, a new directory intake — we'd still be interested; my details are below. Thanks for your time. — [Name]

Chinese

> “您好 [称呼]，关于 **[URL]** 这是最后一封。若无回复，我方按「不开放」处理。日后如有变化（[2027] 版重写、目录重新开放收录），我们仍有兴趣。感谢。—— [姓名]”

No more letters after +21. Relabel that URL's reachability as "reachable (no response)", put it in the monthly recheck column, and send one more round in Q4, the season when listicles are rewritten for the following year. Replace the year in the square brackets with next year's.

Review reply template (reply to every review within 72 hours, negative ones included)

The same on every side. **Never confirm the other person's identity, and never restate the specific service they received.**

Chinese

> “谢谢您留下反馈。我们不会在公开渠道讨论任何个人的服务细节。如果您希望我们跟进，请通过 [电话 / 邮箱] 联系，我们会由 [职位] 直接处理。”

English

> Thank you for the feedback. We do not discuss any individual's service details in a public channel. If you would like us to follow up, please contact us at [phone / email] and [role] will handle it directly.

Three fixed rules: ① the reply never combines a service name, a branch name and personal information ② no price, results, success rate or peer comparison appears ③ every negative review gets the same template: no arguing, no restating what happened.

Why we never confirm identity: Confirming in a public reply that "this reviewer is our client" and restating the service they received could count as disclosing personal data publicly without consent, and on regulated sides it also runs into the industry's confidentiality obligation (rule text in each industry edition's appendix). **The client writing it themselves in their review does not mean you have permission to confirm it publicly.** For whether the strict side can request reviews, and how, see → General Edition 6.5 Reviews: how to ask, how to reply, how to fill the empty slot.

B.4 Headers for work orders, ledgers and monthly reports

What you'll do in this section: Copy the first line of the work order, the assignment rules, the 16 fields of the *Citation-slot table*, the time table for the six dispositions, the three monthly review numbers, the sample monthly ruler readout and the headers for page 1 of the monthly report.

First line of the work order

A one-page work order is ≤30 lines in total; its first line is fixed to these five items:

```
Side: <strictly regulated / lightly regulated / unregulated>
Switches: <each add-on switch that is on, by letter A–F; "none" if none>
Acceptance tier this quarter: <attributable / observable>
Top 3 buyer types by annual opportunity value: ① <type 1> ② <type 2> ③ <type 3>
Languages: by default, an English and a Chinese page for each buyer type
```

For how to decide the side and the switches, see → General Edition 0.2 Decide which side you are on first. If you can't get permission to read the access logs, the red blocker and the date the logs will be supplied also go on the first line. Priority is fixed on this page; every other cell downstream just copies it. One line = one assignable job.

Assignment rules (paste at the end of every task)

1. Keep the URL and the evidence for every finding; label each figure clearly as measured, estimated or unknown. "Done" means all three mechanical gates pass, not that a report was handed over.
2. Don't write small per-question numbers; never promise traffic, conversions, ROI or ranking.

3. Follow the bans in → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers–→ General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions for your side; delete on any hit.
4. Until the compliance memo is signed, don't write any page copy that will go on the site and don't send any outreach letter; once it's signed, check the draft against the memo line by line.
5. The density gate and the ban gate sit side by side; both have veto power. A fact the ban gate hits is always deleted, never kept because "deleting it leaves fewer than 5 items". If fewer than 5 remain after deleting, go back to the brief stage and swap the intent cluster.
6. Any material that names a third-party organisation is for internal eyes only; never delete the footer qualifier (template in B.2, "Internal material").

The 16 fields of the citation-slot table

The main off-site ledger. Fields are fixed — no adding or removing columns. One row = one citation slot (one URL, not one domain); the input must be `domains_cited` from the raw probe data, not the aggregated version. The aggregated version strips out generic domains, and profile-type sites are exactly the target of this step.

Field	How to fill it in
Source URL	The full URL, down to the specific article — never a bare domain
Times cited	The URL-level count from the raw probe data
Form	listicle / directory / editorial review / media / government register / review site / community
Are you present	present / absent
Is a competitor present, at what position	List the competitor names that appear + their position (any export with this column is for internal eyes only)
Reachability	already present / reachable / not reachable (marked per URL, not per domain)
Compliance verdict	can do / correction only / request removal / do not do (per → General Edition 6.2 The send gate, paid listings and the citation-slot table; request removal is used only on the strict side, for a Best / Top N page that lists you)
Contact email	Email or form URL; write "not found" if you can't find one — never leave it blank
Rate-card price S$	free / a figure / quote on request
Charging basis	fixed rate card / per enquiry / commission / unknown (if unknown, do not send)
Submission date	YYYY-MM-DD

Field	How to fill it in
Follow-up date	the actual dates of the two follow-ups, +7 / +21
Live URL	the real public URL, filled in only after you've opened and verified it
Present at monthly recheck	mark once on each monthly retest day: present / dropped / page deleted
Copyable structure	five binary features, marked once per URL: `price table` / `side-by-side comparison` / `author byline + credentials` / `visible update date` / `FAQ section`. Present = 1, absent = 0, written as the five-digit string `1-0-1-1-0`
Competitor's way in	one of five, all observable on the page: bylined submission / labelled sponsored or paid / platform-edited entry / user-submitted profile / cannot tell. Three criteria: whether there's an author byline and date, whether there's a sponsored / advertorial label, whether there's a public submission entry point

Six dispositions: time per item and expected turnaround

For the order and how to decide, see → General Edition 6.3 Six dispositions, bylined contributions and outreach letters; this table is only for scheduling hours.

Situation	Action	Time per item	Expected turnaround
Already present, but ranked low / description out of date	Update letter	20 minutes	1–2 weeks (inferred, 60% confidence, not measured)
Already present, but the information is wrong	Correction letter	20 minutes	1–2 weeks (inferred, 60% confidence, not measured)
Absent but reachable · Group B (government and statutory registers / professional association profiles / official manufacturer links)	Do this group first, and don't miss a single free item: claim the five business profiles → check government and statutory registers → complete association profiles → apply to manufacturer dealer locators	15–30 minutes each	Claiming: instant–3 days; registers and associations: 1–2 weeks
Absent but reachable · earned media	Bylined contribution	20–30 minutes per letter	2–6 weeks (inferred, not measured)

Situation	Action	Time per item	Expected turnaround
Absent but reachable · paid listings and directory submissions	Enquiry and submission, using the B.3 main letter	2.5–4 h to find contacts and rate cards + 2.5–4 h to swap the URL / query / disposition paragraph per letter, sent in two batches, 1 day apart	Free editorial slots: 2–4 weeks; paid slots: 4–8 weeks (both inferred, not measured)
Not reachable, or compliance verdict = correction only / request removal / do not do	Monitor only (correction-only cases still get one correction letter; request-removal cases get one removal request and a check on the due date), enter the monthly recheck column	0–20 minutes	—

Hit-rate basis: bulk outreach reply rate 15–25%; 2–6 weeks from a placement going live to a change showing up in AI citations (**inferred, 55% confidence, not measured**). Both numbers are for scheduling internal hours only. Never say them externally (see B.2, "On this quarter's outreach").

Three monthly review numbers (passive observation on the strict side)

Number	Definition
Review count and star rating	End-of-month snapshot vs. end of last month, by platform
Average reply time	Hours
Reply rate	Replied this month / new this month

On the strict side, always add this one fixed line to this section of the monthly report: "We do not request reviews. The figures above are a passive record of reviews that arose naturally and are not used as a work target or acceptance measure."

Sample monthly ruler readout

All numbers are examples — copy the structure, not the numbers:

```
This month's ruler readout (2026-XX-XX · 30-question pool frozen · 5 rounds per que

① OpenAI API leg        14 seats   (last month 12)
   Gemini model leg      9 seats   (last month 9)
   Google AI feature impressions  1,204   (Search Console, first-party · impression
   → These three lines measure three different systems; they cannot be added togeth
```

```
② On the list and cited this month   7 / 20   (last month 5 / 20 · denominator is t
   Of which on the list but not cited   5 / 20

③ Factual errors (count)   2        (baseline 9 → last month 4 → this month 2)

This month's noise band     ±3 seats  (3 same-week retests: 14 / 17 / 15, range 3)
Verdict                     Did not exceed the noise band → within normal variation
```

Page 1 of the monthly report

All numbers and URLs are examples — copy the structure, not the numbers:

```
This month's off-site asset ledger
① Pages on the list       12 / 20   (last month 9 / 20 · denominator is the frozen
   Of which cited this month   7 / 20   (last month 5 / 20)
② Reviews and star rating  This month 4.6★ → 4.7★ · reply rate 78% (last month 61%)
③ Citation-slot changes    2 new (examples: aaa.sg/directory, bbb.com/list)
                           1 dropped (example: ccc.sg deleted the entry after a red
                           3 new listings outside the baseline Top 20 (none cited b

This month's seat count    14 (last month 12) · this month's total names 1,082 · se
This month's noise band    ±3 seats (3 same-week retests: 14 / 17 / 15, range 3)
Factual errors (count)     2 (baseline 9 → last month 4 → this month 2)
Web control leg            Named-set overlap 60% · source overlap 45% (5 questions
                           → source overlap <50%: this month's seat change holds on
→ Verdict: did not exceed the noise band → within normal variation, cannot be deter
```

Three firm numbers (they don't depend on AI's randomness — they're evidence the month's work actually landed): pages on the list X/20 · reviews and star rating · citation-slot additions / drops. **Two measures that must always be printed alongside**: noise band ±N · the web control leg's two overlap rates.

B.5 Blueprints for unverified page types (for industry editions to include)

What you'll do in this section: When an industry edition needs the pt44, pt45 or pt46 blueprint, `@include` it from here; all three diagrams are labelled, in full, "not measured, spec taken from an external teardown", and not one block is marked cite. For where each of the three types sits and which to validate first, see → General Edition 5.30 The four tier-C types: unverified page types; this section carries only the blueprints themselves, without repeating the judgement calls and examples that section already covers.

title	{NEW} vs {OLD}: the concept gap + the outcome noun
intro	one-line axiom: the two complement each other, they're not substitutes
H2 Why AI doesn't recommend you	200–280 words, 4–5 buyer questions inline
H2 The difference in 30 seconds	180–250 words + comparison table
H2 Why it matters now	250–350 words
H2 What is OLD / What is NEW	250–320 / 320–420 words respectively
H2 The core difference	250–320 words + difference table
H2 Why you need NEW when you already have OLD	450–600 words + figure + 3 "Pattern" subheadings
H2 The most common problems in real projects	evidence section, 550–700 words + figure + numbered subheadings
H2 Four common problems	350–450 words: doesn't exist / known but not recommended / wrong information / buried under old information
H2 Priority differs by engine	300–400 words + figure + comparison table
H2 Will it replace it → How they work together →	180–250 / 300–400 / 300–400 words + 3 "Case" subheadings

Which to do first	
H2 Summary	250–320 words + figure: OLD governs search results, NEW governs AI answers
FAQ	8–11 items, answers must have no bold, no links, no bullets, no brand name

☐ Required

Figure: block order for the concept pillar page (the whole diagram is not measured, spec taken from an external teardown, no block marked cite)

title	An assertion-style title, with an event clause and a claim clause
intro	No heading, 3–5 sentences, 80–120 words: date + event → the bet we're making → headline number → a lesson sentence that leads with a negative
H2 Key findings	4–5 items, one sentence each, 20–40 words, third-party findings first, the last item states our own limits
H2 What happened	120–200 words, past tense, zero adjectives
H2 Why the mechanism matters	150–220 words, one study per paragraph, each with its source
H2 Our own evidence	180–260 words, evidence table
H2 What the reader should do	5–6 bullets, noun-phrase labels, no product names, no using the topic for soliciting `Regulated`
H2 Our own positioning	80–130 words
H2 Method and limitations	60–110 words, one paragraph
FAQ	Exactly 3 items, questions phrased as objections, every answer opens with a negative sentence

Required Optional Regulated

Figure: block order for the news and policy explainer page (the whole diagram is not measured, spec taken from an external teardown, no block marked cite)

title	A market-change statement, with a structural assertion	
intro	90–160 words: statement of fact → evidence → reader-scope constraint → corrective sentence	
H2 Key findings	Exactly 4 items, 20–45 words each, fixed order	
H2 Argument one / Argument two	200–300 words each, mechanism paragraph	
H2 Our own first-hand view	at 40–55% through the article, 200–300 words	Regulated
H2 Framework	150–220-word framework table, column headers below	
H2 Who the market change affects	180–260 words	
H2 What buyers should test	120–200 words, no product names	
H2 Our own positioning	the only sales paragraph, 90–150 words, the third- or fourth-to-last H2, no links, play the brand down	Regulated
H2 Method and limitations	60–110 words, no bullets	
FAQ	Exactly 3 items, 35–70 words, conclusion first	
References	2–4 items	

☐ Required ▭ Regulated

Figure: block order for the market observation page (the whole diagram is not measured, spec taken from an external teardown, no block marked cite)

Full-page specs the three diagrams cannot show

The full-page specs from the external teardown are for reference only (not measured). On length, write what the page type needs, then stop; see → General Edition 5.7 Seven rules that hold for every page type.

Page type	Length and structure	Tables, figures and CTA	Writing constraints
pt44	4,500–5,000 words, 15–16 H2s; the first four H2s use sentence case, switching to Title Case from `What Is {OLD}?` onward	3 tables (30-second difference, core difference, priority by engine; the external teardown gives no fixed column headers), 4 figures, one CTA placed at about 91% through	Contractions allowed; zero em dashes; heavy use of semicolons; hard numbers in bold. The first sentence under every `What Is X?` is written as a definition sentence: the subject is the concept, the verb is "is" / "consists of", and it carries one checkable range (not measured; inferred from pt12's measured pattern)
pt45	1,250–1,400 words, 9–10 H2s	Exactly 1 table, content is our own evidence; no CTA	Strict third person
pt46	1,500–1,750 words, 11–13 H2s	2 tables, framework-table column headers `Layer / What to record / Decision it supports`; no ranking, no CTA	Third person only; body text has zero em dashes, zero contractions, zero we/you

Appendix C

Five industries side by side

Dental and aesthetics, family law, tuition, B2B software and e-commerce across ten factors

C.1 Ten-axis values for the five industries covered, and three things for week one

C.1 Ten-axis values for the five industries covered, and three things for week one

What you'll do in this section: When you're judging a new industry, first find the industry in the table below whose ten-axis values are closest to yours and use it as a reference, then go back to the main text and follow the rules. These are only examples: do not copy another industry's tactics straight onto your own.

① **The ten questions**
8.1 the ten questions — which axis each one fills in

② **Decide the side**
8.2 the five steps that set the side and switches, then back to 0.2 to double-check

③ **Fill in the ten axes**
8.3–8.4 fill in your own value for each axis

④ **Find a reference**
pick the column below whose ten-axis values look most like yours

⑤ **Back to the decision table**
8.5 rule on each axis; do not copy that column's tactics

Figure: This table is only a reference for step 4; tactics are still decided by the main text's decision table

Every cell in the table is a summary, not a work order: the axes are defined in → General Edition 8.3 The ten axes (1): page types, containers, anchors, regulation, question shapes and → General Edition 8.4 The ten axes (2): decision-makers, capacity, attribution, half-life, language; the seven tactics are defined in → General Edition 8.6 Seven tactics: how regulated and unregulated sides differ, and why each moves rankings.

Quick examples only; the rules are set by the main text.

Values on the ten axes

Axis	Dental/aesthetics	Family law	Education and training	B2B software/SaaS	E-commerce
1 Shelf page-type mix	News and long-form articles dominate (`article` 54% / `listicle` 0%)	Government, court and legislation pages dominate the ChatGPT leg (family law 47/60, 49/59); on the AI Mode leg, law firms' own explainer articles and fee pages lead (law-firm sites 28/52); rankings, review platforms and the Law Society directory together make up 10/171 and lead only on "who's best" questions	List pages dominate, with long-form price guides second	List pages + comparison pages together ≥60%; documentation sites about 8%; news about 1%	A mix of list pages and comparison pages. For platforms, read this as two separate sentences, not one: ① a platform's product detail pages feed only the shopping shelf; ② a platform's shop bio, category pages and product description text can be copied by the retrieval leg into the body of an answer (provided that platform's door is open on that leg — decide leg by leg). Sentence ② is the shortest link in path ①

Axis	Dental/aesthetics	Family law	Education and training	B2B software/SaaS	E-commerce
2 Control of citation containers	Mostly third-party slots; you can also take a few own-site slots	Read the two legs separately: the government and court pages the ChatGPT leg cites are not reachable; on the AI Mode leg, own-site slots can be taken (law-firm sites 28/52); directories and rankings are third-party slots and lead only on "who's best" questions	Mostly third-party and platform slots; you can also take own-site slots; 6–8 reachable-enough URLs	The most own-site slots (highest of the five industries); a whole row of platform slots; very few not reachable	Both own-site slots and platform slots are high; whether a platform slot can be used varies hugely from platform to platform — you must verify each with a live curl test
3 Entity anchor layer	A person (the doctor); the organisation second	A person (the lawyer) first, the firm's approved name second (a judgement call, not measured)	Splits by sub-segment: organisation / course / organisation + instructor / person	Product SKU; the founder and engineers second	Dual anchors: parent brand (answer shelf) + SKU (shopping shelf)

Axis	Dental/aesthetics	Family law	Education and training	B2B software/SaaS	E-commerce
4 Side and switches	Strictly regulated side + switch A (agency liability) + switch E (peer comparison banned), already part of the strict side's base layer; add switch C (regulated products) when aesthetic devices are used, or when skincare or supplements are also sold	Lightly regulated side (75% confidence) + switch A (agency liability; the "principal bears it alone" form) + switch B (legally required fields; partly on: only fixed wording that specific conditions trigger) + switch E (peer comparison banned) + switch F (referral commissions banned; any form of consideration counts); C off, D depends on the client; family cases add the "three family-law rules", which tighten how the case-study, testimonial and review-reposting, and review-reply cells are written, without changing the side	Regulated: build the whole site to the strictest segment's standard (switch D (spans two sides)) + switch A (the rule text names the party publishing on the client's behalf) + switch B (ads must carry the registered name, registration number and validity period); whether it falls on the strictly regulated or the lightly regulated side is still to be verified (see below the table)	Unregulated side, with three pass-through red lines: certification wording is a legal claim, the regulation of your buyers' industry passes through to you, and platform terms are enforced faster than the law	Unregulated side; five sub-categories, including supplements and devices, turn on switch C; if any one cluster is hit, build the whole site to that cluster's regulated basis (switch D); switch A is still to be ruled on (see below the table)

Axis	Dental/aesthetics	Family law	Education and training	B2B software/SaaS	E-commerce
5 Question-shape mix	High share of spec type (price)	Spec type + recommendation type (the mix we set when writing the questions; not measured, must be recalculated against the frozen question pool)	Recommendation 4 / spec 3 / comparison 2 / approach 1 (approximate, not measured, must be recalculated against the frozen question pool)	Recommendation 35 / comparison 30 / spec 25 / approach 10 — the only industry where all four fronts must be fought at once	Filled in only after the frozen question pool is signed off; left blank until then — pre-filling is not allowed. Prior assumption (for sequencing the work only; it goes into no deliverable and is never said externally): the share of shopping-type questions is significantly higher than for training, and this must be measured against the frozen questions; do not copy the training column over and treat it as your conclusion

Axis	Dental/aesthetics	Family law	Education and training	B2B software/SaaS	E-commerce
6 Number of decision-makers	Single, over weeks	Depends on the matter (a judgement call, not measured): a divorce is one person (the other party is the opponent); wills and LPAs are often a couple	K12: two people + a third-party user (the child); corporate training: committee	Committee (user / IT and security / finance / legal)	Single, on the spot; for high-ticket durable goods, single over weeks or two people
7 Hard capacity cap / work unit	Practitioner-chair count / one patient type	Lawyer hours / one client type	Class capacity × how often classes run / one course line	No hard cap → switch the numerator / one use case × one role (this is the pricing unit; pages are built one per use case, at most 4 use-case pages a quarter)	Inventory and supply volume / one category cluster
8 Attribution channels	Phone + in-person visits + forms	Forms + phone + WhatsApp (a judgement call, not measured)	WhatsApp (the primary channel) + trial-lesson forms + campus visits	Free-trial sign-up / demo / docs site → sign-up	Add-to-cart and checkout (on platforms there is almost certainly no referrer)

Axis	Dental/aesthetics	Family law	Education and training	B2B software/SaaS	E-commerce
9 Content half-life	Medium/long (driven by procedures and regulation)	Long (driven by regulation)	A mix of three levels, decided field by field	Short (quarterly); compliance evidence: medium; docs site: event-driven	A mix of three levels, decided layer by layer: ① the live-numbers layer, ≤15 minutes — price, stock and promotions, managed by the feed and never entering the body text; ② the spec-and-ranking layer, monthly — spec numbers, list-page ranking, the values in comparison tables; ③ the structure-and-basis layer, quarterly — price ranges, definitions, the variables that decide price, and how buyers choose. The review cadence follows the layer, not the page

Axis	Dental/aesthetics	Family law	Education and training	B2B software/SaaS	E-commerce
10 Language fork	The difference set often holds (medical tourism)	Depends	The difference set holds, with the cleanest evidence	No language fork, but a jurisdiction fork holds	Decided by the buyer mix; for the Singapore market, run the Chinese difference set first by default

Tactics that follow from the ten axes

Tactic	Dental/aesthetics	Family law	Education and training	B2B software/SaaS	E-commerce
Main battlefield	Off-site + own-built price pages	Own-built fee pages and explainer articles (the AI Mode leg mainly cites firms' own pages); off-site directories and rankings only for "who's best" questions	Platform slots + own-built tuition-fee pages + list-page editorial slots	Own-site pages (the highest of the whole table)	Fighting on both shelves at once: own-built media-style pages + the feed
What to build first	Tiered fixed-price page	Tiered fees and billing-method page	All-category tuition-fee guide page (English and Chinese published together)	`/pricing`, eight blocks	An all-category comparison page that includes competitors

Tactic	Dental/aesthetics	Family law	Education and training	B2B software/SaaS	E-commerce
Can you get onto "Best" lists	Do not do it; if you are listed, ask the publisher to remove you (labelling it as paid is not a release from liability)	Self-nomination not allowed; can take part via editorial slots	Possible in some segments (editorial slots); in the strictest segment, entry descriptions must not be laudatory	The main tactic	Can get into editorial review slots; cannot buy a slot
Asking for reviews	Never ask; passive observation only	May ask, on four conditions: no incentives, no ghost-writing, no cherry-picking (no gating), no dictated content; the basis is that no current rule was found banning it, not that a rule allows it (70% confidence); family clients' reviews are not reposted by default	Can do it, split by campus	Mandatory, but the platform's terms are the gatekeeper	Mandatory, with three hard boundaries: ① no gating; ② any incentive must be disclosed, and applied the same way to every invited buyer; ③ never have staff, friends and family, or fake-order accounts write reviews, and every review you display must be visible text
Competitor comparison	Compare procedures only, never clinics	Do not do it	Compare only checkable, objective items	The main tactic	The biggest tactic here, compared with the industries before it

Tactic	Dental/aesthetics	Family law	Education and training	B2B software/SaaS	E-commerce
Weight of earned news	High (`article` 54%)	High	Medium	About 1% (off-site work ranked P3)	Medium (review sites and media listicles are the main source, but not news)
Number of retest rounds	3	3	3	5	3 (but each of the two seat rulers measures its own noise band)
Special door pitfalls	WAF	None	Timetables/price lists/faculty lists rendered in JS	Multiple subdomains + docs-site CSR + hosted trust pages	CSR + numbers with no sentences + the platform's own door
The only change to the ruler	None	None	Split the denominator by grade level × subject	Split factual errors into two columns; measure total names by category	Count seats with two separate rulers; add price errors and stock errors to factual errors

Tactic	Dental/aesthetics	Family law	Education and training	B2B software/SaaS	E-commerce
Net advantage / net disadvantage	Advantage: the official registration number gives you the strongest anchor for free	Advantage: the official register lookup (the LSRA lawyer and law-practice search) gives you an anchor for free; whether it has a publicly citable number field is still to be checked	Disadvantage: instructor turnover breaks path ③ (getting identified correctly)	Disadvantage: namesake-entity pollution; Advantage: the most own-site slots	Disadvantage: one extra door layer that is not in your hands; Advantage: product data supplied directly (the feed) never goes through the door

The four kinds of container in the "control of citation containers" row — own-site slots (you can claim these on your own website), third-party slots (need someone else's agreement), platform slots (fill in the back end and they take effect), and not reachable — are defined in 8.3, axis 2; what the six switches in the "side and switches" row each mean is in 0.2.

The "side and switches" row is decided using the four questions and six switches from → General Edition 0.2 Decide which side you are on first. Dental/aesthetics and family law have published industry editions, with the rule text and sources in each book's appendix; the other three columns are reference results from the four questions — the rule text has not been checked clause by clause, so this is not a final ruling. The family-law cell is now decided; two cells are still unresolved:

- **Family law**: decided as the lightly regulated side, 75% confidence. Switches A (agency liability), E (peer comparison banned) and F (referral commissions banned) are on, B (legally required fields) is partly on, C is off, and D depends on the client; family cases add the "three family-law rules", which only tighten how the case-study, testimonial and review-reposting, and review-reply cells are written, without changing the side. For the basis, the strongest counter-argument and the conditions for changing the ruling, see Chapter 0 of the Law Firms Edition.
- **Education and training**: for the strictest segment (registered private education institutions), testimonials and objective comparisons are both allowed with limits, which points to the light side under question 3; but for tuition fees it requires eight items, each written as a fixed figure, and the original text does not say whether a range can be written (question 2). Which side it falls on is still to be verified in the training edition; until then, write to the stricter, strictly regulated side.

- **E-commerce**: the original text contradicts itself on the agency-publishing cell. One place says there is no agency liability and that set of authorisation paperwork is not needed; another says no rule has been found that names the publishing party, yet every batch must still be signed off. This is still to be ruled on by the industry edition; until then, follow the stricter practice and sign off every batch of external material.

Even if your industry is one of the other three, still check your own industry's original rules against → General Edition 8.2 Five steps to decide the side, and three criteria before you decide the side; when unsure, start on the stricter side.

On penalties for agency liability, the original text gives figures for two industries: dental/aesthetics — a fine of up to S$20,000, up to 12 months' imprisonment, or both, plus up to S$1,000 a day for a continuing offence (checked against the original rule text); for the strictest segment of training, a fine of up to S$5,000 or imprisonment of up to 6 months for false or misleading advertising, and a fine of up to S$10,000 or imprisonment of up to 12 months for not complying with an order to withdraw and correct. Rule numbers and quotations from the original text are not included in this appendix: for dental/aesthetics, see *Dental & Aesthetics Edition*, Appendix A · Rule text and sources; for family law, see *Law Firms Edition*, Appendix A; for the other industries, these will be added once their industry editions are published.

Cells in the table that say "Do not do it" or "Never ask" do not all rest on the same label tier; the clause-by-clause labels in the relevant industry edition decide; where something is banned as Statute text, a client's written sign-off cannot make it publishable — a written sign-off resolves only the three stop-and-escalate situations; it never turns something banned as Statute text into something publishable (see → General Edition 0.4 The three labels for compliance sentences, and the stop-and-escalate rule).

Tactics ruled out and limited

"Ruled out" means that tactic cannot be used at all; "limited" means it can be used but with hard boundaries, not a blanket ban. The table only fills in figures the original text gives; where none is given, it says so — nothing is estimated.

Industry	Ruled out	Limited
Dental/aesthetics	Four tactics	Not listed separately in the original
Family law	Two and a half tactics	Not listed separately in the original
Education and training	Spans tiers; the original gives no figure	Spans tiers; the original gives no figure
B2B software/SaaS	Zero tactics	Zero tactics (plus three pass-through red lines — see the "side and switches" row above)
E-commerce	Zero tactics	Three tactics

Three things for week one (examples)

Examples only: reorder the three things your industry does in week one to match the side-deciding and baseline order in → General Edition 0.1 The book in one sentence, and the 90-day reading order; do not copy the table below directly. For the first item in the table, "the five door gates", week one is a read-only audit only; fixing the door has to wait until after the baseline is saved.

Industry	First	Second	Third
Dental/aesthetics	The five door gates + `/facts`	Tiered itemised fixed-price page	The five business profiles + checking official registers + manufacturer locators
Family law	The five door gates + one page per lawyer	Government and association registers, without missing a single one	Tiered fees and billing-method page
Education and training	The five door gates (including a CSR check on the timetable)	Three platform-slot (C) items: course-catalogue fields + the five profiles for each campus + platform instructor pages	Tuition-fee guide page, English and Chinese launched on the same day
B2B software/SaaS	The five door gates + four special gates	The eight blocks of the pricing page + removing the `nosnippet` tag	A visible HTML mirror of the Trust page on the main domain + checking for namesake conflicts
E-commerce	The five door gates + three tables for CSR / visible text / platform robots (decide leg by leg, not the platform's overall state)	Aligning the dual entity anchors: ruling on the brand name vs the model name + rewriting the shop bio on each platform	Submitting the feed application + a 100-item sample batch (does not go through the door). Week one is the submission and the sample batch, not activation: activation has to clear the other side's business verification, about 1–3 weeks; the feed must not be pushed until the dual anchors are aligned, and the full push is scheduled after the second item

Appendix D

How to write numbers, and the glossary

Rules for labelling numbers, and the book's terms

D.1 Number discipline: how to label numbers, and what stays internal

D.2 Glossary

D.1 Number discipline: how to label numbers, and what stays internal

What you'll do in this section: label every number you plan to publish with its source and confidence, sort material into two kinds (can go public / internal only), and never write a single month's swing up as a long-term conclusion. When you're done, you'll have a self-check list that every number goes through before it goes out.

Which label a number gets

Label	How to write it
External data with a source	Write the source and sample size directly
Inferred: derived from external data, not directly verified	"(inferred, XX% confidence)"; the percentage must be filled in
Self-set number: projection, post-mortem, a self-set threshold	"(self-set number, XX% confidence)" + mark item by item whether it may be said publicly
Confidence below 80%, in any form	Label it in the same sentence; never in a footnote or a separate note

Figure: every number gets exactly one of four labels, each with its own fixed wording

External data with a source is written like this: "84% of AI citations come from **earned media** (news coverage alone is 27%), while **paid placements and advertorials are only 0.3%** (Muck Rack, 25M links / 17 industries)".

Five hard rules

Hard rule 1 · Never hide a known, material uncertainty. Any number used to judge whether something worked must have its uncertainty printed on the same page. Hiding "the seat-count noise floor is unknown" while announcing results with seat counts = lying.

Hard rule 2 · A single-run, single-engine, n ≤ 3 observation must never be used as a result or as acceptance evidence. It can only serve as directional evidence of the kind "this factual error is there however many times you ask", and the sampling spec must be written on the same figure.

Hard rule 3 · The data point below is written only here, once in the whole book, and not a word of it may be changed; any other chapter that needs it points here and does not write it out again:

> Getting onto a list that AI already cites and ranking first on it = visibility **+16.5pp** and a position **0.8–1.8 places** earlier in the answer (Peec, 5.7M data points / nearly 200,000 answers / 8 engines; **+16.5pp is from a B2B SaaS sample; for emerging categories it is +13.4pp; not separately validated in Singapore dental, aesthetics or legal. Use it to order the work, never to promise any result**).

Hard rule 4 · Scouting numbers ≠ your current numbers. A number from a **generic retrieval API, with the region not locked, on a single engine, in a one-off run** can only be used to order your own work and decide where the effort goes; it must never be said publicly as your current state. **Any number calculated from it (for example, "expected N replies per batch") is equally barred from external use.** For the current list of which numbers count as scouting numbers, see → General Edition A.1 Evidence for mechanics, the door and identity. Lifting the ban needs all three conditions at once; if even one is missing, the ban still applies: ① retested on the OpenAI leg (direct connection) + the Gemini model leg, with the target region locked, on your own frozen question pool; ② the retest date and engine are stated (including the number of passes and the region parameter); ③ what you cite externally is **your own number after the retest**.

Hard rule 5 · Any long-term decision based on a single month's swing needs two consecutive months moving in the same direction before it is written into the discipline. A single month's rise or fall may only trigger "does this task go on this month's schedule or not":

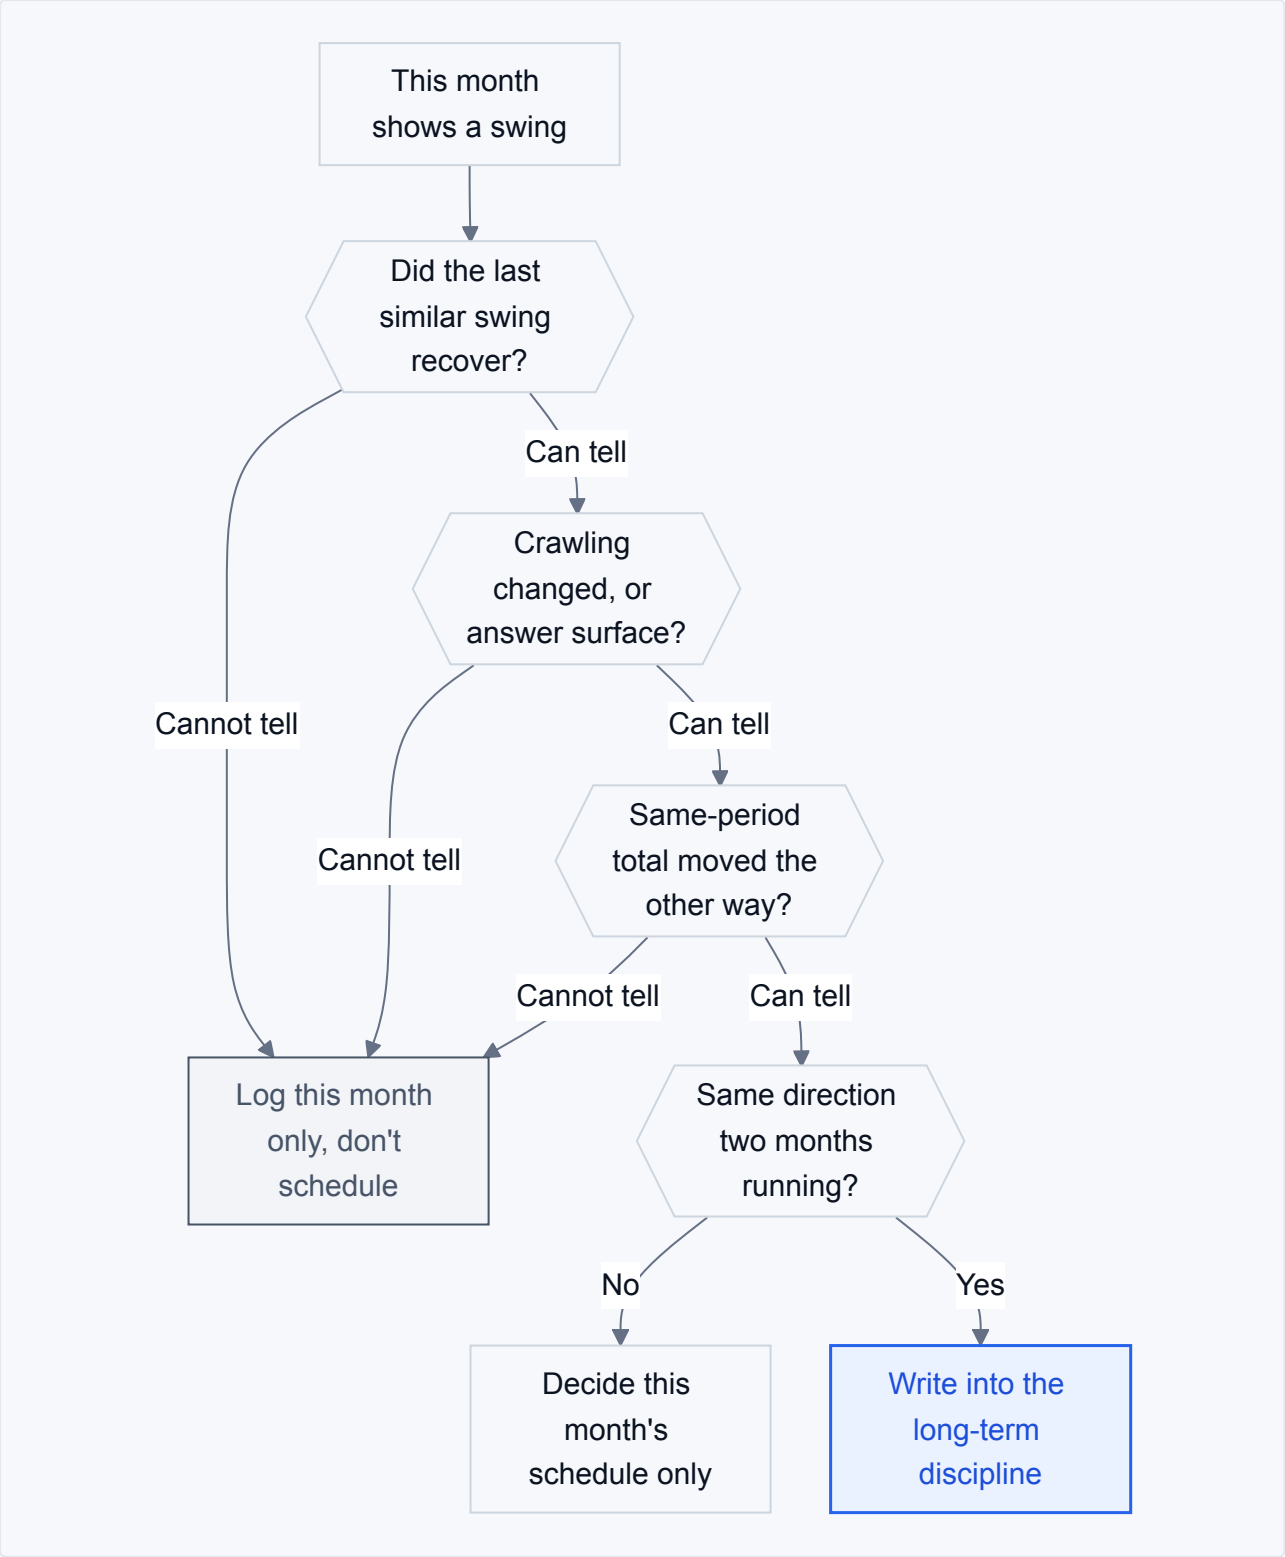


Figure: a single month's swing must pass three checks, then move in the same direction for two consecutive months, before it is written into the long-term discipline; if any check cannot be answered, log it for this month only and do not schedule it

What each of the three checks rules out: ① it recovered last time = this is a swing, not a structural change; ② the source is still being read, it's just not making it into the answer = the answer surface narrowed, the source hasn't failed, and the work still has value; ③ the denominator changed, which

makes a structural change read as an absolute decline. Log all three, item by item, on the same page. **You may not** write long-term conclusions such as "we're dropping this source / this leg is dead".

> Red line **The sentence "an AI answer usually has five to eight names" must never appear anywhere in this book.** It has no source, and it contradicts the rule "record exactly as many businesses as are actually named; never estimate". Use the number measured in your own baseline instead: "I asked K times and counted T company names in total that AI named. **That is X.X businesses per answer on average, fewest A, most B. This is what I counted on this batch of questions, not an industry figure.**"

What material stays internal

Material	Can it be shown externally
Seat shelf map, weekly competitor seat report / competitor board in the monthly report	No, internal only
The "why AI trusts them and not you" comparison table	No, internal only
The parts of sentence-by-sentence before-and-after screenshots that show a third-party name	No; may be shown externally once names are redacted
Your own pages-on-the-list count, profile screenshots, before-and-after site comparisons	Yes (no third-party names)
Facts pages, price pages, thick pages, business profile text, outreach material	Yes; these are your advertising, so check them against your side's bans first

Figure: one test only: if the material names a third-party organisation, it stays internal

Internal only = must not be posted, must not be forwarded, must not be used for any promotion. Print the "Internal reference material" line, word for word, in the footer of this material (burned into the PDF or image itself, not written in the body of an email) and add an `-INTERNAL` suffix to the file name. For the footer wording and the three approved lines you may say publicly, see → General Edition B.2 Approved wording and scripts. For the bans by side that apply to the two "Yes" rows of the table above, see → General Edition 5.2 Quick reference by side (1): identify the advertiser first; prices, promotions and freebies, result numbers and → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.

Why this matters most: regulated industries commonly have a clause along the lines of "no comparing with or disparaging peers", and some also govern promotion that involves a third party. The moment material containing peer comparisons is put up at the front desk or sent out through a public channel, it runs into this kind of clause, and it often takes very little for a named peer to complain. This is the

general pattern; to find exactly which clause applies on your side, check → General Edition 5.3 Quick reference by side (2): testimonials and reviews, comparisons, lists, titles, outbound links, FAQ and captions.

> Note **The most common slip-up**: the first time you show someone, in person, what AI currently says about you, that screen often holds a recommendation-type answer naming a peer. **Only point at the screen and talk it through: no screenshot, no re-saving, nothing put into a PDF, no copy kept.**

D.2 Glossary

What you'll do in this section: when you're not sure what a term means, look it up here instead of guessing from instinct. This is the book's only set of definitions; no chapter may set up its own.

The full definition of each of the following terms appears only here; every other chapter refers back to it without repeating the explanation:

Term	Definition
Names	The total number of business names in the body of one AI answer. **Measured and recorded, never estimated**
Seat / seat count	The ones among these names that belong to one particular business. Your seat count = the number of times you were named across this period's N question-and-answer runs. **"Seat" and "answer-level times named" are one and the same number, not two measures.** Note **The seat count scales up linearly with the number of passes**, so it can only be compared at the same number of passes, on the same engine and with the same question pool. The number of passes is frozen for the whole quarter; to add more passes you must start a new baseline and write, on the same page, "not comparable with the previous baseline" — **the old and new lines must never be added together or plotted on the same trend chart**
Listed as a source	AI lists a given web page as a source underneath its answer. Log each occurrence; it is **never used as acceptance evidence** (during the baseline period this is mostly 0 or 1 — not enough resolution)
Pages on the list X/20	Of the source pages in the frozen baseline Top 20, the number where your name appears **on the page itself**. **The denominator 20 stays fixed for the whole quarter.** It measures off-site work, not AI behaviour
On the list and cited X/20	Of the pages on the list, those that **were actually cited this month**. **On the list ≠ cited** — AI may cite that page without picking up your line
New listings	Third-party URLs newly listed this quarter that are not in the baseline Top 20. **Not counted in pages on the list**; logged on their own line, each marked for whether it has been cited in an AI answer this quarter. One that has not been cited is logged only as work done, never as a result figure

Term	Definition
Noise band ±N seats	The range (max − min) of the seat count when the same batch of questions is run 3 more times within the same week. **Floor fixed at ±1 seat.** A monthly change in the seat count **is only allowed to be written as "increased" when it is larger than the noise band**
Model-API visibility baseline	The seat and source distribution measured by the two model-API legs. In reports, the two legs may only be called "Model-API visibility baseline · OpenAI leg" and "Gemini model leg" — **never written as "what users see in ChatGPT" or "real user visibility"**
Gemini model leg	The leg run through the Gemini model API. **Gemini API ≠ AI Overviews ≠ AI Mode — three different products.** For AI Overviews / AI Mode visibility, **always use Search Console generative AI report impressions instead**
Web control leg	The leg where, once a month, a person manually runs a fixed set of 5 questions, 1 pass per question, on **the ChatGPT web app** (logged out or in a temporary chat, with an IP in the target market). **Calibration only, never acceptance evidence**: it does not go into the seat count, does not go into pages on the list, and is never added together with any ruler
Named-set overlap / source-domain overlap	On the same batch of 5 questions, the intersection ÷ union of the web leg's and the API leg's **named-business sets** and **source-domain sets**. **In any month where either falls below 50%, the report must print the non-extrapolation sentence unchanged**
Coarse-screen top 10	In the entry-ticket coarse screen, copy the source URLs listed under each question's answer into a table of its own, deduplicate, and take the top 10 by number of appearances: **three questions, three tables, never merged**. **n = 1, single engine, done by hand, finished within half a day**, serving only the one judgement of "is it worth doing". Note **It is not a denominator and is never used to report any number externally**; merge the three into one table and the next step, counting "reachable enough" table by table, can no longer be carried out
Baseline Top 20 (= account-level Top 20)	After the baseline has run its full rounds, combine the answers for **every intent**, and take the top 20 URLs by URL-level times cited from the raw `domains_cited`, **frozen, unchanged for the whole quarter**. This is the **only** valid denominator for "pages on the list X/20". Note **It does not judge any single intent** — a shared list always gives the same count, so using it to judge a single intent will either "rule them all out together" or "rule none of them out"
Intent-level top 10	Take the top 10 URLs by times cited, **separately for each intent, from that intent's own answers**. The method: add an intent-number column to the source summary table and group by intent to take the top 10 (the raw data already carries intent information; only how it is grouped when summarised changes). **The abstain line, check 2 of picking targets, and layer 3 of not-moved triage always read this table**, never the account-level one
Intent cluster	A group of buyer questions that are phrased differently but get essentially the same answer. **One cluster, one page**; variants go into H2s on the same page. Building a separate page for every variant only dilutes the signal

Term	Definition
Annual opportunity value	**Price per order × gross margin × monthly capacity cap × 12.** Use it to rank buyer types, which sets the priority for the question pool and the pages

The terms below have their full definition in their own chapter; here you get one line plus a pointer:

Term	One line + pointer
Strictly regulated side / lightly regulated side / unregulated side	The three sides, decided by whether you need a licence or registration to open, whether price can be written as a range, and how far testimonials and peer comparisons are banned; for how to decide, see → General Edition 0.2 Decide which side you are on first
The six add-on switches A (agency liability)–F	Six switches that don't change the side, only add one more hard gate on top of it: A agency liability, B legally required fields, C regulated products, D spans two sides, E peer comparison banned, F referral commissions banned; for how to decide, see → General Edition 0.2 Decide which side you are on first
Content form (professional / consumer)	The second, independent decision after you've decided the side; it decides whether a page's weight goes on process or on tiered prices; for how to decide, see → General Edition 5.9 Content form, FAQ, Chinese pages and other-language pages
Own sentence / market sentence / official-basis sentence	The own sentence is for the ChatGPT leg: the subject is you, with an exact number; the market sentence is for the AI Mode leg: a range + variables + a source. The strictly regulated side and switch E (peer comparison banned) each replace the market sentence with an official-basis sentence. For the full spec, see → General Edition 5.6 Two legs, two kinds of sentence
The three labels	"Statute text", "Conservative line (not statute text)" and "Original text not obtained" — the tier of evidence behind a compliance sentence; for how to label and the stop-and-escalate rules, see → General Edition 0.4 The three labels for compliance sentences, and the stop-and-escalate rule
The six page-type families	The 46 cited page types, grouped into six families by how buyers phrase their questions: the Price family, the Selection and reputation family, the Questions and rules family, the Entity and product facts family, the Primary sources family, and the Commentary family; for the full map, see → General Edition 5.1 The page-type map: 46 types, six families, three evidence tiers, availability by side
Split day	The day you fix the door is the whole project's split day; every page also has its own split day, which is its launch date; for how to decide, see → General Edition 2.6 Fixing the door: nosnippet, the four-step robots merge, the WAF allowlist, → General Edition 2.8 The launch gate for every page

Term	One line + pointer
Baseline saved	All three are saved: the baseline Top 20, the noise band measured by the same-week retest, and the 36 brand-six runs; for the full spec, see → General Edition 4.4 The web control leg and the frozen baseline (the book's only full spec)
Weeks used across the book	D1–D6 are for the audit and freezing the baseline; W1 starts from the day you fix the door (the split day); the D90 settlement falls around week 13; see → General Edition 0.1 The book in one sentence, and the 90-day reading order
The four identities	How the same URL looks through four paths: a 200 in your access logs, Bing's Live Test, GSC's "crawled" status, and ChatGPT quoting it back; for how to read it, see → General Edition 2.3 Gate 2: four identities, fetched live (read-only). The launch gate for every page has a separate check called "the four HTML copies match", which checks whether the source code fetched by four different user agents is the same version — that is not the same thing as the four identities
The four states	Already present / reachable / to ask / not reachable, judged per URL (against the coarse-screen top 10 during the coarse screen, against the intent-level top 10 after the baseline). **"To ask" is only allowed to exist at the coarse-screen stage** (counted as reachable there, with its count logged separately); at the baseline stage it must be cleared to zero; for the full criteria, see → General Edition 4.6 Four states, two denominators and the abstain line
The brand six questions	A fixed 6 questions × 3 rounds × 2 engines = 36 runs / month, frozen for the whole quarter. **Changing the set of questions = changing the ruler**: the factual-error counts of the month before and the month after can no longer be compared; for the questions themselves and what to do after you find an error, see → General Edition 3.5 The brand six questions, and what to do when you find an error

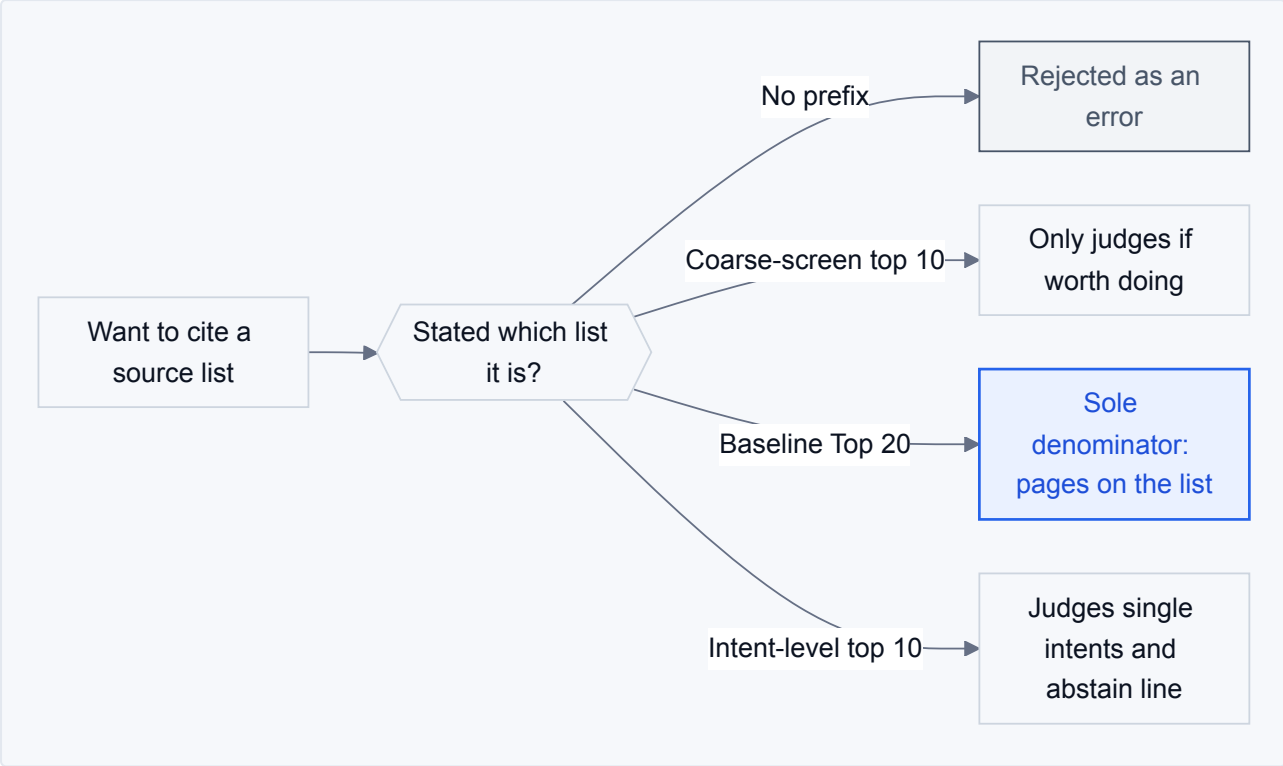


Figure: each of the three source lists answers one question only; anything that doesn't state which one is rejected

Mix them up once, and the basis for the whole quarter's verdicts is void: the coarse-screen top 10 is an n = 1 snapshot taken in half a day; the baseline Top 20 is one shared list that cannot judge a single intent. So anywhere "Top 20" appears it must carry a prefix (baseline / account-level — the same list), and anywhere "top 10" appears it must state whether it's the coarse screen or intent-level. **Writing it bare is always rejected as an error.**

Keep going

Web edition

canlah.ai/playbook/

Version for AI assistants (open-source skill)

github.com/Canlah-AI/geo-playbook

Measurement data (CC BY 4.0)

canlah.ai/data/ai-citations-sg-2026/

Want help with round one?

admin@canlah.ai · WhatsApp +65 8110 0234

Scan for the web edition
https://canlah.ai/playbook/

GEO Playbook · General Edition · v1.0 · CC BY 4.0 · DOI 10.5281/zenodo.23005022