

01 / THE BOOK

Seen isn't *selected*

Marketing's one job in the AI era



Seen isn't selected

Marketing's one job in the AI era · Ben Stokes

A NOTE BEFORE YOU START

I build software based on this thinking, so the argument has to stand without it. The core tools are printed in full, and every template is free in the Selectability Playbooks at rocmode.io/playbooks. RocMode itself appears only at the end of the book, where I tell how it came about. My own stories appear throughout, wherever they illustrate a point.

VERSION	First edition · October 2026
FREE	The Selectability Playbooks and every template: rocmode.io/playbooks
RIGHTS	© 2026 Ben Stokes. All rights reserved.



SUMMARY

The argument in *brief*

Marketing can make more than ever, and is less sure than ever what it achieves.

Marketing never agreed what it was for, so it counted what it could see: awareness, reach, rankings, followers, and now AI visibility. But being seen was only ever a stand-in for being selected. In one large 2025 survey of big business purchases, each buying group weighed about five suppliers and picked one. All five were seen. Four lost.

So this book measures something else. **Selectability** is the odds you're picked. Visibility is the odds you're seen.

A buyer passes through seven stages on the way to a pick: Known, Found, Answered, Believed, Preferred, Bought and Kept. Being seen covers only the first two. A grocery shop turns on two or three of them. A software contract needs all seven, for several people.

AI is the problem, and much of the answer. The machine flooding the market is also the first adviser a marketer can question directly. So Selectability can now be estimated, stage by stage. The estimate is a stand-in too, but a better one.

This book sets out what a choice is made of, then takes the seven stages one at a time, with a toolkit for each you can use the same week.

It's written for marketing leaders, the chief executives and sales leaders who share their targets, and owners of smaller businesses, including those that sell to households.

Every figure is sourced at the back, with a note on who produced it and what they sell. One section near the end sets out what nobody knows yet.

INDEX

Contents

PART ONE	Why it feels like this	6
1	So much, so little	7
2	Marketing's one job	10
3	The visibility trap	14
PART TWO	What a choice is made of	18
4	The anatomy of a choice	19
5	The choice starts without you	24
6	Moments, not funnels	27
7	Five ways people buy	31
8	Why people choose	36
9	The new chooser	39
PART THREE	The way forward	43
	The open ground	44
10	One truth, many answers	47
11	A reason to choose you	51
12	Know your moments	56
13	Be known	61
14	Be found	72
15	Be answered	82
16	Be believed	93
17	Be preferred	102
18	Be bought	113
19	Be kept	123
PART FOUR	The builds	134
20	Content	138
21	The website	142

22	Search and AI answers	147
23	Advertising	151
24	Social	155
25	Creators	159
26	PR and reviews	163
27	The pitch	167
28	Email	171
29	Automation	175
	The rest, briefly	179

PART FIVE

	Knowing it worked	182
30	The Selectability Score	183
	Appendix: the self-score worksheet	193
	Why I wrote this	198
	What we don't know yet	201
	Coda	203
	Sources	204


Four pictures recur through this book: learn them once

1 The ladder




the seven stages of a choice

2 The map




a map of moments:
one buyer at one step

3 The process



a stage's method,
from its toolkit

4 The loop



a routine you repeat:
ask, fix, ask again

Four pictures recur through this book. Learn them once, and the rest read at a glance.

WHERE YOU ARE

01

Why it feels like this

02

What a choice is
made of

03

The way forward

04

The builds

05

Knowing it worked

01 / PART ONE

Why it feels like *this*



CHAPTER 01

So much, so *little*

I've sat in more of these meetings than I'd care to count. The creative agency down one side of the table, the media agency down the other, and the client at the head of it. Each agency brought its own deck, and every deck told the same story. This quarter, we were the hero.

Reach was up. Engagement was up. The cost of something or other was down. If sales hadn't moved, that was a problem for somebody else's slide.

I can't remember anyone ever saying "that didn't work", or saying what would change next quarter beyond a promise to keep optimising.

I'd love to tell you I was the honest one in the room. I wasn't. For the first ten of my twenty-five-odd years in agencies, I was as guilty as anyone. Building a report was a craft, and I was good at it.

I knew it wasn't the whole picture. But the attitude in a lot of agencies back then, mine included, was to keep the client, whatever it takes. The report's real job was to protect the account.

I suspect the clients knew. They just lacked the knowledge to point at where we were dodging, or to say what to hold us to instead.

That meeting is still running, in boardrooms everywhere. Every number is up, the business hasn't moved, and nobody can say which part worked. These days there's one new item on the agenda: what's our AI strategy?

Never before could marketing do so much, and so often achieve so little. You can make a year of content in a week, and you have more tools than your team has people.

The buyer has moved where you can't see

AI has changed how marketing is made. It's changing how people buy too, in a way your reports barely show.

AI is now where buyers go for answers. ChatGPT has 900 million weekly users ([TechCrunch](#)). And Google, still where most searching happens, now answers with AI first: more than a billion people a month use its AI Mode ([Google](#)).

Yet in your traffic AI hardly registers. Assistants mostly answer rather than link, and people rarely click the links they do give. On Google searches with an AI summary, its links were clicked in 1% of visits ([Pew Research Center](#)). Across 101,574 websites, ChatGPT sent about one visit in three hundred in May 2026 ([SE Ranking](#)).

+393%

AI-referred visits to US
retailers, on a year
earlier

ADOBE

That number misleads. The assistant does the comparing and names a few brands, and the buyer then reaches them another way. Semrush surveyed 1,030 US shoppers who use AI. After an AI recommendation, 40% searched Google for more and 28% went straight to the brand's site ([Semrush](#)). Those visits show up as search or direct, not AI.

So AI puts you on the shortlist long before it shows in your analytics. The visits it does send are growing fast and buy more. In Adobe's US retail data, they were up 393% on a year earlier and converted 42% better than other traffic ([TechCrunch](#)).

The urgency is real, but most of it is aimed at the wrong job. The race isn't to make more. It's to be the name in the answer, and there's still time to get this right.

The tools arrived faster than the results

You aren't alone. In Marketing Week's 2026 survey of 2,350 marketers, two in three had felt overwhelmed in the past year. More than half reported emotional exhaustion ([Marketing Week](#)).

Use of generative AI in marketing tripled in two years, from 7% of activities to 22% ([CMO Survey](#)). The return hasn't followed. In the same survey, no technology activity scored above five out of seven, including the return on the technology itself.

More of everything

Graphite checked 55,400 article URLs with three AI detectors. By its estimate, about half of new articles published online in early 2026 were written mainly by machines ([Graphite](#)).

Small businesses feel it most. Among those surveyed in four countries, only 18% feel very confident their marketing works, down from 27% a year earlier. Four in ten have less than an hour a day to spend on it ([Constant Contact](#)).

Nothing is wrong with you

If this is familiar, nothing is wrong with you. Something is wrong with the job you were handed.

AI didn't create the confusion. It multiplied it. To see why, start with a question marketing has never settled: what is it for?

Whose evidence this is. Graphite sells search marketing services. SE Ranking and Semrush sell search marketing software, and Semrush's survey reports what shoppers say they do. OpenAI and Google report their own user numbers. Adobe sells analytics and commerce software. Constant Contact sells email marketing to small businesses. Marketing Week is a trade magazine, and its survey reports opinion, not measured results. Pew Research Center is a nonprofit research organisation.

TRY THIS List every dashboard you report from, and write beside each one what it counts: who saw you, or who chose you.

IN SHORT · CHAPTER 1

WHAT'S TRUE

Marketing can do more than ever and achieve little, while buyers now choose through AI where your reports can't see.

DO THIS QUARTER

List every dashboard you report from, and write beside each whether it counts who saw you or who chose you.

WATCH

How many of the numbers you report count who chose you, not just who saw you.

Marketing's one job

Marketing has one job: to get the brand selected. Not seen, not liked, not followed. Selected.

Everything else either makes that more likely or guesses at whether it happened. So ask one question of every tool, channel and campaign: will this get us selected?

The "so that" test

Take any marketing activity and keep asking "so that what?" We post on LinkedIn, so that people in our market see us. So that they remember us. So that we're on the list when the need arises. So that they choose us.

Try it on a trade stand, a price cut or a new logo. Every chain, long or short, should end on the same sentence. One that ends anywhere else has no job to do.

Most of the work happens before anyone is buying

The Ehrenberg-Bass Institute, whose body of research is the most tested in marketing, has long argued that brands grow by being easy to think of and easy to buy. Professor John Dawes estimates that up to 95% of business buyers aren't in the market for a given product at any one time ([Ehrenberg-Bass](#)). He's clear that the figure is a rule of thumb that varies by category, not a precise count.

So most marketing reaches people who won't buy for months or years. It works by leaving a memory that surfaces when they finally do. A 2025 study fits: in large business purchases, buyers wrote their shortlist on day one, mostly from memory.

The stand-ins

A choice made from memory, months later, leaves nothing to count at the time. So marketing counted what it could see.

Up to
95%

of business buyers aren't
in the market at any one
time

EHRENBERG-BASS

WHAT WE COUNT

WHAT IT STANDS IN FOR

Reach and impressions

They might have noticed us

WHAT WE COUNT	WHAT IT STANDS IN FOR
Awareness	They might remember us when it matters
Traffic	They might be considering us
Engagement	They might care
Leads	They might be choosing
Pipeline	They might choose us

Sales, market share and win rate are missing from that table on purpose. They aren't stand-ins. They're the outcome, counted after the fact.

What nobody could see was the odds beforehand: at this moment, with this buyer, how likely were we to be picked? Every row above is a "might", a way of estimating those odds. Being seen matters, because nobody chooses a brand they've never met. But it's the first step, and it has long been treated as the whole job.

What a choice is

A choice has four parts: one or more buyers, a moment of need, a set of alternatives and a pick. The alternatives always include doing nothing.

Sometimes the alternative isn't another firm at all. In 2023 a major airport asked us to quote for a digital gift finder, a short quiz to help shoppers choose a present in time for Christmas. It came through a warm referral, and we had a quote back within a fortnight of the brief. About a month later, after a wait on a key person's leave, the answer came: they'd bought an off-the-shelf product instead. We hadn't lost to another agency. We'd lost to a different kind of answer, one that was never on our list of competitors.

Two things follow. First, the job is bigger than promotion. Of the old marketing mix, the 4Ps of product, price, place and promotion, only one is promotion. What you sell, what it costs and where it can be bought decide choices as surely as any campaign.

Second, the job doesn't end at the first sale. A renewal is a choice. So is a recommendation to a friend.

Selectability

Selected is the outcome. Selectability is the odds of it.

I define **Selectability** as the odds that a buyer, at a given moment, picks you over the alternatives, doing nothing included. A moment is one kind of buyer at one step of their journey.

Its counterpart is **visibility**: the odds that a buyer, or a machine, notices you. A brand can be seen everywhere and selected almost nowhere.

Selectability is also specific. You can be highly selectable to one buyer at one moment, and barely selectable to another. One number for the whole brand hides that.

People have tried to estimate these odds for decades, through brand-equity models such as Kantar's Meaningful Different framework, which predicts a brand's share from surveys of how buyers see it ([Kantar](#)), conjoint studies and win-loss reviews. Those average a market, stage a choice or look back at one deal. What they lack is the moment. What's new is that AI assistants now advise buyers, and can be asked those buyers' questions directly, as often as you like. Part five turns that into the Selectability Score.

The Selectability Score is itself a stand-in. It samples what buyers and machines say; it doesn't observe the sale. It's better because it measures the choice at the moment, not exposure, and it comes with honesty rules. Your own sales data always outranks it.

Not only marketing's job

Read the first sentence of this chapter again with "sales" in place of "marketing". It still holds. So the job is simple, and the best research has said so for years. How, then, did we build an industry where every team hits its number while the brand loses the choice?

Whose evidence this is. The Ehrenberg-Bass Institute sells research to sponsor brands. Kantar sells brand tracking, including the framework named here. There's no vendor survey data in this chapter.

TRY THIS Run the "so that" test on your five biggest budget lines. Mark any chain that never reaches "so that they choose us".

IN SHORT • CHAPTER 2

WHAT'S TRUE

Marketing has one job: to get the brand selected.
Not seen, not liked, not followed. Selected.

DO THIS QUARTER

Run the "so that" test on your five biggest budget lines, and mark any chain that never reaches a choice.

WATCH

The share of your budget whose chain ends on "so that they choose us".

The visibility trap

Every era of marketing has had its visibility number. The channels changed. The habit didn't.

ERA	THE VISIBILITY NUMBER	WHAT IT COULD NOT TELL YOU
Broadcast	Prompted awareness and recall	Whether you came to mind, unprompted, when the need arose
Media buying	Reach and impressions	Whether anyone noticed
Search	Rankings and traffic	Whether the visitors were buyers
Social	Followers and engagement	Whether the fans were customers
AI	Mentions, citations and share of voice	Whether being named meant being recommended

The first row needs a distinction. Recognising a name from a list, prompted awareness, is the stand-in. Coming to mind unprompted in a buying situation is the real thing. Professor Jenni Romaniuk and her Ehrenberg-Bass colleagues call that second kind mental availability ([Ehrenberg-Bass](#)).

AI visibility is the newest row, not a new idea. A brand can be named in an answer and still not be the one the answer recommends. None of these numbers was foolish. Each was the best view available of a choice nobody could watch.

A stand-in becomes dangerous the moment someone is paid to raise it. Economists know this as Goodhart's law: a measure that becomes a target stops being a good measure. Marketing built an industry on top of it.

How channels became departments

I've worked in most corners of marketing. Digital marketing. Digital experience. Sales process design. Customer experience. A full-service creative agency.

Each had its own scoreboard, and each sensibly chose numbers it could move. There's no point being held to account for something you can't control.

On websites, that meant engagement, time on site and enquiries. What happened after the enquiry wasn't ours. There were times, I'm not proud to say, when we pointed straight at the sales team. Enquiries were up. We'd done our job.

Every other discipline worked the same way. Each reported its own slice fairly, and even with all the slices on the table you couldn't see the whole picture.

The conversation that really mattered was one nobody wanted to have. Should the client be paying for our specialty at all, right now? Not how to improve it. Whether to do it. I can't remember anyone asking it in a meeting I was in.

Most brands did a bit of everything, and quite possibly nothing wonderfully. Nobody was asking whether any of it got the client chosen. And nobody would open that kettle of fish, in case theirs turned out to be the service the client didn't need.

Each new channel arrived with its own number, then a specialist, a team, a budget and an agency, all judged on that number.

TEAM	THE NUMBER IT OWNS	HOW THE NUMBER RISES WHILE THE CHOICE IS LOST
Search	Rankings and organic traffic	Rank for questions nobody asks before buying
Social	Followers and engagement	Win applause from people who will never buy
Paid media	Clicks and attributed sales	Pay to reach buyers who were already coming
Content	Downloads and leads	Collect email addresses from students and rivals
Brand	Awareness	Be known, but for nothing in particular
Sales	Pipeline	Fill it with deals that were never going to close

Now hand every team in that table a machine that makes more of whatever it's measured on. That's what AI has done: multiplied the output and left the aim untouched.

Every team owns a number. Nobody owns the choice.




FIGURE 3.1 Seven stages · which team counts each

The number was real. The cause was not.

eBay ran one of the clearest published tests of this. Its reports showed paid search ads driving sales, because clicks on ads were followed by purchases. Its

economists switched the ads off across 68 US regions for 60 days and compared sales with the regions left on ([Berkeley Haas](#)). Sales showed no measurable difference. Separate tests of ads on eBay's own name found no measurable short-term benefit either. People typing "eBay" were already on their way there.

Ads on other search terms did win some new and infrequent shoppers, but most of the money went on frequent customers who would have bought anyway, so the average return was negative ([Blake, Nosko and Tadelis](#)).

The lesson isn't that paid search fails; for an unknown brand it may be the best money spent. eBay was paying to be seen by people who had already chosen it, and its reports couldn't tell the difference.

That study is a decade old. The trap isn't. An AI visibility score counts how often you're named, which isn't how often you're selected.

Why the trap pulls money the wrong way

Marketing's real work is done before the buyer shows up, where it's hardest to count. The activity nearest the sale is the easiest to measure, and budgets follow what can be measured. So money drifts toward the end of the journey, where the choice has mostly formed, and away from the beginning, where it's won. The better the dashboard, the stronger the drift.

What it looks like from the inside

The monthly review where every slide is green and revenue is flat. The standing argument between marketing and sales about lead quality. The agency report that proves its own channel worked. Each is the same problem: the organisation is shaped like its channels, and the buyer's decision isn't. Every team owns a number, and nobody owns the choice.

The people who fund marketing have noticed. In a 2024 Gartner survey, only 52% of senior marketing leaders said they could prove marketing's value and get credit for it ([Gartner](#)). Chief financial officers and chief executives were the most sceptical.

Visibility can't answer them, because a number that counts who saw you will never show who picked you. It's an early sign of Selectability, never the result.

AI is sorting the selected from the seen

AI is now making the difference between seen and selected visible. Two travel companies show it, with a caution: Tripadvisor is an intermediary people pass through, and Booking is a destination where they book. This shows what AI does to each kind, not a fair test of one strategy against another.

52%

of senior marketing leaders could prove marketing's value and get credit

GARTNER

Tripadvisor built its hotel business on being found: travellers searched, landed on a review page, then clicked through to book elsewhere. In February 2026 its chief executive named the rise of AI overviews as one reason for falling visits ([Skift](#)).

By the second quarter, its hotel revenue was down 23% on a year earlier, though the company blamed a fluctuating economy ([Tripadvisor](#)), and its troubles with search are older than AI. So AI isn't the whole cause.

Booking Holdings is the name travellers go to on purpose. About two-thirds of its consumer bookings come direct. Its chief executive said its brands were "showing up an awful lot" in AI assistants. Yet its finance chief put bookings from them at "less than 1% of total room nights" ([PhocusWire](#)).

Both companies were seen. One was selected.

AI didn't take Booking's sale. It took Tripadvisor's place as the page people passed through. In the few cases so far, AI is sorting the brands buyers pick on purpose from the ones they only pass through.

The way out is to measure Selectability itself. That starts with a question marketing has rarely asked: what is a choice actually made of?

Whose evidence this is. Gartner sells research and advice to marketing leaders. Skift and PhocusWire are travel trade press, and Tripadvisor's figures come from its own filing. Blake, Nosko and Tadelis used eBay's own data, in an NBER working paper.

TRY THIS Ask each team which of the seven stages its number measures, and list the stages nobody owns.

IN SHORT · CHAPTER 3

WHAT'S TRUE

Every team owns a number that counts who saw you, and nobody owns the choice.

DO THIS QUARTER

Ask each team which of the seven stages its number measures, and list the stages nobody owns.

WATCH

Whether every stage has an owner, and whether green reports start moving with revenue.

WHERE YOU ARE

01

Why it feels like this

02

**What a choice is
made of**

03

The way forward

04

The builds

05

Knowing it worked

02 / PART TWO

What a choice is *made* of



CHAPTER 04

The anatomy of a *choice*

Cos Luccitti and I started CXA in December 2015. Not long after, I drafted a model of the customer journey in five stages: Discovery, Deliberation, Decision, Delivery and Devotion. Five Ds.

I took it to Cos and we worked it over together. We changed a few words, but mostly we just looked at it and liked what we saw. We both fell a little bit in love with it.

It made sense of everything we did as an agency, which until then had looked like a list of services. And it gave clients a reason to get excited about marketing and customer experience again. Customer experience in particular had shrunk to a satisfaction score and how well the service desk performed. We wanted it to cover the whole life of a customer, from the first time they heard of you to the day they decided to stay.

The sixth came later, once we'd started asking one question first: what problem is this solving? If every product and service is really a solution to a problem, then the natural way to grow with a client is to solve more of their problems. So we added Design, and put it at the front of the journey, ahead of Discovery.

We didn't get everything right. The irony is that we underplayed Delivery. The rest of the CX industry was so fixated on service delivery that we'd pushed hard the other way, and in doing it we gave too little weight to the stage where you actually keep your promise.

The model in this chapter is that one's descendant. Ten years on, I've rebuilt it around a single question: what has to be true for a buyer to choose you? And this time, Delivery has a stage of its own. It's called Kept.

A choice isn't one event. It's seven things a buyer comes to say about you, and being seen earns only the first two.

Seven statements

STAGE	WHAT THE BUYER CAN SAY
Known	"I've heard of you."
Found	"You show up when I look."
Answered	"You answered my question."
Believed	"I trust that it's true."
Preferred	"You're my first choice."

STAGE	WHAT THE BUYER CAN SAY
Bought	"I bought."
Kept	"I'd choose you again."

The seventh is the first six asked again, when a customer decides whether to stay.

Each one can fail on its own. A brand can be known and impossible to find. It can answer every question and not be believed. It can even be the favourite and still lose the sale to a clumsy checkout or a slow contract. Preferred isn't the same as Bought.

This is not a funnel

Marketing has drawn staged models for a century, from AIDA (attention, interest, desire, action) to Robert Lavidge and Gary Steiner's hierarchy of effects in 1961. They say buyers climb the steps in order. The evidence doesn't support that. Demetrios Vakratsas and Tim Ambler reviewed more than 250 articles and books on how advertising works. They found "little support for any hierarchy, in the sense of temporal sequence, of effects" ([Vakratsas and Ambler](#)).

So the seven stages aren't a sequence. They're conditions. Each has to hold at the moment of choice, in whatever order the buyer meets them. That's why each stage gets its own chapter and its own measure.

One buyer, seven stages

Take an operations director choosing scheduling software. She's an illustration, not a case.

- ✓ **Known.** Four names come to mind before she looks anywhere.
- ✓ **Found.** She searches and asks an assistant. Three of her four turn up, and one new name joins them.
- ✓ **Answered.** Two suppliers publish prices and integration details. The others make her book a call to find out.
- ✓ **Believed.** Reviews and a peer's word back up both of the two that answered.
- ✓ **Preferred.** Her team favours one, for a reason she can repeat to the finance lead.
- ✓ **Bought.** The favourite's contract takes six weeks of legal back-and-forth. The runner-up offers a 30-day pilot she can sign this month.
- ✓ **Kept.** A year later, she asks the first six again at renewal.

The favourite won five stages and lost at Bought, and nothing on a visibility report would have shown it.

Seen is two of seven

Every visibility number measures the first two stages. Awareness is Known. Rankings, impressions and AI mentions are Found.




FIGURE 4.1 Seven stages of a choice · seen is the first two

Both are necessary. But a supplier that makes the shortlist and loses has passed both. It lost somewhere in the other five.

Visibility is the door. The choice is made inside the house.

Three things around the stages

Three other things define the choice itself.

- / **The chooser.** It isn't always the end buyer. A retailer chooses what to stock, a broker what to recommend, a committee what to approve. An AI assistant is the newest name on that list.
- / **The moment.** Every choice starts with a need arising.
- / **The offer.** What you sell, what it costs and where it can be bought: the old marketing mix. No marketing makes a poor offer preferred for long.

A loop, not a line

Buyers don't walk the stages in order. A repeat customer goes straight from Kept to Bought. A happy one tells a colleague, and Kept becomes someone else's Known.

That's why the experience after the sale belongs in a model of selection. Whether you deliver what you promised decides whether you're kept. What your customers say decides whether the next buyer believes you.

A loop, not a line: Kept feeds the next buyer




FIGURE 4.2 The loop · Kept feeds the next buyer

What AI changes here

Each statement now has two audiences: the buyer, and the machine advising the buyer. A person remembers you or doesn't. An AI model has its own version of memory, built from what has been written about you.

So you can ask the adviser directly whether it knows the brand, what it says, and whether it recommends it. Several of the seven can now be inspected at any time, not inferred from surveys.

My reading of the evidence so far is that AI acts most on Found, Answered and Believed, and mostly in purchases that take real thought. For habitual and quick buying it's too early to know.

Every activity has a home

I mapped ninety marketing and sales activities against the model. Sixty-five sit inside one of the seven stages. The rest belong to the chooser, the moment, the offer, or the groundwork that serves them all. The full mapping is in the toolkit at rocmode.io/playbooks.

How the pieces fit

One choice: four parts, seven stages, five modes, three foundations




FIGURE 4.3 How the pieces fit · one page

The ladder is what has to hold. The mode says which rungs carry the weight. The foundations are what every rung stands on.

Seven stages, then. The next question is when they happen, and the answer is earlier than most marketing assumes.

Whose evidence this is. There's no vendor data in this chapter. Peer-reviewed source: Vakratsas and Ambler, in the *Journal of Marketing*.

TRY THIS

Take your last lost deal and mark the one stage where it was lost. The activity mapping is free at rocmode.io/playbooks.

IN SHORT · CHAPTER 4

WHAT'S TRUE

A choice is seven things a buyer comes to say about you. Being seen earns only the first two.

DO THIS QUARTER

Take your last lost deal and mark the one stage where it was lost.

WATCH

Which stages your lost deals keep failing at, beyond Known and Found.

CHAPTER 05

The choice starts *without* you

Being seen feels like winning. Your brand appears in the search results, in the AI answer, on the shortlist. The report turns green.

But by the time a buyer speaks to you, the choice is already under way. The enquiry that marketing celebrates and sales chases comes late in the decision, not at its start.

Go back to the operations director choosing scheduling software. Before she emails a single supplier she has four names in her head. Two she has used, one she evaluated years ago, and one a colleague swears by. Her team spends months reading and comparing. Then she calls the supplier they already prefer.

She's an illustration, but the research below measures the pattern. The choice starts without you, and it isn't finished when you arrive.

The shortlist is written on day one

In 2025, 6sense surveyed about 4,000 people who buy for their organisations (6sense). Every purchase was worth at least US\$25,000. It found:

- Buyers evaluated about five suppliers each time, and chose one.
- The winning supplier was on the buyer's first-day shortlist 95% of the time.
- 94% of buying groups had ranked that shortlist before speaking to any seller.
- The supplier they favoured before first contact won about four deals in five.
- Buyers made the first approach themselves in 79% of cases, roughly 60% of the way through.

Buyers had dealt before with almost four of the five suppliers they evaluated. In this kind of purchase the shortlist is remembered, not discovered.

One vendor's survey is thin ground, so here's a second. Bain and Google asked 1,208 US business buyers in 2022. Some 86% began with a list in mind, and 92% bought from it ([Google, via Marketing Dive](#)).

95%

of the time, the winner was on the first-day shortlist

6SENSE

But the list is not locked

That evidence describes one kind of purchase: large, experienced and decided by a group. Most buying isn't like that.

Where one or two people make a considered purchase, the early list has a strong head start but a real second window. McKinsey's consumer decision journey research, published in 2009, studied almost 20,000 consumers. A brand in the buyer's initial set was up to three times more likely to be bought. Yet car shoppers in that study started with 3.8 brands and added 2.2 while evaluating (McKinsey). The data is old, but it's still the best-known measure of that second window.

Google's "messy middle" research, in a simulated shopping test, found a fictional challenger brand took 28% of preference from an established favourite when it carried strong reviews and benefits ([Think with Google](#)).

In habitual buying there's barely a shortlist at all. The choice is near-automatic, made from whatever comes to mind first and is easiest to buy.

Small purchases follow the same shape, only faster. Take a café owner looking for a bookkeeper, an invented example this book comes back to. He asks two other owners and an assistant, has three names by lunchtime, and calls one that afternoon. His list still forms before any supplier knows he's looking. It just forms in hours, and one good answer can join it. I found no study that measures this for small services, so treat it as illustration.

Even the largest purchases stay open. In the 6sense study, about one buyer in five changed their mind after talking to sellers. So the early list matters everywhere, and nowhere is it final.

What AI has changed

The shortlist habit is older than AI. What AI is changing is who writes the list.

The list is moving from the buyer's memory to the machine. G2 surveyed 1,038 business software buyers in June 2026. More than 80% had taken software recommendations from an AI chatbot in the past two years ([G2](#)).

The list still decides. Among buyers who took recommendations from AI, 80% bought from their first shortlist in at least three of their last five purchases. Among those who didn't, 65% did.

That's an association, not proof, but it suggests AI hasn't loosened the day-one list, and may have started to write it. The largest deals lag: in 6sense's study buyers used AI mostly mid-journey, to compare suppliers already listed.

A list written by a machine can also be rewritten by evidence. In one experiment, AI models dropped a famous brand when a rival offered better specifications

(Chapter 13).

AI has moved the choosing further out of sight, and made counting who saw you an even weaker guide to who chose you.

What that adds up to

You're invited in late. By then you're confirming a preference or overturning one. The choice can still be won, but mostly before you're in the room. That raises the question of where. No buyer comes to know you, believe you or prefer you in general. They do it at a particular moment, with a particular need.

Whose evidence this is. 6sense sells intent-data and account-based marketing software. G2 runs a software review marketplace. Google sells advertising; it partnered on the Bain study and wrote the messy-middle research, and Marketing Dive ran the Bain figures as a Google-sponsored post. Bain and McKinsey sell consulting. There's no peer-reviewed source in this chapter.

TRY THIS Ask your last five new customers which names were on their list before they contacted you.

IN SHORT · CHAPTER 5

WHAT'S TRUE

The choice starts before buyers speak to you. In big purchases, the shortlist is written on day one.

DO THIS QUARTER

Ask your last five new customers which names were on their list before they contacted you.

WATCH

How often you're on buyers' lists before they get in touch, and how often you win from there.

CHAPTER 06

Moments, not *funnels*

A funnel is a picture of your process. A moment is a picture of the buyer's. Only one of them is where a choice gets made.

What the funnel gets wrong

The funnel assumes buyers enter at the top and move down in order. In large purchases they arrive most of the way through, with a ranked list.

It also averages everyone together. It can't say which need brought them, or which need you keep losing.

What a moment is

A moment is one kind of buyer at one step of their journey, with a need and a question. "Which scheduling tool suits a field team of forty?" "Who can refinance my loan before the fixed rate ends?"

The Ehrenberg-Bass Institute calls the triggers behind these category entry points: the cues buyers use to reach into memory when a buying situation arises ([Ehrenberg-Bass](#)). Professor Jenni Romaniuk, who developed the idea, stresses that they're about the buyer and not the brand.

A trigger and a moment are different units. The trigger is the cue in memory: "wages ate the takings again." The moment is a cell on the map: a café owner, comparing options, after that trigger. One trigger can open several moments.

Parts of this aren't new, and the credit matters. Journey mapping has crossed buyer types with journey steps for years. Google named the "Zero Moment of Truth" in 2011: the research a shopper does before buying ([Google](#)). In 2015 it added "micro-moments", "intent-rich moments when decisions are made and preferences shaped" ([Google](#)). What this book adds is the scoring. Each moment is checked stage by stage, and AI makes each one something you can ask.

AI made the moment visible

For two decades buyers squeezed their need into a few keywords. Given a machine that talks back, they say more.

23 vs 4.2

words per ChatGPT prompt, with search off versus on

SEMRUSH

Semrush studied more than 80 million clickstream records from ChatGPT (Semrush). With its search feature switched off, prompts averaged 23 words. With search on, they averaged 4.2, much like a classic search. That's one snapshot, and the gap is still moving.

Those extra words are the moment itself: the situation, the constraints, the worry. People tell a machine what they once told only a salesperson.

The map

Marketers have long aimed for the right message, to the right person, at the right time. Segmentation listed the people and journey mapping listed the steps. Few lists went down to the question each buyer asks at each step. The map is that list. List the steps of the journey down the side and the kinds of buyer across the top. Each square is one buyer at one step. That square is a moment, and in each moment the seven stages either hold or fail.

Here's a filled one, for the bookkeeper for cafés. She's an invented example. Her offer is "Bookkeeping for cafés. Your wage cost against takings, every Monday by nine."

Fifteen moments, three worth winning first

	Owner of one café	Owner of two or more cafés	New owner, just bought a café <i>decide whether she wants these at all</i>
1 Need arises	"Wages ate the takings again this month."	"We're opening a second café."	"I've just bought a café and need to set up the books."
2 Learning	"What percentage of takings should café wages be?"	"How do I compare wage cost across two sites?"	"What does a café bookkeeper actually do?"
3 Comparing	"Who can show me wage cost against sales every week?" chosen	"I own two cafés and wages are eating my margin. Who can show me wage cost every week?" chosen	"Do I need a bookkeeper, or just better software?"
4 Buying	"What does a café bookkeeper cost, and can I stop if it doesn't help?"	"Which bookkeeper can take on both sites from next month?"	"Who can set up the books before my first tax return?"
5 Renew or switch	"Is the Monday report still worth the fee?" chosen	"Our accountant now offers bookkeeping. Should we move it?"	"First year done. Do I still need a bookkeeper?"

Invented example. Every question is illustrative.

FIGURE 6.1 The bookkeeper's map of moments · three buyers by five journey steps

The map shows her where to start. Her offer answers one worry, wage cost, and that worry peaks while owners compare. So she chooses three cells and leaves the rest for later.

That grid is your market. It turns a vague ambition, "grow awareness", into a specific one: win this moment, for this buyer, at this stage.

Start small: fifteen moments is enough to see where you're losing. Chapter 12 shows how to build your own.

In the age of AI the old maxim gains a clause: the right answer, to the right question, at the right time.

The moment nobody maps

One moment appears in every market and on almost no map: renew or switch. Most maps leave it off, because no team owns it. Marketing stops at the sale, and sales moves to the next deal.

Two ways to grow

A brand grows by winning more of the moments that already exist, or by making new ones. Category creation, problem education and seasonal campaigns all make moments. They give a reason to ask to people who had none.

When nobody is searching for what you sell, that's the first job: create the moments.

Seen versus selected

A brand can appear a million times and still be absent from the twenty questions that decide its sales. That's why Selectability is counted moment by moment, never in general.

A map of moments shows where choices are made. It doesn't yet say how they're made, and that differs more than most marketing advice admits.

Whose evidence this is. Semrush sells SEO and AI-visibility tools. Google sells advertising, and named micro-moments and the Zero Moment of Truth in its own marketing. The Ehrenberg-Bass Institute sells research to sponsor brands.

TRY THIS Draw a 3 × 5 grid for your business and write one buyer question, in their words, in each square. The stage tools at rocmode.io/playbooks all start from this map.

IN SHORT • CHAPTER 6

WHAT'S TRUE

Choices are made at moments, one kind of buyer at one step, not in a funnel drawn around your process.

DO THIS QUARTER

Draw a three-by-five grid of buyers and journey steps, and write one buyer question in each square.

WATCH

Whether you're selected in the few moments you chose, not how often you appeared overall.

Five ways people buy

People buy in five recognisably different ways. The mode decides what marketing should do, and how much AI matters.

Three questions that sort every purchase

Three questions separate one purchase from another, each studied for more than fifty years.

QUESTION	THE RANGE	WHERE THE IDEA COMES FROM
How much is at stake?	A snack, up to a six-figure contract	Howard and Sheth, 1969
Is it a first purchase or a repeat?	Something new, through to the usual re-order	Robinson, Faris and Wind, 1967
How many people decide?	One person, up to a committee	Webster and Wind, 1972

Robinson, Faris and Wind’s straight rebuy is what this book calls habitual buying. Their modified rebuy is the renew-or-switch moment, which cuts across every mode.

The five modes

Together, the three questions give five modes.

MODE	LOOKS LIKE	WHO DECIDES	WHAT WINS IT
Habitual	Groceries, fuel, the usual re-order	One person, on autopilot, across a few familiar brands	Coming to mind and being within reach
Quick choice	Something cheap you have not bought before	One person, in minutes	Being the name in front of them at that moment
Considered personal	A car, a home loan, a holiday	One or two people, over days or weeks	Answering their questions and being believed
Owner-led business	A bookkeeper, a CRM, an agency	One to four people, over weeks	Being believed by the owner, plus a business case a second voice can check
Committee	Enterprise software, a major contract	A group, over months	Winning each person, and surviving scrutiny

Habitual doesn't mean loyal to one brand. Ehrenberg-Bass researchers find buyers are loyal "to a repertoire of brands, not to a single brand" ([Ehrenberg-Bass](#)). The habit is a short list, and the job is to stay on it.

From mid-2017 to mid-2019 I was Head of Digital at CUA, the bank now called Great Southern Bank. Most weeks there, I bought in three completely different ways.

The coffee was habit. Every morning, sometimes twice a day, I went next door to the same café, to the same barista, and ordered the same thing with coconut milk. Once you've found your coffee, you don't switch. I didn't compare and I didn't think about it, and I liked it that way. It felt comfortable and reassuring, a bit like home.

Lunch was the opposite. I'd walk out into the city with no plan and duck into the first place that caught my eye. It felt exciting, and I came back to my desk with more energy than I left with.

And then there were the committees. Sigh. I hated them. Slow, drawn out, over-complicated and over-justified. The one I remember most is the company CRM. That process ran for more than a year, and it was still going when I left. It was a big, important decision, and procurement really did matter. But my gosh, the frustration was real. I'd walk out of those meetings tired and confused, and often not much caring which way it went.

Same person, same week, three completely different buyers.

To find your mode, picture your typical sale and follow the tree.




FIGURE 7.1 Find your mode · a decision tree

Not every purchase needs all seven

In all five modes, Known and Found are gates: if the buyer can't recall you or find you, nothing else you do is seen. What differs is what carries the weight once a buyer is through.




FIGURE 7.2 Five buying modes · the stages each leans on

So there's no universal weakest link. A chocolate bar needs to be remembered, then within reach. A software contract has to win all seven, several times over. Knowing your mode, and which stages carry it, is the first act of strategy. Chapter 30 turns these weights into numbers in the Selectability Score.

Where the small business fits

76%

of firms up to 250 staff
decide with four or fewer
people
G2

An owner-led business doesn't buy like an enterprise. In G2's 2025 study, about 76% of businesses with up to 250 staff made software decisions with four or fewer people (G2). Among firms of 1,000 to 5,000 staff, about 37% did. Small firms turned first to Google (47%), ahead of review sites (38%) and AI search (35%).

In a 2026 Capterra study, about half of whose respondents were small businesses, successful software buyers shortlisted three suppliers and decided within three months (Capterra).

So an owner choosing a bookkeeper decides more like someone buying a car than like a committee. That's my reading, not a tested finding. The difference is business risk, and often a second voice: a partner, an accountant, a finance lead. So the owner needs proof from a business like theirs, and numbers the second voice can check.

What AI changes in each mode

MODE	WHAT THE EVIDENCE SAYS AI CHANGES	HOW STRONG
Habitual	Little sign that it changes brand choice yet. My read is that it will reinforce habit before it disrupts it.	Thin
Quick choice	AI summaries cut the clicking, so the brand named in the answer probably gains. No purchase data yet.	Thin
Considered personal	The clearest effect. Complex categories draw the most AI referrals, and there those visitors convert better than most channels.	Moderate to strong
Owner-led business	Google still leads, with AI close behind. It behaves like a considered personal purchase with a business case.	Moderate, by inference
Committee	AI shapes and speeds the shortlist. The day-one list still wins, and evaluation takes longer.	Moderate

AI's hand is strongest where the buyer has the most to work out. In the two low-effort modes, nobody yet knows.

In committee buying, AI has made suppliers easier to find and harder to buy from. In G2's 2026 survey, 40% of buyers named evaluation as the longest stage. And 49% had seen a finance chief veto a purchase already approved (G2).

What is established and what is mine

The three questions are established research. The five modes are my own synthesis of them, tested against the evidence on AI from 2025 and 2026. I found no accepted buying typology for the AI era. Renewal isn't a sixth mode: every mode has one, and the customer can decide differently.

In every mode, somebody advises the buyer. Increasingly, the first adviser isn't a person.

Whose evidence this is. G2 runs a software review marketplace. Capterra is a software review site. The Ehrenberg-Bass Institute sells research to sponsor brands. The three questions come from academic work by Howard and Sheth, Robinson, Faris and Wind, and Webster and Wind.

TRY THIS Take your three biggest products or services through the decision tree, and write each one's mode beside its name.

IN SHORT • CHAPTER 7

WHAT'S TRUE	People buy in five different ways, and the mode decides which stages carry the weight and how much AI matters.
DO THIS QUARTER	Take your three biggest products through the decision tree, and write each one's mode beside its name.
WATCH	Whether your spending follows the stages that carry the weight in each mode.

CHAPTER 08

Why people choose

So far a choice has been a list, a moment and a way of buying. It's also a person, and people don't choose the way a spreadsheet does.

Good enough beats best

People rarely compare every option. They take the first one that clears the bar. Herbert Simon called this satisficing. It's why buyers pick from a list written on day one. So being on that list matters more than being slightly better than whoever is on it.

Familiar feels safe

A name people have met before feels like the safer pick, even when they can't say why. In 1968 Robert Zajonc showed that "mere repeated exposure" to something improves people's attitude towards it ([Zajonc](#)).

Feeling first, reasons after

It was around 2020, and I'd made up my mind about Intercom long before I compared anything.

I'd been on their website a few times and I loved it. The story pulled me in. The design was beautiful, and the product was explained completely without ever going on too long. I'd been building websites for more than twenty years by then, and I was impressed. Some of the suppliers I dealt with used it too, so I'd been on the other end of it as a customer, and I'd enjoyed that.

So when the question came up of which chatbot to put on websites, ours and our clients', it was Intercom's to lose.

We did a proper competitive review. It showed Intercom was well and truly the most expensive, about four times the price of the others in most cases. I argued for it anyway, for my own business and to clients, and most of them took the bite and bought it too.

Being fair to the numbers, it probably shouldn't have won. The alternatives cost much less and were good enough for most of what people needed. But Intercom had already won the emotional argument with me. I'd made the decision on feeling, then gone looking for the rational case to back it up.

Almost 2x

Emotional campaigns' odds of very large profit effects versus rational ones

IPA

The review was real work. The answer was in my head before it started.

Les Binet and Peter Field studied campaigns in the IPA's effectiveness databank. Emotional campaigns were almost twice as likely as rational ones to report very large profit effects ([IPA](#)). Their later business-to-business work found emotion did more there too, over the long term (Chapter 23).

Those are award entries, successes by selection, but they fit what psychologists find: people are poor witnesses to their own reasons.

Reasons still matter, because a buyer has to explain the choice to a partner, a boss or a board. The feeling picks, and the reason defends.

Losses loom larger than gains

A loss hurts about twice as much as an equal gain pleases, though less when the stakes are small. So in a serious purchase, the buyer fears choosing wrong more than they hope to choose right. That's why so many deals end in doing nothing.

We follow other people

In 2006 three researchers built an online music market for 14,341 people. Only some could see what others had downloaded ([Salganik, Dodds and Watts](#)). Where people could see each other's choices, the hits were bigger and harder to predict. Quality sets the range. Other people's choices decide where in the range you land.

A hard choice stalls

The famous finding that too many options paralyse people is weaker than its fame. A review of 50 experiments found an average effect of almost zero ([Scheibehenne and colleagues](#)).

A later meta-analysis of 99 results found when it does bite ([Chernev, Böckenholt and Goodman](#)): complex options, a hard decision, and a buyer unsure what they want. So the number of options isn't the problem. The difficulty of choosing is.

What it asks of you, stage by stage

STAGE	THE PSYCHOLOGY AT WORK	WHAT IT ASKS OF YOU
Known	Familiar feels safe	Show up often, and look the same every time
Found	Good enough beats best	Be on the first list
Answered	A hard choice stalls	Make your answer the easy one to use

STAGE	THE PSYCHOLOGY AT WORK	WHAT IT ASKS OF YOU
Believed	We follow other people	Let other people say it for you
Preferred	Feeling first, reasons after	Give a feeling, and a reason to defend it
Bought	Losses loom larger	Take the risk out of the last step
Kept	Good enough beats best	Keep clearing the bar, so they never start looking

The table is my reading of the evidence, not a measured result.

The machine has no feelings, though it has habits of its own. It can do the comparing, but it's still the person who feels, fears and decides.

Whose evidence this is. The IPA is an agency trade body. LinkedIn, which funded the B2B work, sells advertising. The other sources are peer-reviewed.

TRY THIS Write the feeling your offer should leave, and the one sentence a buyer would use to defend it to their boss.

IN SHORT · CHAPTER 8

WHAT'S TRUE

People don't choose like a spreadsheet. The feeling picks, and the reason defends.

DO THIS QUARTER

Write the feeling your offer should leave, and the one sentence a buyer would use to defend it.

WATCH

Whether buyers repeat that sentence, in their own words, when they explain choosing you.

CHAPTER 09

The new chooser

In October 2025, I asked an AI assistant a simple question, the kind a buyer would ask: "Who are the top digital agencies in Brisbane, Australia?"

It came back with about ten names. I'd heard of most of them. Evolut wasn't one. Not in the list, not in a footnote, nowhere.

I laughed. Then it stopped being quite so funny. I felt we did better work than most of the agencies on that list. But they were clearly doing a better job of being found.

It wasn't for lack of content. Our website is far denser than most agencies' sites, and it's full of case studies. The problem was that almost all of it was us talking about us. There was little or nothing outside our own site to back any of it up, and the machine had no one else's word to go on.

Around then I decided I was going to master this new game of getting brands into AI answers.

For the first time, something other than a person is reading your market. It reads far more of it than any buyer could, and it will tell you what it thinks. It isn't a careful referee, though.

The newest adviser

Buyers have always leaned on advisers: the retailer who decides what goes on the shelf, the broker who recommends a lender, the friend who knows cars. An AI assistant is the newest. It isn't yet the buyer, but it's moving that way. Today people use it to compare options, then decide for themselves. It already writes much of the shortlist, and that's most of the choice. Next it will place the order, starting with the repeat purchase and the routine reorder.

In a 2026 Ipsos study across 15 markets, 27% of consumers who knew of AI tools used them for product research. Only 9% would let them buy without asking ([Ipsos](#)). That's an attitude, not observed behaviour.

Use is growing fast. BCG found nearly a third of consumers now use AI somewhere in a purchase, about triple the rate of 18 months earlier ([BCG](#)).

You can ask it

No marketer could ever interview the retailer, the broker and the friend about every buying moment in the market. The machine will answer all day.

**27% vs
9%**

use AI for product
research, versus would
let it buy

IPSOs

Ask it the questions your buyers ask, many times over, because one answer is one draw. Where it prefers a rival it will often say why, and the why is a list of work to do.

How much it matters so far

The effect is strongest where the purchase is complex. The best evidence is peer-reviewed. Kaiser and Schulze analysed 973 online stores, with more than US\$20 billion in combined revenue, from August 2024 to July 2025 ([Marketing Science](#)).

Sites in complex categories drew 4.6 times the share of ChatGPT referrals that simple categories did. On those sites, visitors sent by ChatGPT converted better than five traditional channels, organic search and email among them. Across all 973 sites, AI referrals were under 0.2% of visits, and overall they converted below every traditional channel except paid social.

It is growing and improving quickly. Adobe’s retail data, reported by Decrypt, shows AI-sent visitors converting 38% worse than others in March 2025. A year later they converted 42% better ([Decrypt](#)). By May 2026 the gap was 54% ([Digital Commerce 360](#)).

Similarweb counted 770.7 million visits a month sent by AI assistants worldwide in the year to May 2026, up 117% on the year before ([Similarweb](#)). It’s still a small share of all visits, and that’s the point: most of the choosing now happens inside the answer, before anyone clicks.




FIGURE 9.1 How big is this for you · AI effect by buying mode

The shortlist moves to the machine

The list is written early, and AI has started to write it. The machine leans on its own memory. In one preprint experiment, AI models picked the famous brand on a tie, but dropped it when a rival offered better specifications.

So AI doesn't level the field for challengers. Where information is thin, the known brand wins the tie. A challenger wins by giving the machine better evidence, not more content. Tally, a ten-person maker of online forms, says AI assistants now bring it a large share of its new users.

There is no page one

When volunteers asked AI assistants the same request for brand recommendations, they almost never got the same list twice. So there's no position to hold, and a tool that reports your rank in an AI answer is mostly reporting noise.

What did hold steady was how often each brand appeared across many runs. Selection by a machine is a rate, not a rank. Chapter 14 shows how to measure it.

Named is not recommended

An AI can name five brands and recommend one. Counting how often you're named is visibility. Counting how often you're the one recommended, and why, is Selectability.

The gap is wide: in one adviser's sample of 31,200 answers, the brands it tracked were the first choice in under a third. And it changes what people do. In one 2026 preprint, a recommendation moved people to search for a brand two to three times more than a neutral mention ([Iannelli and Ai](#)).

The seven stages apply to the machine as they do to a person: does it know you, find you, believe you and prefer you?

The buyer checks the machine

The machine advises and the buyer audits, harder when more is at stake. In a 2026 TrustRadius survey of 1,862 technology buyers, 63% used AI in their research. And 94% fact-checked what it told them at least some of the time ([TrustRadius](#)).

So in a considered purchase, the machine's answer starts the checking, often on your own website. What the buyer finds there has to match what the machine said.

Two memories

Marketing works by leaving a memory. There are now two memories to leave: one in the buyer, one in the machine. The buyer's is built by experience and

advertising. The machine's is built from what has been written about you, much of it by other people.

Studies of what AI cites disagree widely on how much. My reading is that your own site supplies the facts, and others supply the verdict.

The machine can be fooled

In March 2026 China's state broadcaster showed how easily. An insider used software bought openly online to invent a smartwatch, with fake reviews and experts who didn't exist. Soon a popular chatbot was recommending it ([China Media Project](#)).

Nobody yet knows for how long that will work. Two things still favour real evidence. Buyers check the machine when the stakes are high, and the platforms now treat manufactured evidence as spam.

The machine reads what you have published, and what's said about you, in one sitting. So what should it find?

Whose evidence this is. Ipsos sells market research. BCG sells consulting. Adobe sells analytics and commerce software. TrustRadius runs a software review site. Similarweb sells web analytics. Peer-reviewed source: Kaiser and Schulze, in *Marketing Science*. Iannelli and Ai is a preprint.

TRY THIS Ask two AI assistants the question your best customer asked before buying, ten times each. Count how often you're named, and how often you're the one recommended.

IN SHORT • CHAPTER 9

WHAT'S TRUE

An AI assistant is the buyer's newest adviser. Being named by it isn't the same as being recommended.

DO THIS QUARTER

Ask two AI assistants your best customer's question before buying, ten times each.

WATCH

How often you're named, and how often you're the one recommended, across many runs.

WHERE YOU ARE

01

Why it feels like this

02

What a choice is
made of

03

The way forward

04

The builds

05

Knowing it worked

03 / PART THREE

The way forward



INTERLUDE

The open *ground*

It was 2025, and AI had started eating into search.

We'd just done some of our best work for a lender we worked with. A new content strategy, an upgraded website, tools that helped buyers work out what they needed. By every early sign we watched, it was working. People were finding the site, spending time on it and using the tools.

Then the leads started to dry up, and the deals went with them.

I couldn't work it out. Nothing we'd built had changed. Every engagement number we reported was still healthy. We were doing the hard work, and doing it well. But the numbers that mattered to the client were falling, and they pressed us hard to account for it. Fair enough, too.

When they asked, I told them the truth: "I'm honestly confused, and a little bewildered." Soon after, they halved our budget.

Shortly afterwards I found the culprit. AI answers were taking the searches our content used to win. Buyers hadn't all moved, but enough of them were getting their answers from an AI assistant to cut the client deeply.

The answer came too late to save the account. Within a few months we'd handed the work back.

I had to admit I was out of my depth. Getting a brand into AI answers was a new discipline, and I set out to learn it.

Before the foundations, an admission. I looked hard for a company that has published a measured case of winning at every stage. I didn't find one. That doesn't mean nobody does it well. It means nobody publishes it: every team owns a number and nobody owns the choice.

Why winning the whole choice is rare

The arithmetic shows why it's rare. Take a considered or committee purchase, where the buyer needs all seven stages to hold.


Suppose a brand wins each stage four times in five. It wins all seven about one time in five.

In a shortlist of five suppliers, one in five is only fair share. A brand that's good at every stage is merely average.

A single weak stage does more damage. Let Believed slip to one in two and the whole falls to about one in eight.

Good at every stage wins about one choice in five: only fair share

Of 100 buyers, those for whom every stage so far has held



An illustration for one considered or committee purchase, not a funnel.

Stages are conditions that must all hold at the moment of choice, not steps taken in order. They are not independent.

One considered purchase · seven stages that must all hold, an illustration

Treat that as an illustration, not a model. The stages aren't independent, and buyers don't pass through them in order. Not every purchase needs all seven, either. But for considered and committee purchases, the point stands. A brand that's good everywhere can still lose most of the time.

What AI changes

AI raises the stakes, because a machine now writes more of the shortlist before any supplier is called. It also makes the stages easier to see: you can ask the machine what it knows, believes and prefers. It doesn't make them easier to win. Your rivals can ask the same questions, and the machine leans toward names it already knows.

So the ground is open to whoever manages the whole choice, rather than one stage of it. The rest of this book is how.

What it takes: three foundations and seven stages

Being selected isn't a campaign. It rests on three foundations, which every stage draws on, and then on seven stages won one moment at a time.

- ▮ **Hold one truth.** A single account of who you're for, what you claim and your proof, told the same way everywhere. Chapter 10.
- ▮ **Have a reason.** A reason to choose you that a sensible rival wouldn't claim. Chapter 11.

▮ **Know your moments.** Which buyers, at which steps, in which way of buying. Chapter 12.

Then come the seven stages, one chapter each, each built on one line.

STAGE	CHAPTER	THESIS
Known	13	You're never just known. You're known for something, at a moment.
Found	14	Found is a rate, not a rank.
Answered	15	The questions you avoid are the ones that decide.
Believed	16	You can't claim belief. You can only supply the evidence for it.
Preferred	17	A reason counts only if it survives being repeated by someone else.
Bought	18	Most lost sales are lost to nothing.
Kept	19	Today's customers write the evidence tomorrow's buyers read.

Each foundation chapter ends with something to do this afternoon and this quarter. Every stage chapter ends with its core tool, printed in full; the rest are free in the Selectability Playbooks at rocmode.io/playbooks.

IN SHORT • THE OPEN GROUND

WHAT'S TRUE

Winning the whole choice is rare. A brand that's good at every stage can still be merely average.

DO THIS QUARTER

Start on the three foundations: hold one truth, have a reason, and know your moments.

WATCH

Your weakest stage, because a single weak stage drags down the whole choice.

CHAPTER 10

One truth, many *answers*

The map multiplies your messages. The machine reads them all at once. One thing survives both: a story that's true of you and of nobody else, told the same way everywhere.

The arithmetic problem

Go back to the map of buyers and journey steps. Ten kinds of buyer and twenty journey steps make 200 moments. Seven stages in each makes 1,400 answers to get right.

For most of marketing's history that number was absurd. AI has changed the sum. Drafting 1,400 answers now costs little. Checking and approving them still costs what it always did.

Nobody should start with all 1,400. Chapter 12 shows how to choose the few moments worth winning first.

Which leaves a new problem: if you can draft them, so can everyone else.

When everyone can say everything

When any brand can produce any answer, answers stop setting brands apart. The scarce thing becomes what's true only of you: your data, your proof, your method, your history, your point of view.

There's early evidence that machines reward this. A Princeton-led team tested which rewrites made a page more likely to be used in AI-generated answers, across 10,000 queries ([Aggarwal and colleagues](#)). They called the practice generative engine optimisation, or GEO.

Adding statistics, quotations and cited sources lifted a page's visibility by up to 40% on their simulated engine. On Perplexity, a live AI search engine, the gain was up to 37%. Stuffing in keywords did worse than changing nothing.

The tests ran on 2023-era systems, and measured being used in an answer, not being bought. Even so, specific and evidenced beats generic, and the most specific evidence of all is the kind only you hold.

It also has to be checkable. When the stakes are high, buyers audit the machine, and only proof they can verify survives.

Up to 40%

lift in AI-answer visibility from adding statistics, quotes and sources

AGGARWAL AND COLLEAGUES

The living answer

How does one story become 1,400 answers without becoming 1,400 stories? Every cell in your body carries the same DNA and expresses it differently. The identity never changes, and the expression always does.

What stays and what changes

STAYS THE SAME EVERYWHERE	CHANGES WITH THE MOMENT
Who you are for	Which question you answer
What you claim	Which proof you lead with
The facts	The words, the format, the length

Personalisation that changes the right-hand column is good marketing. Personalisation that changes the left gives you a different brand every time. Marketers have called this discipline integrated marketing communications since the 1990s. What’s new is the number of answers, and a reader that sees them all.

One truth, held in one place

Saying one thing in 1,400 forms can't be done from a brand deck. It needs a single structured account of who you’re for, what you claim and the proof, which every answer is drawn from.

The evidence for structure comes from inside a company’s own systems. In one benchmark, GPT-4 answered insurance-company database questions correctly about 17% of the time. Once the same data was organised as a graph of facts and relationships, it was right 54% of the time ([Sequeda and colleagues](#)).

The test was internal business data, not brand content. I read it as a reason to keep your own facts structured, not as evidence that public AI assistants read your structured data. Google says no special markup is needed for its AI answers. In one test, assistants treated a site’s hidden structured data as ordinary text ([Search Engine Journal](#)).

So the one truth is the foundation, not the answer. Its job is to make everything you publish agree, so the machine meets one consistent brand, in words it can read and others can confirm.

Fit the question, keep the facts

The right-hand column matters too. Researchers at Columbia and Yale built a simulated marketplace and let AI shopping agents choose among products

(Allouah and colleagues).

A seller's own agent then made "simple, query-conditional description tweaks": it reworded the product description to fit the question being asked. The facts stayed fixed. Across six AI models, the product's average share of the agents' choices rose by between 0.32 and 14.89 percentage points.

The gains were uneven: a single round of rewording produced a significant increase in only 33% of the experiments. Fit to the question helps, by an amount that depends on which model is choosing.

It was a simulation. But it shows both halves of this chapter at once. Change the words to fit the moment. Never change the facts.

Why the machine raises the stakes

A human buyer only ever saw their own moment. The machine reads many at once, but it's a poor referee. Given two sources that disagree, AI models mostly fail to notice the conflict, and pick one version to state with confidence (Chapter 16).

Tell one audience your software saves 20% and another 30%, and the machine may repeat either, to anyone who asks. The buyer who checks and finds the other version now doubts both.

Nothing yet shows that consistency on its own makes a machine recommend you more. The case rests on what inconsistency costs: a lapse can now be repeated at scale.

This afternoon, this quarter

This afternoon: write one page. Who you're for, the three things you claim, and the proof for each. Then hold your home page, your best-selling product page and your last proposal against it, and mark every place they disagree.

This quarter: turn that page into a structured record that your team and your tools draw every answer from. Make it somebody's job to keep it true.

The next chapter gives that truth an edge.

Whose evidence this is. Sequeda and colleagues work for data.world, which sells data-catalogue and knowledge-graph software. Search Engine Journal is a trade publication. Academic sources: Aggarwal and colleagues (presented at KDD 2024, a peer-reviewed conference); Allouah and colleagues (a preprint, run as a simulation).

TRY THIS write your one-page truth, then log each claim in the claim, proof and source inventory (Chapter 16), free at rocmode.io/playbooks.

IN SHORT • CHAPTER 10

WHAT'S TRUE

Tell one story that's true only of you, the same way everywhere. Change the words, never the facts.

DO THIS QUARTER

Turn your one-page truth into a structured record every answer draws from, and make keeping it true someone's job.

WATCH

How many places your pages, proposals and AI answers disagree about who you're for, what you claim and your proof.

CHAPTER 11

A reason to *choose* you

Being seen needs content. Being selected needs a reason.

One truth says what's so. This chapter is about the reason a buyer, or a machine, would pick you over someone just as capable. It's the most opinionated chapter in the book, and I mark where the evidence runs out.

Distinctive is not different

Distinctive means you're recognised: a name, a look or a sound people know as yours. Different means there's a reason to pick you that isn't true of the rest.

The Ehrenberg-Bass researchers, whom this book has leaned on, argue that distinctive matters more. Their evidence is real. In their surveys, most of a brand's buyers don't see it as different, and they go on buying it ([Romaniuk, Sharp and Ehrenberg](#)). In one, 77% of Apple's users didn't rate it as unique ([Ehrenberg-Bass](#)).

They don't say difference is worthless, only that it plays a smaller part than the textbooks assume. For habitual purchases I think they're right. Nobody weighs the unique selling point of a chocolate bar.

Where a reason is needed

But where a buyer, or a machine, compares options on purpose, the comparison needs a reason. An AI asked which supplier is best for a field team of forty has to say why. A brand with no reason to give leaves it nothing to say.

The machine evidence agrees. AI models pick the famous name when products are identical, then mostly drop it once a rival has better specifications (Chapter 13).

Being distinctive gets you Known and Found. Being different gets you Preferred. A brand needs both, and most work only on the first.

Why it matters more each year

Here's the part I believe and can't yet prove.

There are more businesses competing for each buyer. The US counted 30.4 million non-employer establishments in 2023, businesses with nobody on the payroll, up from 29.8 million a year earlier ([US Census Bureau](#)). Those figures predate most AI tools, so they show the direction, not the cause.

AI also makes competent work cheap, and more alike. By one search firm's estimate, about half of new articles are now written mainly by machines. People take the first option that clears the bar. When every option clears it, something else has to decide.

And the newest chooser has to give reasons. On a tie, a machine leans to the name it knows, but a recommendation still has to say why.

I think that makes a real point of view worth more every year. No study yet shows it. It's a conviction, not a finding.

What a reason is made of

Al Ries and Jack Trout called it positioning, and April Dunford has brought it up to date. A reason to choose you is usually one of four things.

- ▮ **A problem you solve that others treat as a sideline.** You're for someone in particular, and you can name them.
- ▮ **A way of working you believe in and rivals don't.** An opinion about how the job should be done.
- ▮ **Something you refuse to do.** The work you turn down says who you're for.
- ▮ **Proof only you hold.** Your data, your results, your history: what's true only of you.

I've turned down somewhere between six and a dozen projects over the years. It was rarely about the industry. It was almost always about the person across the table.

The red flags usually show up in the first meeting. Someone talks over the top of me, or my team. They don't take my advice in my own field. They care about the budget to the exclusion of the outcome. Or they promise good things down the road if we look after them now.

I can't remember ever regretting walking away. I have regretted not walking away.

One project still stands out. A professional services firm came to us with a fixed budget and a big ambition: new branding and a website, designed and built from scratch. They wanted something industry-leading, and we could have delivered it.

The flags went up during the negotiation. They pushed back on price again and again. They kept quietly adding to the scope and telling us it had already been

discussed, when it never had. Twice I nearly walked away, and I told them so. Each time they backed off, nudged the budget up a little and gave some ground. And each time I knew they'd be trouble. I signed anyway.

They were. They blew through the agreed rounds of changes every time. They contradicted their own feedback. They kept bringing in new ideas and new scope.

They weren't bad people, but they were bad clients for us. As a small agency, I pride myself on going above and beyond. The wrong client is very good at taking advantage of that, and they did.

The project took twice the time we'd budgeted, and we didn't make an extra cent. It was an expensive lesson, and a valuable one. These days, when I see one of those flags, I listen to it.

The strongest reasons are opinions. Vanguard was built on the view that picking stocks is a mistake. Fidelity was built on the opposite view, that the stock-picker is the heart of the firm. Both succeeded ([Roger Martin](#)). Each gave a certain kind of investor a reason.

A reason that travelled

The best reasons are short enough for customers to pass on. Wise, the money-transfer firm, began in 2011 with two friends in London who were tired of their banks' fees and poor exchange rates ([Wise](#)).

Its reason was an opinion about the job: move money at the real, mid-market exchange rate, and show the fee. Most banks did the opposite.

Customers repeated it for the company. In its 2023 results, Wise said around two-thirds of new customers joined "as a result of someone having told them about Wise". It served about 10 million active customers that year, up 34% ([Wise plc](#)).

These are the company's own figures, and a cheaper price did much of the work. Even so, the reason was true, checkable and easy to say in a sentence.

The test

Roger Martin has a test for this. State your claim, then state its opposite. If the opposite is plainly stupid, you haven't made a choice, and every rival is saying the same thing.

"We put customers first" fails. Nobody puts them last. "We only work with cafés and restaurants" passes, because a good rival could sensibly choose to serve every trade.

2 in 3

new Wise customers
joined because someone
told them

WISE PLC

Take a bookkeeper for cafés. She's an invented example, and she returns in every stage chapter. Her proof is one café's twelve months of numbers.

CLAIM	WHAT IT SAYS	ITS OPPOSITE
Generic	Friendly, reliable bookkeeping for small business	Unfriendly and unreliable. Nobody would say it
A reason	Bookkeeping for cafés. Your wage cost against takings, every Monday by nine	Every trade, reported monthly. A sensible rival does exactly that

The second one costs something: the plumber down the road won't call. Most brands won't pay that price, which is why most sound the same.

Then check it three ways. Is it true, with proof a buyer could check? Is it wanted, by a real buyer at a real moment? Is it yours, or could a rival say it by Friday?

Chapter 17 applies the same opposite test once the reason has left your hands, to what customers, colleagues and machines say when they repeat it.

By mode

MODE	HOW MUCH A REASON MATTERS
Habitual	Little. Be distinctive and easy to find
Quick choice	Some. One clear reason, visible at a glance
Considered personal	A lot. The buyer is comparing on purpose
Owner-led business	A lot. The owner asks whether you are for a business like theirs
Committee	Most. The reason has to survive being repeated in a room you are not in

This is my reading, like the other mode tables.

This afternoon, this quarter

This afternoon: finish the sentence "choose us when". Then write its opposite. If the opposite is absurd, start again.

This quarter: ask ten customers why they chose you, and a machine why it would recommend you, and why it wouldn't. Where those accounts and your own differ, that's the work.

Next come the moments where the truth and the reason both have to be heard.

Whose evidence this is. Graphite, the source in Chapter 1, sells search marketing services. The Ehrenberg-Bass Institute is funded by sponsor brands. Wise's figures are its own results. Peer-reviewed: Romaniuk, Sharp and Ehrenberg (*Australasian Marketing Journal*). Official statistics: the US Census Bureau.

TRY THIS fill in the "choose us when" worksheet from Chapter 17, free at rocmode.io/playbooks, and strike any claim whose opposite no rival would choose.

IN SHORT • CHAPTER 11

WHAT'S TRUE

Being seen needs content. Being selected needs a reason a sensible rival wouldn't claim.

DO THIS QUARTER

Ask ten customers why they chose you, and ask a machine why it would and wouldn't recommend you.

WATCH

Whether customers and machines repeat your reason when they explain why they'd choose you.

CHAPTER 12

Know your *moments*

You can't be selected everywhere. So decide where, and by whom, before you make a single thing.

Most marketing starts with what to say and where. This chapter starts earlier: which buyers, at which moments, in which way of buying. Every stage chapter starts from the moments you choose here.

Get the offer right first

No marketing gets a wrong offer selected. The product, its price and where it can be bought decide more choices than any message.

Marketing can't fix those, but it's often first to see they're broken. Buyers who find you, believe you and still choose someone else are telling you about the offer. A moment you can't serve well isn't worth winning.

Derive them, don't brainstorm them

Most maps of buyers are brainstormed in a workshop, and they show it: the buyers the team already talks to, the questions sales already hears, and nothing else. Getting the moments right is the most important step in this book, because everything after it inherits the map. So don't invent it. Four things you already know generate most of it.

- **Your industry gives the structure.** Research on how organisations buy has named the roles in a buying group for fifty years: the people who use the purchase, influence it, decide, buy, and can block it ([Webster and Wind, 1972](#)). Later writers added the person who starts it ([Bonoma, 1982](#)). Gartner's research describes six jobs every business buying group has to get done, in no fixed order: identify the problem, explore solutions, build requirements, select a supplier, validate the choice, and build agreement inside the business. It now puts those groups at five to sixteen people ([Gartner](#)). Robinson, Faris and Wind's buy classes say whether it's a first purchase, a changed repeat or a straight repeat. And the triggers come from Jenni Romaniuk's category entry points: the situations that bring a category to mind, found by asking why, when, where, while, with whom, with what and how the buyer is feeling ([Oxford University Press](#)).
- **What you sell gives the needs.** Every offer is hired to get a job done, and every job runs through the same steps: define, locate, prepare, confirm,

execute, monitor, modify and conclude ([Bettencourt and Ulwick, 2008](#)). Each step raises its own questions. Most maps stop at the purchase, and miss everything the buyer needs after it.

- / **Your alternatives give the switching moments.** For each alternative, doing nothing included, four forces decide whether a buyer moves: the push of today's problem, the pull of the new option, the anxiety about it, and the habit of the present ([The Re-Wired Group](#)). Each force is a question, or an objection someone in the group will raise.
- / **Your difference gives the problem insight.** A real difference implies that buyers misread their problem. The moments where they frame it are where your reason decides the choice. If no moment needs your difference, it isn't one.

For the bookkeeper, the industry adds a buyer she'd left off the map: the café's accountant, who influences the choice and can block it. The job map adds what happens after the sale: checking the Monday report each week, and changing what's tracked when the menu or the roster changes. Her alternatives (the accountant, software, doing it herself) each bring their own objection. And her difference, weekly wage figures, explains why "wages ate the takings again" is the moment she wins first.

Check every moment, and check the map is whole

A derived moment is still a guess until something backs it. Give each one a source: a sales call, a lost deal, a tender, a review, a forum thread, a question people put to an assistant. Mark each one sourced, inferred or a hypothesis. Take the hypotheses to interviews, and keep going until new conversations stop adding needs. In one well-known study that happened within twelve interviews, with the basics clear within six ([Guest, Bunce and Johnson, 2006](#)). Treat that as a guide, not a rule.

Then test the whole map with five questions:

1. Was every person in the buying group considered at every step, with any absence stated?
2. Does every product have its job mapped to the end, past the purchase?
3. Does every alternative, doing nothing included, have its switching moments?
4. Does every part of your difference decide at least one moment?
5. Does every stage, from Known to Kept, have moments?

Expect the map to grow. A single job can hold fifty to a hundred and fifty separate needs ([Strategyn](#)). A moment keeps that manageable: it's one person at one step, with the needs held inside it.

The quick version still has its place. Three kinds of buyer and five journey steps, filled from your customers, your sales team and an assistant, will show you where you're losing this afternoon. The full derivation, worked through for one firm, is in the playbooks.

Pick the moments you can win

Not every moment deserves the effort. Three tests sort them.

- ✓ **Is the buyer worth winning?** Some moments belong to customers you can't serve at a profit.
- ✓ **Do you have a true reason to be selected there?** A claim only you can make. Where you have none, being seen is wasted.
- ✓ **Is the moment open?** Ask the machine the buyer's question, and see who it recommends today, and why.

To test properly, use Chapter 13's moment memory map and Chapter 14's presence sample.

The bookkeeper's shortlist

Back to the bookkeeper for cafés. She filled her fifteen squares, then ran the three tests on the six that looked most promising.

MOMENT, IN THE BUYER'S WORDS	WORTH WINNING?	TRUE REASON THERE?	OPEN?	VERDICT
"Wages ate the takings again this month"	Yes	Yes: the Monday wage report	Yes: assistants name general bookkeepers	Win it first
"My accountant only sends numbers at year end"	Yes	Yes: weekly figures	Partly: two online firms appear	Win it
"Our accountant says the books are a mess before tax time"	Yes	Partly: tidy books, but not weekly ones	Yes	Win it, with the accountants who refer her
"We're opening a second café"	Yes	Yes	Unclear: few answers name anyone	Test it next quarter
"I've just bought a café and need to set up the books"	Small jobs, little margin	No	Yes	Leave it
"Good bookkeeper for a small business near me"	Some	No: every bookkeeper says the same	No: dozens compete	Stop spending here

The tests cut six moments to three. All three start from a café owner’s problem, in her terms. Her old spending ran the other way: most of it went on the generic search, where she had no reason to be chosen.

Small and specific beats broad

Shopify, the online-store platform, told investors in August 2026 that AI-driven orders to its stores had tripled in a year. Three-quarters of AI-attributed purchases came from outside its hundred biggest product categories (Retail TouchPoints).

That’s a long tail of categories, not proof that small sellers win. My reading is that a buyer who describes exactly what they need is matched to whoever describes exactly what they sell. Machines reward specifics, and a precise offer gives a small seller a way into moments a famous brand never thought to enter.

Know how they buy

The same moment is won differently in each of the five buying modes.

MODE	WHERE TO CONCENTRATE
Habitual	A few moments, reached again and again
Quick choice	The moment of need itself
Considered personal	The questions asked while comparing
Owner-led business	The owner’s worries: risk, cost and time
Committee	A separate map for each person in the group

This afternoon, this quarter

This afternoon: write down three kinds of buyer and five journey steps, then run the five completeness questions over them. For each moment, ask one AI assistant the buyer’s question a few times, and note who it recommends.

This quarter: choose the five moments you most want. Write down your reason to be selected in each, and stop spending on the moments where you have none. To defend that cut with numbers, score those moments with the Selectability Score (Part five).

A truth, a reason and a set of moments. Now the stages, one at a time.

Whose evidence this is. The Shopify figures are the company’s own, from its earnings call, reported by a trade publication. Gartner sells research and advice to the sellers it studies. Strategyn sells outcome-driven innovation consulting, and The Re-Wired Group sells jobs-to-be-done training. Webster and Wind, Robinson, Faris and Wind, and Guest, Bunce and Johnson are peer-reviewed or academic sources.

3 in 4

AI-driven Shopify purchases came from outside its top 100 categories

RETAIL TOUCHPOINTS

TRY THIS run the three tests on your fifteen moments, then score the five you keep with the self-score worksheet, free at rocmode.io/playbooks.

IN SHORT • CHAPTER 12

WHAT'S TRUE

You can't be selected everywhere. Decide which buyers and moments to win before you make anything.

DO THIS QUARTER

Choose the five moments you most want, write your reason in each, and stop spending where you have none.

WATCH


Who AI assistants recommend when you ask your buyers' questions in the moments you chose.

CHAPTER 13

Be known

SEEN

SELECTED



“I've heard of you.”

DEFINITION

Known means your name comes to mind, for buyers and machines, when a buyer is in a moment you chose.

THE BUYER ASKS

“When this happens to me, who comes to mind?”

MEASURED BY

Machine memory: AI recognises you by name and ties you to the moment. Buyer memory: your share of search.

FAILURE LOOKS LIKE

Healthy awareness, flat sales. The machine describes you when named, but never names you for a need.

You're never just known. You're known for something, at a moment.

My favourite from-zero story is my second agency, Cru Digital. It's probably the one people still remember.

It started with two of us: my close friend Damien Bredberg, a photographer, and me, a web designer. It was July 2004, and I was working full-time at Ballistix, a sales process agency. Great company; I was sad to leave. But a bunch of Damien's photography clients wanted websites, and I was itching to use everything I'd just learned about marketing and sales process. So off we went.

Damien's clients gave us a few months of runway at most. We had to bootstrap hard.

Ballistix had taught me the power of a good client newsletter, so that's where we started. We paid a freelance designer to create a set of postcards, and they really were cool. On the back was an offer to sign up for our newsletter. We dropped tens of thousands of them into local business PO boxes. It cost something ridiculous, maybe twenty cents a card, delivered.

About a thousand people signed up, nearly all of them business owners. Then I went to work. I wrote the most compelling, best-designed newsletters I could, packed with insights, marketing ideas, competitions, anything I thought an owner would actually want to read.

For the first few months the work only trickled in. It was barely enough to live on, even with no running costs at all. But I kept the newsletter going, every month, no exceptions.

About three months in, the trickle became a stream. We were busy. We hired two designers and a web designer to keep up.

Within a year we'd started on Google AdWords and moved out of the home office into a fairly big space in West End, Brisbane. We had about fifteen people by then, and twenty-five within eighteen months. We ran our own local events and turned up to community and industry ones. We ended up on just about every ad agency's supplier list for websites.

For a few years Cru was the go-to digital agency in Brisbane, known for the most creative, most immersive websites around. We won more than a hundred international awards over about five years, before we merged with Junior to form JuniorCru.

What it taught me is that you can manufacture awareness, if you're committed enough. And you can go from nothing to something substantial very quickly.

One caveat. I was in my early to mid twenties, and I routinely pulled all-nighters and worked eighty-hour weeks. Once I became a father, I went looking for a quieter agency life, and I've kept Evolut deliberately small to keep some balance. I haven't always managed it.

Known is the first gate. A buyer can't choose a brand they can't bring to mind, and neither can a machine.

Most brands measure it with one number: awareness. That number is useful, but it says nothing about which buying situations bring your name to mind, and which leave it out.

So this chapter asks a narrower question, twice, of the buyer and of the machine. When this buyer is in this moment, does your name come up?

Memory is filed under moments

Buyers don't search their memory by brand. They search it by situation.

Professor Jenni Romaniuk of the Ehrenberg-Bass Institute calls them category entry points. They're the cues buyers use to reach into memory when a need arises ([Ehrenberg-Bass](#)). This book calls each one a moment. Hunger at four in the afternoon is one. A letter from the tax office is another.

Moments matter after the sale too. Romaniuk surveyed US business-insurance decision-makers for LinkedIn's B2B Institute ([Romaniuk](#)). The more buying situations customers linked to their insurer, the less likely they said they were to leave, in 16 of 17 products.

The report's reading: "for each 2 additional CEP... probability of defection lowers by 5%". That's stated intention, not actual switching.

A single awareness number misses this. Ehrenberg-Bass researchers compared prompted awareness with the share of buyers who linked a brand to buying situations. They found "high stability and little congruence" between the two, over time and at brand level ([Ludwichowska and colleagues](#)). The awareness number isn't wrong. It hides where you're remembered, and where you aren't.

Fame breaks ties

Recognition still counts. In a classic experiment, "brand awareness was a dominant choice heuristic" ([Hoyer and Brown](#)). People picked the brand they knew, even when it was the lower-quality one.

Machines do the same. Xi Chu and Yupeng Hou had three AI models choose between one real skincare brand and nine invented ones, all with identical specifications. The models picked the real brand in all 670 valid trials ([Chu and Hou](#)).

Then a rival got a slightly better rating. A lead of 0.075 stars wiped out the famous brand's advantage. Once quality information was present, brand explained just 1.2% of how products were ranked.

So in this set-up, fame was a tie-breaker. It decided when nothing else did, and gave way to a small, real difference. The rivals were invented, though, and in the wild fame may well do more than break ties.

Either way, a tie-breaker needs a tie. To be in the tie at all, the machine first has to retrieve you for the situation the buyer described.

The machine knows you by what's written and searched about you

The best evidence on how machines retrieve brands comes from Edward Malthouse, a Northwestern marketing professor, and his colleagues ([Malthouse and colleagues](#)). They asked six AI models for 2,400 lists of recommendations in five categories. In the 1,200 lists from plain category prompts, many established brands never appeared.

L.L.Bean and REI were missing for hiking jackets, Craftsman and Black+Decker for drills, Braun and Philips for coffee makers.

An unwritten moment needs coverage, reviews and talk where the machine reads: PR, review and community work, with a named owner, in-house or agency. An unfelt moment needs reach and repetition among buyers.

A few moments, not one

The usual advice is to pick the one thing you want to be known for. Ehrenberg-Bass argues for breadth. In Romaniuk's words, "Brands with wider, fresher CEP networks have higher acquisition... also have higher retention" ([Romaniuk](#)). The insurance result points the same way.

So my working rule is a few chosen moments. Start with the one that matters most, win it in both memories, then widen.

One caution. Bigger brands get linked to more situations simply because more people buy them. The useful signal is whether you're linked more often than your size predicts. Romaniuk's work on Dirichlet norms sets out how to judge that ([Romaniuk](#)).

What ties a name to a moment

Snickers tied itself to a moment, hunger, with "You're not you when you're hungry". Its awards paper reports sales up 15.9% in a year, share growth in 56 of 58 markets and global sales up US\$376.3 million ([IPA](#)).

The meerkat shows the other half: a distinctive way of being remembered. In 2009 a British comparison site, hard to tell from its rivals, launched a Russian meerkat called Aleksandr Orlov ([AdNews](#)). More than fifteen years on, the campaign still runs in Britain and Australia. It didn't describe the product better. It made the name hard to forget.

A brand needs both halves: a moment to be recalled in, and a cue that makes sure it's noticed. And once a brand is the name people ask for, it stops waiting to be recommended, as Booking's direct bookings show (Chapter 3).

How to tell it is failing

- Ask customers what was happening when they first thought of you, and the answers scatter. None match the moments you market to.
- AI assistants describe you accurately when named, but never name you when a buyer describes a need.
- Your marketing promotes a different moment every quarter.

The method: the moment memory map

The moment memory map checks each chosen moment in both memories: buyers' and the machine's.

1. **Choose the moments.** Take three to five from your Chapter 12 map, written in the buyer's own words. Romaniuk suggests prompting with the W's: why, when, where, while, with or for whom, with what, and how they feel.
2. **Rank them by the 3C's.** Romaniuk's report also offers three tests. Credibility: can you claim this moment believably? Competition: who holds it now? Commonness: how often does it happen?
3. **Ask customers what was happening.** Interview recent customers, and code the answers with the guide below. Count how many name each moment unprompted.
4. **Read your branded search.** Look for your name searched alongside a moment's words. This is my working rule for a cheap proxy, not a tested measure.
5. **Ask the machine three ways.** Malthouse and colleagues used three tiers of prompt: category only, needs-based, and positioning. The gap between them separates "the machine does not know us" from "it knows us but is not cued".
6. **Sample properly.** Fix the list of competitors before you start. Run each prompt many times and record the share of runs that name you. The researchers used 40 runs per model; an owner can start with five.
7. **Ask the machine about you by name.** "What is this brand known for? When would someone choose it?" Compare the moments it gives with the ones you chose.
8. **Mark each moment's state, and assign the work.** Held, unwritten, unfelt or open. Repeat it each quarter with the same prompts.

Sampling matters. A remodeler once showed Marcus Sheridan, the marketing author, a ChatGPT answer naming it "one of the best remodelers in our city". The same question on his account didn't ([21 Hats](#)). One answer is one draw.

Rate each memory on a three-point scale.

RATING BUYER MEMORY		MACHINE MEMORY
Strong	Most customers interviewed name this moment unprompted	Named in at least half of needs-based runs
Weak	A few customers name it	Named in some runs, under half
Absent	Nobody names it	Never named

This scale is my own working rule, not tested against sales. Strong in both is held. Strong for buyers and weak or absent for the machine is unwritten, and the

reverse is unfelt. Weak or absent in both is open.

Ask what was happening

The interview is the cheapest part of the map, and the part most teams skip. "Why did you choose us?" gets you a reason. "What was happening when you first thought of us?" gets you the moment. "Our accountant said the books were a mess" is a moment. "You seemed good" isn't, so ask what happened just before. Ten interviews is enough to start. Code each answer on five lines:

- ✓ **Cue:** what brought the need to mind, such as a letter, a date or a bad month.
- ✓ **Trigger:** what turned the need into action that week.
- ✓ **Person:** who was involved, or who they asked.
- ✓ **Words:** their phrase for the problem, written down exactly.
- ✓ **Moment:** cue, person and words combined into one line. Count matching lines across interviews.

The bookkeeper's moment memory map

The bookkeeper for cafés is an invented example: "Bookkeeping for cafés. Your wage cost against takings, every Monday by nine." She interviewed ten clients and ran needs-based prompts for four moments.

MOMENT, IN THE BUYER'S WORDS	BUYER MEMORY	MACHINE MEMORY	STATE	NEXT STEP
"Wages ate the takings again this month"	Strong	Weak	Unwritten	Ask clients to describe the Monday report in reviews
"Our accountant says the books are a mess before tax time"	Strong	Absent	Unwritten	Write with the two accountants who refer her
"We're opening a second café"	Weak	Absent	Open	Leave it for now
"I've just bought a café and need to set up the books"	Absent	Weak	Unfelt	Decide whether she wants this moment at all

Her strongest moments live in people's heads, but almost none of it is written where a machine would read it. The machine knows her for a set-up guide she once posted, which isn't her chosen moment. Her work this quarter is to get others writing about wage cost, not to publish more of her own.

Fieldline's moment memory map

Fieldline is an invented software firm of 200 people. It sells scheduling software to field-service teams: plumbing, heating and cooling, and electrical

contractors. A committee buys it: operations, finance and IT. Its signature claim is cutting technicians' travel time, based on a 2023 pilot.

From here to Chapter 19, one invented deal runs through every stage. The buyer is Northgate Heating & Cooling, a contractor with 80 technicians. Dana, head of operations, wants the change and champions it. Ravi leads finance. Mei is the IT manager. The owner, who is also chief executive, signs.

Dana remembers Fieldline from a trade talk on travel time. Ravi and Mei have never heard of it.

PERSON	MOMENT, IN THEIR WORDS	BUYER MEMORY	MACHINE MEMORY	STATE
Dana, operations	"My technicians spend half the day driving"	Strong	Strong	Held
Ravi, finance	"Job costs never match what we invoice"	Absent	Absent	Open
Mei, IT	"Another tool to connect and keep secure"	Absent	Weak	Unfelt
Owner-CEO	"Will this pay for itself?"	Absent	Absent	Open

Fieldline is known for one moment, by one person. That gets it onto Dana's list, but each person in a committee arrives with a moment of their own. Its next quarter's work is a job-costing story, told by a customer's finance lead, where finance leads read.

By mode

What the memory is tied to changes with the mode. Habitual buying hangs on an everyday cue, so build broad reach repeating one cue. A quick choice hangs on the moment of need itself, so be present where the need strikes. Considered personal buying starts with life events; owner-led buying with the owner's worry about cost, risk or time. A committee brings a different moment for each person, so earn one moment per role, in the places each role reads.

How Known connects to the other stages

The moments come from your map, and distinctiveness, from your reason, gets you remembered. Known feeds Found: the machine can't find you for a moment it hasn't filed you under. And it shares its sources with Believed, because the coverage others write builds the machine's memory too.

What not to do

- Don't trust a single answer. Measure a share of many runs.

- Don't try to buy the machine's memory.** In the best study so far, ad spend didn't predict AI prominence.
- Don't read a big count as a win.** Big brands get linked to more moments because they're big.

How you know it worked

The Selectability Score (Part five) checks whether the buyer, or the machine, already knows you. Your own numbers go further, moment by moment:

- The share of new customers who name a chosen moment when asked what was happening.
- Searches for your name alongside each moment's words.
- The share of needs-based prompt runs that name you, for each moment.

For the board: priority moments held this quarter, against last quarter. "Three of five held, up from one" is a line a board can follow.

This afternoon: write three moments in your buyer's words. Turn each into a needs-based prompt, and run it five times in two AI assistants. Then ask each what your brand is known for. It's a sketch, not a measurement.

This quarter: interview ten recent customers about what was happening. Build the full moment memory map, running the three prompt tiers properly. Pick the one unwritten moment that matters most, and get other people to write it down.

TOOLKIT

1. Building the moment memory map. The process, start to finish. Run it once to set a baseline, then each quarter with the same prompts and questions.

The moment memory map: check both memories, one moment at a time




FIGURE 13.1

Read it left to right, then down. Each moment is checked in both memories, and any gap loops back for a retest next quarter.

2. Moment memory map template. One row per chosen moment. Rate each memory strong, weak or absent, using the scale in the method.

MOMENT, IN THE BUYER'S WORDS	WHO IS IN IT	CUSTOMERS NAMING IT (OF THOSE INTERVIEWED)	BRANDED SEARCH WITH MOMENT WORDS	NEEDS-BASED RUNS NAMING YOU (SHARE)	BUYER MEMORY	MACHINE MEMORY	NEXT STATE STEP	OWNER
------------------------------------	--------------------	--	--	---	-----------------	-------------------	--------------------	-------

Every template for this stage is in the free Selectability Playbooks at rocmode.io/playbooks.

What we don't know

- No peer-reviewed study I found shows that moment-based memory measures predict actual sales better than prompted awareness. The strongest link uses stated intention, from one B2B insurance survey.
- No published study tracks a brand causally from more coverage to more AI recommendation. The AI evidence is correlational, or uses invented brands, and both AI studies are preprints without independent replication.
- The one small-firm case here that was built from zero is my own, Cru Digital, and it's self-reported and from before AI assistants. VOCAL has a plan but no results yet.
- We don't know whether the two memories reinforce each other.

Whose evidence this is. Ahrefs sells SEO tools and an AI-visibility tracker. LinkedIn, which commissioned the insurance study, sells B2B advertising, and Ehrenberg-Bass is funded by corporate sponsors. The Snickers figures come from the agencies' own awards paper. Peer-reviewed sources: Hoyer and Brown; Romaniuk on Dirichlet norms. Chu and Hou, and Malthouse and colleagues, are preprints not yet peer reviewed, and Ludwichowska and colleagues is a conference paper.

IN SHORT · CHAPTER 13

WHAT'S TRUE

You're never just known. You're known for something, at a moment, in the buyer's memory and the machine's.

DO THIS QUARTER

Interview ten recent customers, build the moment memory map, and get others to write down your most important unwritten moment.

WATCH


The share of needs-based prompt runs naming you per moment, branded search, and customers naming your moment unprompted.

CHAPTER 14

Be found

SEEN

SELECTED



“You show up when I look.”

DEFINITION

Found means you appear when a buyer describes a need without naming you, on search, assistants or marketplaces.

THE BUYER ASKS

“Who can help me with this?”

MEASURED BY

A presence sample: the share of unbranded answers that name you, per moment, per assistant, with a margin.

FAILURE LOOKS LIKE

A green SEO report, flat enquiries. Found on learning questions, missing where buyers choose.

Found is a rate, not a rank.

Found means you appear when a buyer asks without naming you. They describe a need, and a search page, an assistant or a marketplace puts names in front of them. Either yours is among them or it isn't.

Known was about memory. Found is about the moment. A buyer can know your name and still never be shown it when the need arrives.

Most brands get Found wrong in two ways. They measure it as a position, which no longer holds still. Worse, they're found in the wrong moments, where nobody is buying, so the numbers rise while selection doesn't.

There is no position to hold

An AI answer has no fixed order, so a rank inside it means very little. The same question, asked twice, gets a different list.

<1 in 100

Chance two AI answers give the same list of brands

SPARKTORO

Rand Fishkin of SparkToro tested this with Patrick O'Donnell of Gumshoe ([SparkToro](#)). Six hundred volunteers ran twelve requests for brand recommendations 2,961 times. They used ChatGPT, Claude and Google's AI Overviews and AI Mode, in November and December 2025.

For ChatGPT and Google's AI, the chance of the same list of brands in two responses was under one in a hundred. The chance of the same list in the same order was about one in a thousand. Lists ran from two or three names to more than ten.

The authors are open about the limits. The study isn't peer reviewed, and they couldn't control chat history, device or the model's randomness. O'Donnell joined Gumshoe during the work, and Fishkin flags the possible conflict himself.

Google's own AI surfaces move constantly. Ahrefs watched more than 43,000 keywords for a month ([Ahrefs](#)). An AI Overview had a 70% chance of changing between observations, and changed every 2.15 days on average.

The meaning held steady, though. Ahrefs scored the Overviews' meaning at 0.95 for consistency between updates. Yet each update replaced 45.5% of the citations, and only 54% of the named brands and other entities survived from one version to the next.

Profound found that 40% to 60% of cited domains changed within a month for identical prompts ([Profound](#)). That counts source pages, not brand mentions, but the pattern is the same.

Part of the reason is that one question is no longer one search. Google says AI Mode works by "breaking down your question into subtopics and issuing a multitude of queries simultaneously on your behalf" ([Google](#)). A buyer types one question, and many hidden searches decide the answer.

So "rank for a keyword" has no fixed meaning on these surfaces. Fishkin's own verdict is blunt: any tool claiming "a 'ranking position in AI' is full of baloney".

Leaders appear almost every time

There's no stable position in any single answer. Across many answers, though, how often a brand appears holds steady, and the leaders appear almost every time.

In the SparkToro study, City of Hope appeared in 97% of 71 ChatGPT answers to an unbranded request for the best cancer hospital. Bose, Sony, Sennheiser and Apple appeared in 55% to 77% of 994 headphone answers.

That held even though the volunteers wrote their own prompts. Their wording barely overlapped, with a mean similarity of just 0.081. Very different questions still surfaced the same leaders.

Small firms can lead too. Smartsites, a digital marketing agency, appeared in 85 of 95 answers to a request for e-commerce consulting. Its moment is narrow, and it owns it. The Northwestern study of AI recommendations found the same shape at the top: five drill brands appeared in every run.

So the stable number is a rate. Leaders are found almost every time; everyone else is found some of the time, and only a rate shows the difference. A rank of “third” tells you nothing when third place moves on every asking.

A rate has limits of its own. It’s a rate over the questions you chose, not over every buyer. Choose the questions well and it tells you where you stand in those moments.

One vendor claims more stability than this. BrightEdge reports that 96.8% of cited domains didn’t change week to week (BrightEdge). It doesn’t publish its method. My reading is that weekly totals smooth out the noise in single answers, which is the point of a rate.

Most brands are found in the wrong moments

Most brands that are found are found where nobody is buying. They rank for questions people ask while learning, not while choosing. Traffic rises, and selection doesn’t move.

The Northwestern study shows the clearest form of it. Famous brands such as Craftsman and L.L.Bean were never named on plain category prompts. They were named most of the time when the prompt used their own positioning. They went missing whenever buyers didn’t describe the need in the brand’s terms.

Buyers now ask long questions. Prompts to assistants run far longer than search keywords, and most match no keyword at all (Chapter 6). A long prompt carries the buyer’s situation: who they are, what went wrong, what they need. That’s where positioning words live.

Learning questions lose their clicks. When an AI summary appears on an informational search, people click through far less, though brands cited inside it fare better (Chapter 21).

HubSpot said it plainly. On its 2025 second-quarter earnings call, chief executive Yamini Rangan said: “Just 10% of our leads come from blog traffic” (MarketBeat transcript). She added that “clicks from LLMs convert much better than organic search because buyers are deep in the research”.

That doesn’t prove blog readers never buy. It does show that a firm’s biggest source of traffic and its source of leads had come apart. Outside estimates put a large fall in HubSpot’s blog traffic around then. I couldn’t verify them, so I leave the number out.

Put together, the diagnosis is simple. Traffic measured where people learn tells you little about presence where people choose. You can win the first and lose the second.

Measure it like a poll, not a league table

Found is measured over many askings, in the moments you chose, the way a pollster samples. A ranking reports one result as if it were the truth. A poll asks a defined group many times and reports a share, with a margin.

The pollster's hardest choice is who to ask. Poll people who can't vote, and the answer is precise and useless, which is what you get when you measure Found on learning questions.

So the moment comes first. The moments you chose to win are your sample frame, and nothing outside them counts.

Branded questions never count. A question that names you tests Known, not Found. "Is Fieldline any good?" will always name Fieldline. It tells you nothing about whether a buyer who had never heard of you would meet you.

Each assistant is its own electorate. In one study of 250 queries, three AI models agreed on the top brand 41.6% of the time ([Zatuchin](#)). Ahrefs found that AI Overviews and AI Mode shared only 13.7% of their citations for the same question. They agreed on what to say about 9 times in 10 ([Ahrefs](#)). So you take a rate for each assistant, never one rate for "AI".

Classic search still runs on rank. On the page of blue links, position one gets 27.6% of clicks and position ten gets 2.44% ([Backlinko](#)), so rank still matters there. The claim here is narrower: on AI surfaces, rank has no stable meaning.

How to tell it is failing

- ✓ Your SEO report is green, and enquiries are flat.
- ✓ An agency or tool reports your "AI rank" or "position", with no number of runs behind it.
- ✓ The questions you track contain your own brand name.
- ✓ You appear when a buyer uses your words, and vanish when they describe their problem in theirs.

The method: the presence sample

A presence sample is a small, repeated poll of the places buyers ask. It tells you how often you're named, and recommended, in the moments you chose. It takes an afternoon to set up and an hour or two each month to run.

1. **Choose the moments.** Take three to five moments from your map in Chapter 12. Tag each one: learning, comparing or buying. Most of your sample should sit in comparing and buying.
2. **Write the questions as buyers would.** For each moment, write five or more unbranded phrasings, long and specific, with the buyer's situation in them. Real buyers barely overlap, so vary the wording.
3. **Strip every brand.** No question may name you. Keep questions that name a rival in a separate list, tagged as comparisons.
4. **Pick the surfaces.** Use the two or three assistants your buyers use, run separately. Add classic search for the short form of each question, and any marketplace, review site or map your buyers search directly. For "near me" questions, set the location to where your buyers are, and record it.
5. **Run each question several times.** Use a fresh session each time, logged out where you can, so history doesn't tilt the answer. Spread runs across days, because AI Overviews change every couple of days.
6. **Record two things per answer.** Were you named at all? Were you recommended? Count "recommended" only when the answer picks you, or puts you first with a reason, rather than listing you among others. Note any rival named in your place.
7. **Report shares, with a margin.** For each moment and assistant, divide the answers naming you by the answers collected. Do the same for recommended. Never report a position.
8. **Repeat on a schedule.** Run the same sample monthly, and weekly for any moment you're actively working on. Change the questions only between periods, never within one.

How many runs make a stable rate. No published guidance says. Northwestern used 40 runs per model, SparkToro 60 to 100 runs per prompt, and Ranqo 50 prompts run 10 times each ([Ranqo](#)).

Plain statistics gives a rough guide. With n answers, a 95% margin is at most about one divided by the square root of n . That's standard maths, not a finding from any study.

ANSWERS PER MOMENT, PER ASSISTANT	ROUGH MARGIN	WHAT IT CAN TELL YOU
20 (e.g. 5 phrasings × 4 runs)	±22 points	Absent, sometimes there, or almost always there
100	±10 points	Whether a quarter's work moved the rate
400	±5 points	Small shifts against a close rival

My own working rule, from standard binomial maths; not tested against sales.

The maths assumes each answer is an independent draw, and repeat runs of one phrasing aren't. They lean the same way, so the true margin is wider than the table shows. Add phrasings before you add reruns: twenty phrasings run once beat five phrasings run four times.

Twenty answers sounds thin, but it's enough for the question most brands need answered first. The Northwestern gap ran from 0% to more than 80%. Any sample sees a gap that size.

Reading the sample fairly

A presence sample is noisy by design. A few habits keep it useful.

- ✓ **Compare like with like.** Same questions, same assistants, same number of runs. A rate that rose because you added easier questions hasn't risen.
- ✓ **Read the margin before the headline.** A move from 30% to 40% on 20 answers is within noise. The same move on 400 answers is real.
- ✓ **Watch the rivals too.** A share of answers is a share of a contest.
- ✓ **Keep named and recommended apart.** Being named is Found. Being recommended is Preferred. Both come from the same answers, so record both now.
- ✓ **Add share of search as an early signal.** Les Binet, of the agency adam&eveDDB, found that share of search tracked market share in cars, energy and phones. It could lead market share by "up to a year for cars" (IPA). He calls it "not a perfect predictor".

The bookkeeper's presence sample

The bookkeeper for cafés is an invented example. She sells "Bookkeeping for cafés. Your wage cost against takings, every Monday by nine." She ran a small sample across two assistants. The figures are invented to show the method.

MOMENT	EXAMPLE UNBRANDED PHRASING	STAGE	NAMED (CHATGPT, 20 ANSWERS)	NAMED (GOOGLE AI MODE, 20 ANSWERS)
Owner wants a bookkeeper	"Good bookkeeper for a small business near me"	Buying	0%	0%
Owner learning about wage cost	"What percentage of takings should café wages be?"	Learning	0%	5%
Owner losing margin to wages	"I own two cafés and my wages are eating my margin. Who can show me wage cost against sales every week?"	Comparing	45%	30%

MOMENT	EXAMPLE UNBRANDED PHRASING	STAGE	NAMED (CHATGPT, 20 ANSWERS)	NAMED (GOOGLE AI MODE, 20 ANSWERS)
Owner switching from a general accountant	"My accountant only sends numbers at year end. I run a café and need weekly figures. What are my options?"	Buying	35%	20%

The generic buying question gave her nothing. Dozens of firms compete for it, and she says nothing that sets her apart there. The learning question gave her nothing too, and it would bring her little if it did.

Where the café owner described the problem in her terms, she appeared often. Her positioning was doing the work, much as it did for Craftsman. Her job now is to be the clearest public answer in those two moments, and to ignore the generic one.

Fieldline’s presence sample

Fieldline’s deal with Northgate, the invented contractor, moves on. Dana has put Fieldline on her list. Before the first meeting, Ravi, the finance lead, does his own homework.

He asks two AI assistants which field-service software makes job costs match invoices. He asks it several ways, over a couple of days. Fieldline is absent from his runs. Two rivals appear in most of them.

Ravi downloads one rival’s comparison sheet and brings it to the first meeting. Fieldline now starts behind with the person who controls the budget. Nobody at Fieldline can see why.

After the meeting, Fieldline’s marketing team ran a presence sample, one moment per committee member. The figures are invented.

COMMITTEE MEMBER	MOMENT	NAMED (60 ANSWERS ACROSS 3 ASSISTANTS)	RECOMMENDED
Dana, operations	Technicians lose too much time driving	70%	40%
Ravi, finance	Job costs that match invoices	10%	0%
Mei, IT	Fits our existing systems, secure enough	5%	0%

Fieldline was found by the person who already cared most, in its own words. It was missing for the two people who could block the sale.

The fix wasn’t more content about travel time. It was public, specific answers on job costing for finance, and on integration for IT.

By mode

Where Found happens changes with the mode. Habitual buyers are found on the shelf and in the reorder list, and Found matters most when a habit breaks. Quick choices happen on maps, marketplaces and quick questions to an assistant, so accurate listings come first. Considered personal buyers ask long questions, then search to check. An owner asks an assistant, a peer or a local search, so own one narrow moment in the owner's words. In a committee each member asks separately, so run a presence sample per member.

How Found connects to the other stages

Found needs Known and the moments: the machine names brands it has filed under a need, and your map of moments is the sample frame. It feeds Answered, which makes sure what buyers find answers their question. And it feeds Preferred, because the same sample records whether you were recommended.

What not to do

- ✦ **Don't mistake traffic for presence.** Visits from learning questions aren't visits from buyers.
- ✦ **Don't change the sample mid-period.** New questions in the middle of a period make the rate meaningless.
- ✦ **Don't abandon classic search.** Rank still decides clicks on the page of blue links.

How you know it worked

The Selectability Score (Part five) checks one thing for this stage. Do you appear when they ask without naming you? It counts the share of answers to unbranded questions that name you. Branded questions never count.

Your own numbers:

- ✦ Named rate per moment, per assistant, with its margin.
- ✦ Recommended rate on the same answers.
- ✦ The gap between learning, comparing and buying moments. Most of your rate should sit in the last two.
- ✦ Your rate beside your two closest rivals' rates.
- ✦ Share of search in your category, as an early signal.

This afternoon: pick one buying moment. Write five unbranded questions a buyer might ask in it, long and specific. Ask each one twice in two assistants, in

fresh sessions. Count how many of the 20 answers name you.

This quarter: build the full presence sample for three to five moments. Run it monthly. Choose the one moment with the biggest gap between you and a rival, and work on it. Judge the work by the rate, not by traffic.

An agency can sell this sample as a monthly managed service. Run well, with fixed questions and honest margins, it's worth more than another rank report.

TOOLKIT

1. Running a presence sample. Use this process each month, or whenever you change positioning or publish for a moment. It turns a pile of answers into a rate you can act on.

The presence sample: Found is a share of many askings, not a place in one answer




FIGURE 14.1

Read it along the top row, then back along the bottom. Choose the moments before writing any question, and loop back each period with the same questions.

2. The presence tracking sheet. Log one row per question, per assistant, per period. Read the shares by moment, never by single answer. Margin = $1 \div \sqrt{(\text{answers})}$, as a rough upper bound; widen it when most answers come from reruns of the same phrasing.

PERIOD	QUESTION MOMENT ID	ASSISTANT OR SURFACE	ANSWERS NAMING		ANSWERS RECOMMENDING		SHARE NAMED	SHARE RECOMMENDED	RIVALS NAMED MOST	MARGIN (± = 1/ √RUNS)
			RUNS	YOU	RUNS	YOU				

Every template for this stage is in the free Selectability Playbooks at rocmode.io/playbooks.

What we don't know

- ▮ No independent, peer-reviewed study yet says how many runs make a stable rate. My table is arithmetic, not evidence.
- ▮ The strongest support for “found in the wrong moments” is the Northwestern preprint from September 2026. It should be rechecked before it carries more weight.
- ▮ No study yet links a presence rate to pipeline or sales. The rate measures the moment of choice, not the sale.
- ▮ There's no verified small-business case of traffic falling while sales rose. The pattern is plausible, and the proof so far comes from large firms.

Whose evidence this is. SparkToro sells audience research, and Gumshoe, Ahrefs, Profound, BrightEdge and Semrush sell AI-visibility or SEO tools; Backlinko is owned by Semrush. Ranqo's and Žatuchin's authors work with tracking vendors. HubSpot sells a product for appearing in AI answers. Their incentives point towards measurement looking necessary. Binet's share-of-search work was published by the IPA, a trade body. None of the AI-visibility studies here is peer reviewed.

IN SHORT · CHAPTER 14

WHAT'S TRUE

Found is a rate, not a rank. There's no position to hold in an AI answer, only how often you appear.

DO THIS QUARTER

Build a presence sample for three to five moments, run it monthly, and work on the moment with the biggest gap.

WATCH


Named and recommended rates per moment, per assistant, with margins, beside your two closest rivals' rates.

CHAPTER 15

Be answered

SEEN

SELECTED



“You answered my question.”

DEFINITION

Answered is where the buyer's real question, and every follow-up, meets what you've published, before any sales call.

THE BUYER ASKS

“What does it cost, what goes wrong, and is it for me?”

MEASURED BY

The share of the buyer's questions and objections your site answers in public, moment by moment.

FAILURE LOOKS LIKE

"Contact us" pricing, calls opening with the same questions, and assistants quoting someone else's figure.

The questions you avoid are the ones that decide.

Answered is the stage where the buyer's real question meets what you've published. Being found brought them to you. Being answered decides whether they take the next step, or leave to ask someone else.

Most brands answer the questions they like: what we do, why we're good, how to get in touch. What they want is what it costs, what goes wrong, who it suits and how it compares.

Those are the questions most businesses leave for the sales call. The buyer, and now the machine, rarely waits for the call. If you don't answer, someone else will, and the buyer takes their answer with them.

The avoided list

Every business has a list of questions it would rather not answer. That list is the content plan.

The idea isn't mine. Marcus Sheridan set it out in *They Ask, You Answer* (2017), with his "Big 5": cost, problems, comparisons, reviews and "best of" lists. This chapter builds on his framework. What it adds is the chain behind each question, and the machine that now asks the chain for the buyer.

The list is much the same everywhere. It rarely runs past six items:

- ✓ **What does it cost?** Including the parts that aren't on the quote.
- ✓ **What goes wrong?** The common problems, and what happens when they do.
- ✓ **Who is it not for?** The buyers you would rather turn away.
- ✓ **How do you compare?** Against the named rival, and against doing nothing.
- ✓ **How long does it take?** And what the buyer has to do along the way.
- ✓ **What if I want to leave?** Contracts, exit costs, getting my data back.

Each is avoided for a reason that feels sensible from the inside. But the buyer must settle every one before they can say yes.

River Pools, Sheridan's own firm, is the case most often told. It sold fibreglass pools in Virginia, and the 2008 crash nearly sank it. The New York Times reported that orders fell from about six a month to barely two in 2009. The company overdrew its bank account ([NYT reprint](#), [SupplyChainBrain](#)).

Its owners cut the marketing budget from about \$250,000 to about a tenth of that. They started writing answers to the questions buyers asked, beginning with the awkward ones. The company later passed its pre-crash revenue.

The best-known article answers the question every pool company dodged: what a fibreglass pool really costs. It's still live, updated in December 2025, and gives turn-key ranges of \$85,000 to \$200,000 ([River Pools](#)).

In 2014 Sheridan tweeted that the article had generated "over \$2 million" ([Heidi Cohen](#)). That's his own figure, and later talks quote larger sums I couldn't trace. What holds up is the pattern: the pages that mattered were the ones the industry avoided.

Yale Appliance, a Boston store, did the same with candid comparisons. Its president reported visitors rising from 40,000 to 150,000, monthly leads from 800 to 2,300, and revenue up 40% since 2009 ([MarTech](#)). Those figures are self-reported too.

Why disclosure works, and where it stops

Telling buyers what you would rather not tell them can raise sales, under conditions.

21%

Lift in purchase probability from showing how a price breaks down

MOHAN, BUELL AND JOHN

Bhavya Mohan, Ryan Buell and Leslie John ran six studies, two of them in the field. Showing buyers how a price broke down raised purchase probability by 21% and wallet sales by 22% ([INFORMS; paper](#)).

The condition matters. The effect held when the firm chose to disclose, and not when the law required it. My reading: a voluntary answer signals good faith, and a compelled one signals nothing.

Martin Eisend pooled the published experiments on two-sided messages, the kind that admit a weakness ([Eisend](#)). They raised the credibility of the source, with small positive effects on brand attitude and purchase intention.

But the same meta-analysis found the curve bends. Past a point, more negative information made attitudes worse. How much you admit, where you put it, and whether the firm chose to say it all changed the result.

Shai Ein-Gar, Baba Shiv and Zakary Tormala found something sharper. A minor flaw, placed after the positives, made a product more attractive ([Ein-Gar, Shiv and Tormala](#)). The effect held only for people reading lightly. When readers scrutinised, or the flaw came first, it reversed.

So “confess everything” is the wrong rule. Ayn Crowley and Wayne Hoyer set out when a drawback needs refuting and which ones to disclose ([Crowley and Hoyer](#)). My working rule from all four papers:

- ✓ Disclose by choice, before anyone forces you.
- ✓ Keep the drawbacks a minority of what you say.
- ✓ Link each one to a benefit or a fix.
- ✓ Be specific. A vague admission reads like a dodge.
- ✓ Put them after the strengths, but don't rely on the order.

That last point needs care. The order effect is a low-involvement effect. A committee or a considered buyer reads closely, which is exactly where Ein-Gar's result reverses. For those buyers, what carries the drawback is the fix and the proof, not where it sits on the page.

Buyers also say they want the answers. TrustRadius found that 71% of software buyers would be more likely to buy if a vendor's site showed pricing. And 81% want to find pricing themselves ([TrustRadius](#)). The sample size isn't published.

Gartner reported in 2019 that sellers who helped buyers sort conflicting information closed a high-quality, low-regret deal 80% of the time ([Gartner, September 2019](#)). The release gives no method, but the point is useful: settle the doubt, rather than add to it. An open doubt is how a serious purchase ends in no decision at all.

An answer is finished only when it removes the next question

Buyers don't ask one question. They ask in chains: a question, then a follow-up, then an objection. An answer is finished only when it removes the next question.

Take "how much does a fibreglass pool cost?". Then "why such a wide range?", then "what pushes it to the top?", then the objection: "concrete looks cheaper, so why fibreglass?". Every link left open is a reason to keep looking.

AI conversations make these chains visible, and the machine now asks some of them for the buyer. Google says AI Mode "uses our query fan-out technique, breaking down your question into subtopics and issuing a multitude of queries simultaneously" ([Google](#)). Sundar Pichai said AI Mode queries are on average twice as long as ordinary searches ([Alphabet](#)).

That doesn't mean buyers sit typing follow-ups for an hour. Semrush studied 5,054 shared ChatGPT conversations from 2023, and the median was three messages ([Semrush](#)). WebFX studied 13,252 shared conversations published in 2025. Sessions averaged about 1.7 messages, but opening prompts averaged 103 words ([WebFX](#)).

My reading is that the chain is often packed into one long question, or the engine unpacks it by fanning out. Either way, you have to answer the follow-ups in advance.

Buyers are doing more of this alone. Of 646 business buyers Gartner surveyed in 2025, 67% preferred to buy without a sales rep. And 45% had used AI in a recent purchase ([Gartner, 2026](#)). The chain your sales team answered on the call is now answered before the call, or not at all.

The answer page

Each avoided link gets one page, built the same way. I call it the answer page, and the build chapters in Part four use the same anatomy:

1. **The question**, in the buyer's words, as the heading.
2. **The answer**, in two sentences: the number, the range or the yes or no, first.
3. **The detail**: what drives the answer, and the typical case.
4. **The proof**: a case, a sourced figure, a customer's words.
5. **The risks**: what can go wrong, each with its fix, and who this isn't for.
6. **The next step**: one small, clear action.

The machine repeats what's specific. Pages with sourced figures and quotations are cited more often than vague ones, and keyword stuffing does worse than nothing (Chapter 10). A price range, a failure rate and a named comparison are exactly what an answer page contains. The pages you've been avoiding are also the most quotable ones you could write.

How to tell it is failing

- ✓ Sales calls open with the same three questions every time, and none of them are answered on the site.
- ✓ Your pricing page says "contact us" and nothing else.
- ✓ You ask an assistant what you cost, and it quotes a competitor, an old listing or a guess.
- ✓ Deals stall at "we need to think about it", and nobody knows which question is still open.

The method: the question chain map

A question chain map follows each buyer question to the end of its chain. It shows which links you answer, which you avoid, and which to write first. You build one per moment from your map in Chapter 12.

1. **Pick one moment.** One kind of buyer at one step. "Café owner, after a bad quarter, wondering whether to get help" is a moment. "Small businesses" isn't.
2. **Write the first question in their words.** Take it from sales notes, support emails, call recordings or assistant prompts.
3. **Follow the chain.** For each answer, ask "and then what would they ask?". Write each follow-up beneath the one before it.
4. **Add the objections.** Mark the links where the buyer pushes back: "too expensive", "we tried that", "what if it fails?". Objections are questions with a doubt attached.
5. **Stop when nothing is left.** The chain ends when the next step is to buy, book or rule you out. If it never ends, the moment is too broad.
6. **Mark what you avoid.** Flag every link you don't answer in public, or answer vaguely. If the sales team says "we explain that on the call", it's avoided.
7. **Answer the avoided links first.** Write each as an answer page. Agree the price, risk and comparison pages with legal and any channel partners before you start, not after.
8. **Test with a machine.** Put the first question to two or three assistants and note the follow-ups they raise. Any follow-up your pages don't answer is the next page to write.

The price question, without a price list

Answering the price question doesn't mean publishing a price list. It means the buyer leaves knowing roughly what they will pay and why.

There's a real case for hiding price. Glenn and Sara Fisher Ellison found that where comparison engines made prices easy to compare, retailers obscured them to protect margins ([Ellison and Ellison](#)).

So firms that obscure price can be acting rationally, and I concede it.

The cost is the buyer who leaves to find an answer. Increasingly the machine fills the gap with someone else's figure. River Pools shows the middle way: ranges, the factors that move the cost, and typical totals. A rival learns little from a band that one sales call wouldn't tell them. That's my judgement, not a finding.

The numbers show the trade. HockeyStack looked at 80 software firms and 31 million visitors ([HockeyStack](#)). Pricing pages that showed prices converted 2.8% of visitors to a form, against 4.6% for pages that hid them.

But leads from the transparent pages turned into sales pipeline at 17.5%, against 10.3%. Multiply through and pipeline per visitor comes out about the same, near 0.5% either way. The gain is the same pipeline from fewer, better leads, and fewer wasted sales calls.

Partners and lawyers raise the other objection. Bring both in at step 7, and publish what they can live with: ranges rather than list prices, and real problems with their fixes.

The bookkeeper's question chain map

The bookkeeper for cafés is an invented example: "Bookkeeping for cafés. Your wage cost against takings, every Monday by nine." Her moment: a café owner after a bad quarter, wondering whether a bookkeeper would help. Here's the chain.

	LINK THE BUYER'S QUESTION	KIND	ANSWERED TODAY?
1	Do I need a bookkeeper, or just better software?	First question	Yes, on the home page
2	What does it cost each month?	Follow-up	Avoided: "get in touch for a quote"
3	Why is it more than my accountant charges for the year-end?	Objection	Avoided
4	What exactly do I get on Monday morning?	Follow-up	Partly: a screenshot, no explanation
5	What if my till system doesn't connect?	Objection	Avoided
6	Is this for a café my size?	Follow-up	Avoided: no "who it's not for"
7	Has it helped a café like mine?	Follow-up	Yes: the twelve-month case

LINK	THE BUYER'S QUESTION	KIND	ANSWERED TODAY?
8	How do I start, and can I stop?	Follow-up	Partly: a contact form, no terms

Four of the eight links are avoided: the price, the risk and the fit among them. She has been answering them on the phone, to people who already rang. The owners who never rang were lost at link 2.

So she writes three answer pages first. One gives her monthly fee in bands by café size, and what drives the band. One says which till systems connect and what happens if yours doesn't. One says plainly that she isn't for a café with no staff, because with no wage cost there's nothing to watch.

Fieldline's question chain map

At Northgate, the invented contractor, Dana wants Fieldline and Ravi has a rival's sheet. Now Mei, the IT manager, starts her own checks. Her chain is the one Fieldline had avoided.

LINK	MEI'S QUESTION	KIND	FIELDLINE'S SITE BEFORE
1	Does it integrate with our accounts system?	First question	A logo wall, no detail
2	What happens when the integration fails?	Objection	Avoided
3	Where is our data held?	Follow-up	Avoided: "contact sales"
4	What is the total price, with setup and training?	Follow-up	Avoided: "request a demo"

Mei put the chain to an AI assistant. It was vague on link 1, cited a rival's integration guide on links 2 and 3, and quoted an old reseller listing on link 4. Mei wrote "integration risk unclear" in her notes, and the committee slowed.

Fieldline answered the chain in public. It published an integrations page naming each accounts system, with what syncs and how often. It added a page on what happens when a sync fails: how the customer is told, what's queued, and how fast it's fixed. It stated where data is held. And it published price bands by fleet size, with setup and training shown.

A fortnight later the assistant answered each link from Fieldline's own pages. Mei's objection dropped, and she told Dana so.

By mode

The chain's length changes with the mode. A habitual chain is often a single link, so accurate product facts and price wherever they shop come first. A quick choice runs two or three links: does it do the job, what does it cost, can I have it now? Considered personal chains are long, with comparisons and "what goes

wrong". An owner asks about cost, time and risk, then "will it work for a business like mine?". A committee runs several chains at once, so write one answer page per role's avoided question, linked from a single hub.

How Answered connects to the other stages

Answered needs Found and one truth: an answer nobody reaches answers nobody, and the price on your answer page must match the price everywhere else. It feeds Believed, because an answer with proof and an honest drawback is the raw material others repeat. And it feeds Bought: the last links in every chain are "how do I start?" and "can I stop?". The question chain map is also the brief for the website and content builds in Part four.

What not to do

- ✓ **Don't confess everything.** A long list of flaws, or flaws placed first, turns the effect around.
- ✓ **Don't build a FAQ dump.** One page per real question, answered to the end of its chain, does more than a list of short answers.
- ✓ **Don't borrow someone else's numbers as your own.** If a case figure is self-reported, as River Pools' is, say so.

How you know it worked

The Selectability Score (Part five) checks one thing for this stage. Does what you publish answer what this buyer needs to know? It counts the share of the buyer's questions and objections your site answers, moment by moment.

Your own numbers add to that:

- ✓ The share of links in each chain that you now answer in public.
- ✓ How many sales calls still open with an avoided question.
- ✓ Lead quality, not just lead volume, after the pricing page changes.
- ✓ Whether the assistants cite your answer pages when asked your price, your problems and how you compare.
- ✓ How many deals end in "no decision", before and after.

This afternoon: write down your avoided list. Ask your sales or service team which questions they always answer on the call and never in writing. Pick the one that comes up most.

This quarter: build a question chain map for your three most important moments. Write an answer page for every avoided link, price and risk first. Then

put each first question to the assistants, log what they ask next, and close the gaps.

TOOLKIT

1. Building a question chain. The process behind the question chain map. Run it once per moment, with someone from sales or service in the room.

Building a question chain: answer the avoided links first




FIGURE 15.1

Read the top row left to right, then follow the loop. The chain ends only when an assistant raises no follow-up you leave open.

2. The answer page template. One page per avoided question, in the anatomy above. Fill every section, even the uncomfortable ones. Before it goes live, get

sign-off from legal and any partner whose prices or terms it touches. A month after, ask the assistants the question again and note whether they cite the page.

[The buyer's question, in their words]

The answer: [Two sentences that settle the question.

Give the number, range or yes/no first.]

The detail

[What drives the answer. The factors, the ranges, the typical case.]

The proof

[A case, a figure with its source, a quotation from a customer.]

The risks

[The real problems customers meet, each with what we do about it.

Then who this isn't for, and why.]

The next step

[One small, clear action. What happens after they take it.]

Sign-off: legal [name, date] · partners [name, date]

Citation check: [date] · assistants asked · cited? (Y/N) · what they quoted

Every template for this stage is in the free Selectability Playbooks at rocmode.io/playbooks.

What we don't know

- ▮ No study I found measures which questions settle a purchase. The thesis rests on disclosure research and on cases, and the best-known cases are self-reported.
- ▮ Nobody has tested whether answering price, risk or comparison questions on your site changes how often AI recommends you. The GEO paper tests how content is written, not which topics it covers.
- ▮ How long buying conversations with AI really are is unclear. The two datasets disagree and use only shared conversations.
- ▮ Where the drawback curve turns from helpful to harmful has no number. My rule of "a minority, with the fix" is a working rule, not a measured threshold.

Whose evidence this is. Marcus Sheridan sells the method through an agency and a book, and his agency promotes Yale Appliance as a client story. TrustRadius sells vendors paid pricing pages. Gartner sells advisory research on buyer behaviour. HockeyStack sells attribution software, Semrush sells search software and WebFX is a marketing agency. Peer-reviewed sources: Mohan, Buell and John; Eisend; Ein-Gar, Shiv and Tormala; Crowley and Hoyer; Ellison and Ellison.

IN SHORT • CHAPTER 15

WHAT'S TRUE

The questions you avoid are the ones that decide. An answer is finished only when it removes the next question.

DO THIS QUARTER

Map the question chains for your three key moments, and write an answer page for every avoided link, price first.

WATCH


The share of chain links answered in public, calls opening with avoided questions, and deals ending in no decision.

CHAPTER 16

Be *believed*

SEEN

SELECTED



“I trust that it's true.”

DEFINITION

Believed is where buyers and machines weigh what you say against what everyone else says, and go with everyone else.

THE BUYER ASKS

“Who else says so, and does it all add up?”

MEASURED BY

Whether AI answers about you match your one truth, cite outside support, and agree with each other.

FAILURE LOOKS LIKE

AI answers hedge, quote an old price, or repeat a number that differs across your own channels.

You can't claim belief. You can only supply the evidence for it.

At Cru Digital, most of our work came from our work.

Our own website did more selling than any of us. It didn't look like every other corporate site, and prospects drew their own conclusion: if we could build something that bold, we could certainly build theirs. Years later I put it to a colleague plainly. I was convinced we'd won most of our work because of the website we had.

The next most common reason was a small line at the bottom of somebody else's website: “Site by Cru Digital”. We believed in our craft completely, and every site we built carried the proof to whoever visited it.

Most of the rest came through referrals, and there were a lot of them. Clients had a good experience with us, and when someone asked who'd built their site, they told them.

Evolut is a much smaller agency, but the pattern hasn't changed. Almost all of our work still comes from referrals.

Which is exactly why an AI assistant had never heard of us. A recommendation over coffee is powerful, but it leaves nothing behind for a machine to read.

Believed is the stage where your own words count for least. A buyer, and now a machine, weighs what you say against what everyone else says. Where the two differ, they go with everyone else.

Belief comes from two things: proof the buyer can check, and voices that aren't yours. And you have to say the same thing everywhere, because the machine won't catch your contradictions. It will repeat them.

Two kinds of claim

Facts are things only you can say with authority: your price, your hours, your terms, what the product does, who you serve. Buyers want these from you, and they want them stated the same way everywhere.

Judgements are what you would like buyers to conclude: that you're reliable, fast, easy to deal with, the best choice. Nobody accepts a judgement from the party who gains from it.

The evidence splits the same way. In Nielsen's 2015 survey across 60 countries, 70% trusted brands' own websites, slightly more than online reviews.

Recommendations from people they know led, at 83%, and 88% in 2021 (Nielsen). This is stated trust, and the figures are dated.

The machines seem to split the same way. McKinsey reported in 2026 that only 1% of the sources AI models cite on consumer goods are brands' own websites (McKinsey). Yext analysed 6.8 million citations and put brand websites at 44% (Yext).

My interpretation, not a tested finding, is that they answer two different questions. Yext's questions were local and practical, such as opening hours. Those are facts, and the machine went to the brand. "Which is best" asks for a judgement, and the machine went to everyone else. Neither study publishes its full method.

The pattern is the chapter in one line. **On facts, the buyer and the machine read you. On judgements, they read others.**

Your lawyers know the first half. In the US, advertisers must have "a reasonable basis" for objective claims before they run (FTC). In the UK, marketers must "hold documentary evidence" for them (CAP Code). So legal often holds a ready-made proof file. Start there, and give the inventory one owner across every product line.

One contradiction anywhere can be repeated everywhere

It would be convenient if the machine caught your contradictions. It mostly doesn't, and that's worse.

IBM's team found that "all models struggle" to reflect a conflict between two passages, especially when it's implied rather than stated ([Hou and colleagues](#)). A Stanford team found models took wrong retrieved information over their own correct knowledge more than 60% of the time ([Wu, Wu and Zou](#)).

So the machine doesn't referee your inconsistencies. It picks one version, often whichever is most relevant to the question, and states it with confidence. If your pricing page says one thing and an old listing says another, the buyer may be told either.

Air Canada learned what that costs. Its website chatbot told a grieving customer he could claim a bereavement fare after travelling. The policy page on the same website said he couldn't.

A Canadian tribunal ruled that the chatbot was "still just a part of Air Canada's website". A customer couldn't be expected to check one part of the site against another ([Moffatt v. Air Canada](#)). The damages were about C\$650. The coverage went around the world.

People react differently. When a source sends mixed signals, they start checking the argument ([Ziegler, Diehl and Ruther](#)). So the machine repeats whichever version it finds, and the buyer who checks finds the other one. Only the brand that says one thing everywhere is safe.

Why other people count more

A source that gains from your belief is discounted, and one that gains nothing isn't.

Martin Eisend and Franziska Küster combined the published experiments comparing editorial coverage with advertising. Coverage's lead on credibility was about three times its lead on how the message itself was rated ([Eisend and Küster](#)).

There's a real limit, and it matters to large brands. Coverage won for unfamiliar products, but for established brands advertising did better. A famous brand borrows belief from familiarity, so this chapter's case is strongest for challengers.

Reviews show the same thing in behaviour. Northwestern's Spiegel Research Center studied clickstream data from a handful of retailers ([Spiegel](#)). Four findings are worth knowing:

270%

More likely to be bought
with five reviews than
none

SPIEGEL RESEARCH
CENTER

- ✓ A product with five reviews was 270% more likely to be bought than one with none.
- ✓ Displaying reviews lifted conversion by 190% for cheaper items and 380% for dearer ones.
- ✓ Reviews from verified buyers lifted purchase likelihood by 15% over anonymous ones.
- ✓ Purchase likelihood peaked at an average rating of about 4.2 to 4.7, and fell towards a perfect 5.0.

The data is from 2017, and the direction hasn't been seriously challenged since.

Small firms gain most. Michael Luca of Harvard Business School matched Yelp ratings with Washington State tax records for restaurants. "A one-star increase in Yelp rating leads to a 5-9 percent increase in revenue", and the effect was "driven by independent restaurants" ([Luca](#)). Chains saw no effect, and lost market share as Yelp spread. Other people's words did for the independents what a famous name does for a chain.

The machines lean on other people too, but on no single source. Profound analysed 680 million citations, and the largest single source was a few per cent of the whole. Reddit had 2.2% in Google's AI Overviews, and Wikipedia 7.8% in ChatGPT ([Profound](#)).

The shares also move: Reddit's share of ChatGPT's citations fell from 3.8% to 0.5% in a few weeks of August 2026 ([Semrush](#)). There's no single source to win, only a spread of independent ones to be present in. Paying to be shown doesn't count: AI shopping agents mark down "sponsored" tags.

How to tell it is failing

- ✓ AI answers name you, then hedge: "some users report", "reviews are mixed".
- ✓ AI answers quote an old price, a wrong policy or a product you no longer sell.
- ✓ Ask a machine where its view of you came from, and it names sources that never mention you, or only your own site.
- ✓ On the platform your buyers actually use, you have fewer than five recent reviews. Or you have a perfect 5.0, which looks managed.
- ✓ Your proof page shows logos with no names, numbers or dates.
- ✓ The same number differs between your website, your sales deck and your listings.
- ✓ Salespeople are asked for references only after the shortlist is set.

The method: a claim, proof and source inventory

The tool for this stage lists everything you claim, with what supports each claim. It takes an afternoon to start and a quarter to fill.

1. **List every claim.** Take them from your site, sales deck, listings and ads; most businesses find twenty to forty. Ask legal for the substantiation file they already hold.
2. **Mark each as a fact or a judgement.** If a rival would dispute it, it's a judgement.
3. **Make every fact consistent.** One wording, one owner, everywhere. Ask several AI assistants, your own chatbot included, about your price, terms and policies, and compare with your one truth from Chapter 10. Where they differ, fix the source the machine read, not the answer.
4. **Give every judgement its proof.** Attach the evidence you own: a named customer, a measured result and the date it was true.
5. **Count the confirmations.** For each judgement, count the independent sources that say the same thing. My working rule is that a judgement needs at least two, of different kinds, such as a review platform and the trade press. Score each from 0 to 3 using the table below, and take a point off for any live contradiction.
6. **Cut, or earn.** A judgement scoring 0 is a slogan. Cut it, or make it next quarter's brief for reviews and press. Or turn it into a fact you can prove: "trusted" becomes "300 cafés, 94% still with us after a year".
7. **Find the machine's sources.** Ask the buying questions for your category and read the sources the answers cite. That list is where your confirmations need to live.

SCORE	WHAT IT MEANS
0	Only you say it
1	One independent source says it
2	Two or more independent sources, of different kinds
3	As 2, and AI answers repeat it while citing those sources

The scoring is my own working rule, not tested against sales. Use it to rank the work, not to report to a board.

Admit the weakness, with the fix

By 2009 people said Domino's pizza tasted bad. After about eighteen months reworking the recipe, it ran ads showing its own focus groups calling the crust "like cardboard" and the sauce "like ketchup".

In the first quarter of 2010, sales at its established US stores rose 14.3% on a year earlier (Domino's). The company credited “a remarkably better pizza, and our honesty in how we told consumers about it”. That’s one quarter, against a weak year, as reported by the company itself.

The order matters. The product changed first, and the admission made the fix believable. In a considered purchase, a flaw admitted without a fix is just a flaw. Admit it with the fix, and with evidence that the fix works.

Reviews, done properly

Reviews are the cheapest confirmation most businesses can earn.

- ✓ **Ask when the job is done.** At delivery or completion, while the detail is fresh.
- ✓ **Get from none to five first.** That step moves purchase likelihood most.
- ✓ **Show that reviewers bought.** Verified buyers lifted purchase likelihood by 15%.
- ✓ **Never filter, and never steer.** Fashion Nova paid US\$4.2 million to settle a US Federal Trade Commission case over holding back reviews below four stars (FTC). Ask every customer the same way, happy or not.
- ✓ **Let the average be human.** Somewhere between 4.2 and 4.7 is more believable than 5.0. Answer the bad reviews in public, and fix what they name.

The bookkeeper's inventory

The bookkeeper for cafés is invented. She has four claims on her home page.

CLAIM	KIND	HER PROOF	WHO ELSE SAYS IT	SCORE	ACTION
Your wage cost against takings, every Monday by nine	Fact	The report itself, with a sample on the site	Not needed	n/a	Same wording on the site, the listing and the proposal
Café clients cut wage cost as a share of takings	Judgement	One café's twelve months of numbers	Nobody yet	0	Publish the case with the owner named, then ask two clients to say it in their own reviews
Certified adviser for her accounting software	Fact	Her certification	The software's own partner directory	n/a	Link to the directory from the site
Fast and friendly	Judgement	None	Three Google reviews from 2024	1	Cut the words. Ask every client at the end of their first quarter, until there are five recent reviews

Two of her four claims needed no outside voice, only the same wording everywhere. The other two were requests to be believed, and only one had anything behind it. Her quarter's work is one named case, five reviews and one listing, all of it evidence someone else can check.

Fieldline's inventory

At Northgate, the invented contractor, Mei's objection has dropped. Ravi, the finance lead, now checks the number the whole case rests on: travel time.

Fieldline's figure comes from one customer's 2023 pilot. The website says 20%, the sales deck 18%. Its review-site profile, written by a former product manager, says 25%, and the AI assistant repeats the 25%, citing it.

No machine flagged the contradiction. One version simply won. Ravi set the three figures side by side and asked which was true. The deal paused for three weeks while Fieldline found the pilot data.

Fieldline fixed the source, not the answer.

- ▮ **One number, everywhere.** The pilot result, with its date and how it was measured, on the website, in the deck and on every listing. The review-site profile was corrected at the source.
- ▮ **A second voice.** The pilot customer's operations manager agreed to be named, and to take reference calls.
- ▮ **Proof for each person.** In a committee, a claim that convinces operations does nothing for IT. Mei got a security report. Ravi got a return-on-investment case checked by a customer's finance lead. Dana got peers in the trade forum.

Asked again, the assistant gave the one number and cited the pilot.

By mode

Habitual buyers believe their own last experience, so belief there is won or lost at Kept. Quick choosers read the star rating, the count and the first few reviews: build five recent, verified reviews where they buy. Considered personal buyers check reviews, friends and comparison sites. Owners believe other owners and their trade, so one named peer in the same trade goes furthest. Committees need one proof for each role in the buying group.

How Believed connects to the other stages

- ▮ **It needs Answered and one truth.** A claim must be on the page before anyone can check it. Facts stay consistent only if they come from one record.
- ▮ **It feeds Preferred.** A reason without proof is a slogan. The machine can repeat a reason only when it can find someone else saying it.
- ▮ **Kept feeds it.** Every confirmation is written by a customer, a partner or a journalist who had a reason to say it.

What not to do

- Don't claim adjectives.** "Leading", "trusted" and "award-winning" are judgements with nothing behind them.
- Don't filter or fake.** In 2017 a journalist used fake reviews to take a restaurant that didn't exist to number one of 18,092 in London on Tripadvisor (CBC). Platforms and AI firms now treat faked evidence as an attack.
- Don't let your own chatbot improvise.** What it says, you've said.

How you know it worked

The Selectability Score (Part five) asks whether what's said about you is accurate, and backed by others. It checks the AI answers that name you: do the facts match your one truth, is outside support cited, and do the answers agree?

Your own numbers sit beside it:

- the number and recency of reviews on the platform your buyers use;
- the share of your judgement claims scoring 2 or more;
- the number of open contradictions, which should be zero.

This afternoon: ask an AI assistant whether your brand can be trusted, and where it got that. Then ask your price and main policy three times each, and compare with your website.

This quarter: build the inventory, fix every contradiction at its source, and take your three most important judgements to a score of 2.

TOOLKIT

The contradiction sweep. Run it once a quarter, and after any change to price, terms or policy. It finds the facts the machines get wrong about you, and fixes them where they started.

The contradiction sweep: fix the source, not the answer




FIGURE 16.1

Where an answer is wrong, the fix is never to argue with the machine. Find the page, listing or profile it read, correct it there, and ask again.

Every template for this stage is in the free Selectability Playbooks at rocmode.io/playbooks.

What we don't know

- No study yet links consistency across sources to how often AI recommends a brand, or to sales. The case here rests on the lab evidence about how models handle conflict, and on what buyers do when they check.
- The share of AI citations that come from brands' own sites ranges from 1% to 44%, depending on the study and the question. My reconciliation of the two is an interpretation.
- The trust figures are from 2015 and 2021, and the review figures from 2017. The direction is well established. The exact sizes may have moved.

Whose evidence this is. Nielsen sells audience measurement. Yext sells listings management, and McKinsey's figure comes from a third-party tool it doesn't describe. Spiegel's review data was supplied by PowerReviews, which sells review software. Profound and Semrush sell AI-visibility tools. Peer-reviewed sources: Eisend and Küster; Ziegler, Diehl and Ruther. The model-conflict studies are preprints, and Luca's study is a Harvard Business School working paper. Domino's figures are its own.

IN SHORT · CHAPTER 16

WHAT'S TRUE

You can't claim belief, only supply evidence. On facts, buyers read you. On judgements, they read others.

DO THIS QUARTER

Build the claim, proof and source inventory, fix every contradiction at its source, and take three judgements to a score of 2.

WATCH


Recent reviews on your buyers' platform, judgement claims scoring 2 or more, and open contradictions, which should be zero.

CHAPTER 17

Be preferred

SEEN

SELECTED



“You're my first choice.”

DEFINITION

Preferred means that when you appear, you're the one recommended, for a reason someone else can repeat.

THE BUYER ASKS

“Why them, and could I explain it to someone else?”

MEASURED BY

Your recommendation rate in unbranded answers, and the reason-match: how often they give your reason.

FAILURE LOOKS LIKE

Named, then a rival recommended first, or a generic reason: "well reviewed", "popular", "good value".

A reason counts only if it survives being repeated by someone else.

Evolut's first big win, in 2019, is still one of my favourites.

I can't remember exactly how we got to the table. We were brand new. I think our fresh new website did a pretty compelling job, and it ranked reasonably well, which can't have hurt.

We found ourselves pitching against two of Brisbane's larger digital agencies. I knew both well from the Cru Digital days. They were big and impressive. We were new and nimble.

So we pulled out all the stops. We met the client several times. We designed mock-ups. We thought hard about the strategy, the user experience, the structure of the site, the design and the pitch itself. Even by the standards I'd been used to at Cru, it was impressive work.

But that isn't what won her over. What won her over was us.

I'd brought together the best people from my old Cru and CXA teams: my best digital designer and my best CX strategist. Our new lead developer was hugely

experienced too. We presented well, and the experience showed.

She told me afterwards, as I remember it: "I chose you because you were clearly senior and experienced. All three of you. You think clearly, and I can see myself running a project with you." Then she added that the pitch and the concept hadn't hurt our chances either, and that they'd helped win over her other stakeholders.

The project was worth several hundred thousand dollars, and it quickly put Evolut on the map.

Preferred is the stage where you aren't in the room. When you appear, are you the one recommended, and for a reason? The person or machine giving that reason is never you.

Most brands get this wrong because they write their reason for themselves. Then a customer, a colleague or an assistant tries to pass it on, and it falls apart or turns into "they're good".

People choose what they can defend

People pick the option they can justify, and they lean on that harder when someone else will hear the justification. That finding is decades old and well tested.

Eldar Shafir, Itamar Simonson and Amos Tversky reviewed the experiments in 1993. People, they wrote, "often seek and construct reasons" to settle a hard choice. They do it to "justify their choice, to themselves and to others" ([Shafir, Simonson and Tversky](#)).

Simonson had already tested the "to others" part ([Simonson](#)). Two known effects, compromise and attraction, push people towards the option that is easiest to explain. Both grew stronger when people expected to justify their choice to someone else. It's the closest lab test I know of the room you aren't in.

The counter-evidence. Richard Nisbett and Timothy Wilson showed that people often can't report why they chose ([Nisbett and Wilson](#)). Shoppers picked the right-most of identical stockings nearly four times as often as the left-most. "No subject spontaneously mentioned position." They praised the quality instead.

I accept that. The reason a buyer gives is often not the cause of the choice: the feeling picks, and the reason defends. That makes the reason more useful, not less. It's the currency of the choice, even when it isn't the cause. So leave the hidden cause alone. Your job is to supply a true reason that is easy to borrow.

I learned what that means from a loss. In 2020 a mining research organisation asked us to redesign its website. We sent a proposal within a fortnight and followed up for two months. When the answer came, it came with an apology:

"You were actually still my preferred choice but I was overruled." The other agency had more examples from research and technology.

Our champion liked us. But "I liked them" doesn't survive a conversation with the boss. "They've done this for organisations like ours" does.

Three people who repeat your reason

Three re-tellers decide whether you're preferred: the machine, the champion and the customer. Each one has to say why, and none of them is you.

The machine has to say why. An assistant asked for a recommendation doesn't just name a brand. It gives a reason, and the reason shapes what the buyer does next.

A 2026 study tracked non-customers after an assistant mentioned a brand to them ([Iannelli and Ai](#)). Within seven days, a recommendation raised searches for the brand by 4.3 percentage points. A neutral mention raised them by 1.8 points. That's about 2.4 times the effect.

The panel size is withheld and the study is observational, so treat it as a strong hint.

Most brands sit in the weaker group. Leoprd looked at 31,200 answers about the brands it monitors ([B&T](#)). The brand was absent from 36%, "mentioned but not preferred" in 35%, and the primary recommendation in 29%. The full report isn't public.

The champion argues in a room you aren't in. In a business purchase, someone inside the buyer has to win over everyone else. CEB researchers, writing in HBR, analysed more than 700 business purchases ([CEB in HBR](#)). The contacts they called "Talkers" were four to six times less able to build consensus for a purchase.

Over 30% of those Talkers were senior executives, so a friendly, senior contact isn't enough. The CEB follow-up said suppliers must "arm them with the skills and materials needed to win over fellow decision makers" ([HBR](#)). Hand your champion a reason that travels.

The customer tells a friend. Fred Reichheld wrote that customers recommend only when they will "risk their reputations" ([Reichheld](#)). Whether his score predicts growth is contested, so I use only the mechanism. A recommendation spends the recommender's credibility, so they need a reason they can stand behind.

What gets repeated is also what gets cued. Jonah Berger and Eric Schwartz studied 335 word-of-mouth campaigns covering over 300 products ([Berger and Schwartz](#)). Interesting products got talked about at first, but not later. Products

2.4x

Search lift from an AI recommendation, against a neutral mention

IANNELLI AND AI

that everyday situations brought to mind kept being talked about. So a reason tied to a recurring moment travels further than a clever line.

One small firm shows the machine and the friend together. Tally, a ten-person company, makes online forms. Its onboarding survey found that “43% of new users discover Tally through ChatGPT, Claude, Gemini, Perplexity, and friends” ([Tally, 22 September 2026](#)).

Its reason is short: “Anyone can use Tally for free with unlimited forms and submissions” ([Tally, June 2025](#)). It credits years of answering questions in forums and communities, plus deliberate work on how AI describes it. The figures are self-reported, and whether assistants repeat the “free, unlimited” reason is untested.

To be someone's first choice, you must be wrong for someone else

A reason that suits everyone gives nobody anything to repeat. My working rule: if no sensible buyer would turn you down for it, it isn't a reason.

Roger Martin's opposite test still applies. State your claim, then its opposite. If the opposite is plainly stupid, every rival is saying the same thing.

This stage adds a second question. Who does your reason rule out, and would you say so out loud? A champion needs to be able to say “they're right for us because we are X”. That only works if someone else isn't X.

This doesn't shrink your market. You rule out a reason, not the category's buyers. You still want every buyer in the category to know you, because growth comes from reaching them. The reason only has to win the ones whose situation it fits.

The counter-view. The Ehrenberg-Bass researchers argue that being distinctive matters more than being different. In their surveys across 17 categories, only 17% of a brand's users saw it as different or unique ([Romaniuk, Sharp and Ehrenberg](#)). Of Apple's users, 77% didn't see it as unique.

Those surveys measured established users, mostly in habitual categories, where nobody had to justify the choice to anyone. This chapter's claim applies where a third party has to explain it.

A reason doesn't have to be a specification. Vanguard's is an opinion: buy the whole market at low cost, rather than pick stocks. A survey for an advocacy group, Friends of Jack Bogle, asked investors to describe Bogle. A top phrase was “investing in the entire market with low cost index funds is better than stock picking” ([Rockland Dutton survey](#)). The sample wasn't disclosed and the sponsor was sympathetic. But it shows investors repeating the founder's reason, and a stock-picker would sensibly choose the opposite.

The reason gap

Every brand has three accounts of why it would be chosen: its own, its customers' and the machine's. Where they differ is the work. I call that difference **the reason gap**.

Your account is what you would say if asked "why you?". It's the most polished, and the least repeated. **Your customers' account** is what they say to a friend: rougher, often truer, and what they'll repeat. **The machine's account** is what an assistant says when it recommends you, or says about a rival when it doesn't. It's built from what other people wrote.

Line the three up and one of four things is true.

- ✓ **All three match.** Rare. Keep checking.
- ✓ **Customers and machine agree, and you differ.** They have found your real reason. Change your claim to fit it.
- ✓ **You and customers agree, and the machine differs.** The reason exists, but not where the machine reads. Fix the sources.
- ✓ **All three differ.** You have no reason that travels yet. Start with Chapter 11.

How to tell it is failing

- ✓ An assistant names you, then recommends a rival first, or recommends nobody.
- ✓ When it does recommend you, the reason is generic: "well reviewed", "popular", "good value".
- ✓ Ten customers give ten different reasons, and none of them is yours.
- ✓ A champion asks you for "something I can send to finance", then rewrites it.
- ✓ Your reason, read aloud, could be on any rival's home page.

The method: the "why them?" test and the reason gap sheet

The "why them?" test collects the three accounts, lines them up on one sheet, closes the biggest gap and asks again. Run it once per moment.

1. **Write your own account first.** Finish "choose us when" in one sentence. Add the proof and who it rules out. Do this before you hear anyone else.
2. **Run the opposite test.** If a sensible rival wouldn't choose the opposite, rewrite before going further.
3. **Collect the customers' account.** Ask ten recent customers how they would describe you to a friend, what they would say if asked "why them?", and who they would not send to you. Write their words down exactly. Don't tidy

them. Ten is a sample for spotting a pattern, not a rate, so report it as “n of 10”.

4. **Collect the machine’s account.** Run your buyers’ unbranded questions through several assistants, several times each. Collect at least 20 answers per moment, the minimum, which still carries a margin of about 22 points either way. Note whether you were absent, named or recommended, and the reason given.
5. **Add the champion (owner-led and committee sales only).** Send a buyer contact one paragraph. Ask them to forward it to their finance lead, partner or accountant unedited. Ask what came back.
6. **Score each account against yours.** Use the scale below. Your **reason-match** is the number of customer answers, and the share of AI answers, that score 2 or 3. Keep the two apart.
7. **Close the biggest gap at its source.** If customers say something truer, change your claim. If the machine is vague, fix what it reads. If customers lack the words, give them a line and a moment to use it.
8. **Retest.** Same questions, new customers, after each change. Watch the recommendation rate and the reason-match together.

SCORE	WHAT THE ACCOUNT SAYS
0	No reason, or one every rival could claim
1	A real reason, but not yours
2	Your reason, in other words
3	Your reason, with your proof

This scale is my own working rule, not tested against sales.

What makes a reason easy to repeat

A reason travels when it’s short, specific and tied to a situation. My checks, drawn from the evidence above:

- ▮ **Can it be said in one breath?** Tally’s fits on one line.
- ▮ **Is there a moment that brings it to mind?** Cued products kept being talked about.
- ▮ **Would it defend the choice if things went wrong?** That is what Shafir and Simonson’s subjects reached for.
- ▮ **Does it carry its proof, without your jargon?** A number, a name or a result someone else can check, in words that aren’t your product names.

The bookkeeper's reason gap sheet

The bookkeeper for cafés is an invented example. Her offer: "Bookkeeping for cafés. Your wage cost against takings, every Monday by nine." The answers below are invented too.

ACCOUNT	WHAT IT SAYS	SCORE
Hers	Bookkeeping for cafés. Your wage cost against takings, every Monday by nine	—
Customers (10 asked)	Six said "I get my numbers every Monday". Three said "she only does cafés". One said "she's nice"	Six at 2, three at 1, one at 0
The machine (20 answers to "bookkeeper for a café")	Named in 4. Described as "a bookkeeper for small hospitality businesses"	1 where named
Reason-match	Answers scoring 2 or 3	6 of 10 customers; 0 of 20 AI answers

The customers have her reason, roughly. They say "Monday" but drop "wage cost against takings", the part that makes her worth choosing. So her Monday email now opens with "This week, wages were X% of takings".

The machine's gap is bigger. Nowhere it reads says "Monday" or "wages". So she publishes the twelve-month case as a page, asks every client for a review at the end of their first quarter, and reruns the same questions a month later.

Fieldline's reason gap sheet

Fieldline, the invented scheduling-software firm, is still working on Northgate Heating & Cooling. By now Mei's integration questions are answered and Ravi's travel-time contradiction is fixed. One person is left: the owner-CEO, who signs.

ACCOUNT	WHAT IT SAYS	SCORE
Fieldline's	Cuts your technicians' travel time, proven in a 2023 pilot	—
Customers (10 asked)	Seven ops leads: "the dispatchers love the map". Two finance leads: "it replaced two other tools". One: "less driving"	1 of 10 at 2
Dana, to the CEO	"It's the best scheduling tool we looked at"	0
The machine (20 answers)	"A well-reviewed scheduling tool with a strong mobile app"	0 of 20 at 2 or 3

Dana, Northgate's head of operations, is the champion. She wants Fieldline. The CEO asks her one question in the corridor: "Why Fieldline?"

She reaches for the line she has: "It's the best scheduling tool we looked at." "They all say that," he says, and the decision slides to next month's meeting.

That's the reason gap in one sentence. Dana had a reason and a mood, and she could only repeat the mood.

So Fieldline's account manager asks Dana the customer questions, turned round: "How would you explain it to him?" Dana talks about her eighty technicians. "Half of them spend a day a week in the van between jobs."

Fieldline rewrites its claim to pass the opposite test. "Choose Fieldline when travel, not paperwork, is what eats your technicians' day." A sensible rival chooses paperwork, and says so. It rules out in-house facilities teams whose staff never leave one site.

Then it hands the reason back in Dana's words. Her version for the CEO reads: "We pay eighty people to sit in vans for a day a week. Fieldline is built to cut that, and its pilot shows by how much." Fieldline also publishes the 2023 pilot where the machine can read it.

At the next meeting the CEO didn't say "they all say that". He asked, "And what does it cost us if it doesn't work?" That is a reason surviving the room. His new question belongs to the next stage.

By mode

In habitual categories almost nobody has to repeat a reason, so a look and a name people know carry more weight. For a quick choice, the machine has to lift your reason whole, in one line. In a considered personal purchase, the buyer has to repeat it to a partner or friend. An owner has to repeat it to a partner, an accountant or staff, so give them "for a business like yours", with proof from one like theirs. In a committee, the champion needs one forwardable paragraph per role. This is my reading of the evidence, like the other mode summaries.

How Preferred connects to the other stages

- ▮ **It needs a reason and one truth.** State the reason three ways and the machine picks one.
- ▮ **It needs proof.** A reason with no proof is a judgement, and nobody accepts a judgement from the party who gains from it.
- ▮ **It feeds Bought and Kept.** The landing page has to repeat the reason, and at renewal the buyer asks whether it still holds.

What not to do

- ▮ **Don't chase the top slot.** Iannelli and Ai found it was "being named at all, not where in the list" that moved behaviour. Earn the reason instead.

- ▮ **Don't pay for a reason, or rewrite your pages for one model.** Write for the question, and retest across assistants.

How you know it worked

The Selectability Score (Part five) checks Preferred with one question. When you appear, are you the one recommended, with a reason? It scores how strongly you're recommended, and whether your reason is the one given.

Your own numbers:

- ▮ **Recommendation rate.** Of at least 20 unbranded answers per moment, the share where you're the first choice.
- ▮ **Reason-match.** The share of AI answers that give your reason, scoring 2 or 3, reported apart from the customer count.
- ▮ **Customer retell.** Out of ten recent customers, how many say your reason unprompted. Report it as "n of 10", never as a percentage.
- ▮ **Forward test.** Whether your champion's paragraph comes back with a question, not a rewrite.

This afternoon: ask one assistant to recommend a supplier like you, without naming yourself, five times. Write down the reason it gives each time. Then write your own in one sentence, and see whether any of the five match.

This quarter: run the full "why them?" test on your most valuable moment, close the biggest gap at its source, and retest.

TOOLKIT

The "why them?" test. Use it once a quarter per moment, and after any change to your claim. Collect the three accounts, compare them, close the gap and retest.

The “why them?” test: close the gap where it starts, then ask again




FIGURE 17.1

Read it left to right, then follow the loop from “Retest” back to “Compare them”, because the test never finishes.

Every template for this stage is in the free Selectability Playbooks at rocmode.io/playbooks.

What we don't know

- / **No study links a retold reason to win rate.** The thesis is my inference from work on reasons, champions and word of mouth. Nobody has tested it end to end.
- / **Whether assistants repeat a brand's own reason.** I found no independent record of assistants giving Tally's or Vanguard's stated reason.
- / **How much the machine evidence holds.** Iannelli and Ai and Leoprd have undisclosed panels or unpublished reports.
- / **Whether the machine's stated reason is its real one.** People give reasons that aren't causes. The reason gap may measure what's said more than what decides.

Whose evidence this is. Iannelli and Ai both work for Scrunch AI, which sells AI visibility tools. Leoprd is an AI visibility and reputation advisory. CEB, now part of Gartner, sells sales advice. Reichheld went on to sell Net Promoter, and the Ehrenberg-Bass Institute sells research to sponsor brands. Tally's figures are its own. Peer-reviewed sources: Shafir, Simonson and Tversky; Simonson; Nisbett and Wilson; Berger and Schwartz.

IN SHORT • CHAPTER 17

WHAT'S TRUE

A reason counts only if it survives being repeated by the machine, the champion and the customer. None of them is you.

DO THIS QUARTER


Run the "why them?" test on your most valuable moment, close the biggest reason gap at its source, and retest.

WATCH

Recommendation rate and reason-match together, plus how many of ten customers say your reason unprompted.

CHAPTER 18

Be bought



“I bought.”

DEFINITION

Bought is where the buyer has picked you and now has to act: sign, pay or tell their boss.

THE BUYER ASKS

“What if I’m wrong, and can I undo it?”

MEASURED BY

Whether the landing page repeats the reason, shows proof and gives a clear next step; plus no-decision rates.

FAILURE LOOKS LIKE

Proposals go quiet after good meetings. “No decision” beats any rival in your lost column.

Most lost sales are lost to nothing.

The decisions I’ve put off longest have always been hires. I’ve been notoriously bad at hiring, in three ways that feed each other.

I wait too long to hire for senior roles, usually three to six months past the point I knew I needed someone. When I do hire, I rarely land the right person, and I suspect that’s because I want too much from one person. And because I’ve waited so long, the new hire gets rushed through their first weeks and thrown in at the deep end. Too often they’re gone within three to six months, sometimes by their choice and sometimes by mine.

The worst one still makes me wince. I needed what was, in effect, a studio manager. I found someone who’d done that exact job in-house, with a team about our size and much the same scope. Great references. On paper, a legend.

In reality, the worst hire I’ve ever made. They couldn’t juggle more than one client. They treated clients as the enemy and talked constantly about how stupid they were, and it showed. They routinely put clients offside. Before it ended, they’d cost me clients and I’d paid out refunds.

Looking back, the waiting was fear, and that hire is why. A bad senior hire costs recruitment fees, delayed projects, refunds, lost clients, lost sales and, worst of all, my focus. I could feel all of that far more sharply than I could picture the upside of a good hire. So I put it off, and putting it off made the bad outcome more likely.

What would have made it safe to say yes? Proof that the person had done the exact jobs I needed, in an agency, not just in-house. Most candidates had done some of them, and rarely all. Or, better still, some kind of guarantee that it would work. As far as I know, nobody offers one.

Every buyer I've sold to has felt some version of that fear. These days I try to give them what I wanted back then: proof from someone in their position, and a first step small enough to undo.

Bought is the stage where the buyer has picked you and now has to act. Then they have to sign, pay or tell their boss, and many of them stop.

Most brands read that stop as a lost sale to a rival. Often there was no rival. The buyer simply did nothing. So they push harder at the moment the buyer is most afraid, which is exactly the wrong time to push.

The rival is doing nothing

In a considered purchase, your toughest competitor is the buyer's own inaction. Every proposal sits next to a free option: carry on as before.

Matt Dixon and Ted McKenna analysed about 2.5 million recorded sales calls from several dozen companies. Between 40% and 60% of qualified deals ended in no decision ([Dixon](#)). The buyer didn't pick anyone.

They then split those losses by cause. In 44% the buyer preferred the status quo. In 56% the buyer wanted to act but couldn't decide, out of fear of getting it wrong ([Dixon and McKenna's deck](#)).

One caveat. This comes from the authors' commercial book, *The JOLT Effect*. The sample mixes quick and complex sales, and nobody has replicated it. So I wouldn't claim that most lost sales everywhere are lost to nothing. Many are, and in complex business purchases it may be about half.

The reframing matters more than the exact figure. If the rival is doing nothing, you can't beat it with a better feature list. Doing nothing has no features. It wins by being safe.

Why doing nothing feels safe

Doing nothing wins because people tend to weigh a loss more heavily than an equal gain. A review of 607 estimates from 150 articles put the ratio at 1.955

40–60%

Qualified deals that ended in no decision

DIXON AND MCKENNA

([Brown, Imai, Vieider and Camerer](#)). A loss hurts just under twice as much as the same gain pleases.

That idea isn't settled. David Gal and Derek Rucker reviewed the evidence and concluded it doesn't show that losses generally loom larger than gains ([Gal and Rucker](#)). I use the meta-analysis as the average, and their critique as a warning that the effect depends on the situation.

William Samuelson and Richard Zeckhauser named the second force in 1988. "Most real decisions," they wrote, "have a status quo alternative", and people choose it far more often than its merits suggest ([Samuelson and Zeckhauser](#)).

A third finding explains the fear. Ilana Ritov and Jonathan Baron found people would rather suffer harm from not acting than from acting. Their subjects refused a vaccine whose risk was smaller than the disease ([Ritov and Baron](#)). The scenarios were hypothetical, but the pattern is well known.

Put together, the three explain the stall. Buying and being wrong feels like your fault. Not buying and being wrong feels like bad luck. A nervous buyer will take bad luck every time.

One limit. For small amounts, one study found the pattern reversed: "gains loom larger than losses" ([Harinck, Van Dijk, Van Beest and Mersmann](#)). My reading is that fear matters most when the stakes are high. For a cheap, quick purchase, the problem at the last step is mostly friction.

Pressure loses sales. Safety wins them.

When a buyer stalls from fear, pressure makes the fear worse. Safety removes it.

Dixon and McKenna report that sellers win about 6% of deals with deeply indecisive buyers, while high performers win 31% ([Dixon](#)). Their deck adds that the usual pressure tactics backfire 84% of the time. That figure comes from the same dataset, and it's unreplicated.

The high performers judged the indecision, offered a recommendation, limited the exploration, and took risk off the table. None of those is a push.

I've seen it work from the seller's side. In 2021 the founder of a small consultancy reviewed our proposal for a new website with us, then went quiet for six weeks. When she finally wrote, it wasn't to say she'd chosen a rival. A long lockdown had stalled her main source of revenue, and she needed to keep cash in reserve. Doing nothing was the safe choice.

I didn't push. I told her I'd probably have made the same call in her shoes, and mentioned, with no pressure, that the cost could be spread over twelve months. Ten months later her team wrote back: a new financial year, a better position, were we still keen? She signed the original proposal unchanged, and she's still a

client. Pressure would have won nothing that year, and might have cost us the next.

But scarcity usually works. The evidence doesn't say urgency always fails. Barton, Zlatevska and Oppewal combined 416 effect sizes from 131 studies. On average, scarcity raised purchase intention, and time-limited offers worked best for high-involvement products (Barton, Zlatevska and Oppewal). That's exactly where I argue against pressure. Two distinctions reconcile the finding with the method.

Real scarcity is information. If the price really rises on the first of the month, or only two install slots are left, say so. That helps them plan.

Manufactured pressure is a lie, and now it's unlawful. Arunesh Mathur and colleagues crawled about 11,000 shopping sites. They found 393 countdown timers, and 157 of them, on 140 sites, were deceptive: they reset, or the offer outlived the clock (Mathur and colleagues). In April 2026 an Australian court ordered Emma Sleep to pay AUD 15 million over false savings and countdown timers that reset (HWL Ebsworth).

So my working rule is this, and it's mine rather than a tested finding. A real deadline is fine, as long as a way out sits underneath it. Put the way out first and the deadline second.

Where the last step lives now

The last step increasingly happens on your own ground, after a machine has done the comparing.

Walmart let shoppers buy inside ChatGPT or click through to its own site. Purchases made inside the chat converted at a rate "three times lower" than those that went to the retailer's site (CNBC). CNBC also cites inaccurate item data and onboarding problems, so this isn't purely trust.

My reading: the machine suggests and the buyer pays where they trust. So the landing page has to repeat the machine's reason, and state the way out next to the price.

In business buying, the last step has also grown more crowded. G2 surveyed 1,038 software decision-makers in June 2026. 49% had seen a finance chief veto a purchase that was already approved. Finance involvement rose from 31% to 46%, and budget approval delayed 32% of purchases (G2).

How to tell it is failing

- ✓ Proposals go quiet after a meeting that felt positive.

- ✓ Your CRM's "lost" column says "no decision", "timing" or "budget" more often than a rival's name.
- ✓ Buyers ask "can we get out of it?" or "what if it doesn't work?" near the end.
- ✓ Checkout or sign-up analytics show people reaching the price, then leaving.
- ✓ The finance or IT person first appears in the last week, with questions nobody prepared for.

The method: the risk removal ladder

The method has two parts. First, audit the last step for friction and fear. Then climb the ladder: remove the fears one rung at a time, starting with the lowest-risk step a buyer could take. The ladder is my working rule; each rung rests on evidence noted below.

The audit.

1. **Name the rival.** Write "doing nothing" at the top of your competitor list. For each lost deal last quarter, mark whether it went to a rival or to nobody.
2. **Walk the last step yourself,** as a first-time buyer, on a phone. Baymard Institute puts average cart abandonment at 70.22% across 50 studies ([Baymard](#)). Excluding "just browsing", 40% of US shoppers left over extra costs. 19% didn't trust the site with their card, and 18% were forced to create an account.
3. **Map the fears, by person.** Wasting money, looking foolish, being locked in, a failed rollout. Write each fear in the buyer's words.

The ladder. Five rungs, then an optional cap, in the order of the diagram.

1. **A small first step.** A trial, a pilot, a first month. Keep it short and active. One randomised trial covered 337,724 users at a large software firm. A 7-day free trial produced 5.59% more subscriptions than a 30-day one ([Yoganarasimhan, Barzegary and Pani](#)). Longer trials left people idle.
2. **A way out.** A money-back guarantee, lenient returns, or no lock-in. Experiments found guarantees raised purchase intention and willingness to pay a premium, more so for experience goods ([Suwelack, Hogreve and Hoyer](#)). A meta-analysis of 22 papers found refund leniency lifted purchases, though also returns ([AMA summary of Janakiraman, Syrdal and Freling](#)).
3. **A plain recommendation.** "If I were you, I'd pick this one, because..." A hard choice stalls, and a recommendation turns it into an easy one.
4. **Proof from someone like them.** One named customer in the same position, with a result.
5. **Someone to call.** A real person, with a name and a number, who answers. This rung is my addition.

6. **The cap: an honest deadline, only if one exists.** Test it first. Would it survive a regulator asking “what happens when the timer hits zero?” If not, remove it.

Sales, by hand

In considered and business purchases, a person still closes the sale. The salesperson’s work is to confirm the reason, supply the proof and remove the risk, not to add urgency.

In owner-led sales, one person decides, and the risk is personal. They want a price, a timeline and someone to call. A short call with the owner often does more than a long proposal.

In committee sales, each person needs their own fear answered. Someone who walks out of the meeting frustrated, tired and confused defaults to doing nothing. Give your champion a one-page case they can forward without editing. G2’s 49% suggests finance will read it, so write it for finance first.

When a deal goes quiet, the instinct is to chase. My working rule is the opposite. Make the next step smaller, and say so.

The bookkeeper’s ladder

The bookkeeper for cafés is an invented example. Café owners are busy, and many have been burnt by an accountant before.

RUNG	WHAT SHE OFFERS	THE FEAR IT REMOVES
A small first step	One free Monday report, built from last month’s takings and rosters	“It won’t be worth the effort to hand over my numbers”
A way out	Monthly, no lock-in; cancel by email and keep your files	“I’ll be stuck with someone useless”
A plain recommendation	“For one café, start with the weekly report. Add payroll when you trust me.”	“I don’t know which package I need”
Proof from someone like them	One café’s twelve months of numbers, shown with the owner’s permission	“That works for big places, not mine”
Someone to call	Her mobile number on the page; she answers before the morning rush ends	“I’ll be talking to a call centre”

None of these cost her much. The free report takes an hour, and the monthly terms only cost her clients who would have left anyway.

Fieldline's last step

Fieldline, the invented scheduling-software firm, has reached the last step with Northgate Heating & Cooling. Dana now has a reason she can repeat. The CEO's question was "what does it cost us if it doesn't work?" Then nothing happened for five weeks.

The account manager's first draft was a "price rises next month" email. Instead, she mapped the fears.

PERSON	THEIR FEAR AT THE LAST STEP	THE RUNG THAT REMOVES IT
Dana, head of operations (champion)	"If the technicians hate it, I'll wear it."	A small first step: a pilot in one depot, with the travel-time measure agreed before it starts
Ravi, finance lead	"This becomes a three-year cost we can't cancel."	A way out: an exit after the pilot with nothing owed beyond it, in writing
Mei, IT manager	"Security review will take months, and I'll get blamed for the delay."	Someone to call: Fieldline's security lead, plus the review questionnaire answered in advance
The CEO, who signs	"We change systems in our busy season and lose a quarter."	A plain recommendation: start after the peak, on one plan, named with the reason
All four	"Will this work for a firm like ours?"	Proof: a named contractor of similar size, and its pilot result

She sent Dana a one-page case written for Ravi and the CEO. It led with the cost of doing nothing: the hours lost to travel each week.

It didn't work. The CEO read the case, agreed the numbers were real, and said no. Northgate would stay with its spreadsheets for another year. The busy season was coming, two technicians had just left, and a new system felt like one more thing.

Fieldline did the one thing left. It sent a short email offering a smaller first step, ending "if now isn't the right time, just say so. I won't chase." Dana replied that it wasn't, and thanked them.

The story is invented, and I have made it end this way on purpose. Every rung was in place, and doing nothing still won. The ladder raises the odds of a yes; it doesn't make one certain. A method that claims to win every deal is selling you something.

What Fieldline kept was a champion who could repeat its reason, a CEO who had seen the numbers, and a door left open.

By mode

Each buying mode stalls at the last step for a different reason. Habitual and quick-choice buyers stall mostly on friction: a stock-out, hidden costs, a forced account, a long checkout. Fix the top two items from your walk-through first. Considered personal buyers fear a costly mistake that can't be undone, so lead with a way out. Owners carry the risk personally and have no one to ask: give them a plain recommendation and someone to call. Committees need each person's fear answered, a one-page case for finance and a scoped pilot. The friction and fear split is my reading of the small-stakes evidence, not a tested rule.

How Bought connects to the other stages

- ✓ **It needs a reason.** The landing page must repeat the reason the buyer, or the machine, was given. Without it, the buyer arrives and finds a stranger.
- ✓ **It needs proof and answers.** Proof from someone like them makes the risk feel small, and many last-step fears are questions that should already have been answered.
- ✓ **It feeds Kept.** A way out only builds trust if you honour it, and keeping the customer starts the moment the first step ends.

What not to do

- ✓ **Don't fake the clock or the crowd.** Resetting timers cost Emma Sleep AUD 15 million in Australia.
- ✓ **Don't make the first step long,** or hide the total.
- ✓ **Don't treat a lost-to-nothing deal as dead.** Leave the door open in writing, and keep the champion's reason alive.

How you know it worked

The Selectability Score (Part five) checks Bought on the page a buyer lands on. Does it repeat the reason they were given? Does it show proof? Does it give a clear next step? Where you have your own sales data, the score adds that too.

Then watch your own numbers.

- ✓ The share of qualified deals ending in "no decision", tracked each quarter.
- ✓ The time from proposal sent to yes or no.
- ✓ The share of first steps (trials, pilots, free reports) that become paying customers.
- ✓ Checkout or sign-up abandonment at the price step.

This afternoon: pull your last ten lost deals. Count how many went to a rival and how many went to nobody. Then walk your own last step as a buyer would, and note every surprise.

This quarter: build the ladder for your main moment. Add one small first step, state the way out next to the price, and name one recommended option.

TOOLKIT

The risk removal ladder. Use it once per moment, after the audit has mapped the buyer's fears at the last step. Start at the bottom rung and add each one before the next.

The risk removal ladder: make each yes a little bigger, and safe




FIGURE 18.1

Read it from the bottom up. Each rung is a slightly bigger yes, and any deadline sits on top of the way out.

Every template for this stage is in the free Selectability Playbooks at rocmode.io/playbooks.

What we don't know

- How big “no decision” really is.** The 40–60% and the 44/56 split come from one proprietary dataset. No one has replicated it, and it mixes simple and complex sales.
- Whether safety beats urgency head to head.** I found no study comparing the two in high-stakes business buying. “Way out first, deadline second” is my working rule.
- How much guarantees move real purchases.** Most guarantee evidence is from lab studies of intention. I found no large field trial in business buying.

Whose evidence this is. Dixon and McKenna’s data came from Tethr, a conversation-analytics firm where Dixon worked, and the authors sell JOLT training. G2 runs a software review marketplace. Baymard sells e-commerce research. Peer-reviewed sources: Brown and colleagues; Gal and Rucker; Samuelson and Zeckhauser; Ritov and Baron; Harinck and colleagues; Barton, Zlatevska and Oppewal; Suwelack, Hogleve and Hoyer; Janakiraman, Syrdal and Freling.

IN SHORT • CHAPTER 18

WHAT’S TRUE

Most lost sales are lost to nothing. Doing nothing wins by being safe, so pressure loses and safety wins.

DO THIS QUARTER


Build the risk removal ladder: a small first step, a way out next to the price, and one recommended option.

WATCH

The share of qualified deals ending in no decision, time from proposal to answer, and first steps that convert.

CHAPTER 19

Be kept



“I'd choose you again.”

DEFINITION

Kept is where a customer chooses you again, or someone else, and more of them ask a machine first.

THE BUYER ASKS

“Are they still the best option for me?”

MEASURED BY

Renewal or repeat rate, referrals, fresh reviews, and whether AI answers about you cite your customers' words.

FAILURE LOOKS LIKE

Nobody owns the renewal, reviews go stale, and assistants quote complaints you never answered.

Today's customers write the evidence tomorrow's buyers read.

Kept is the stage where a customer chooses you again. They can also choose someone else, and more of them now ask a machine before they decide.

Most brands treat Kept as the point where marketing's job ends. The customer passes to service, and nobody owns the moment they decide whether to stay.

That gap costs more than one renewal. A kept customer writes much of the evidence the next buyer reads: the referral, the review, the forum post. Lose customers quietly and the earlier stages weaken, often with a delay and no obvious cause.

Every renewal is a fresh sale

A date arrives, a price rises or a rival launches, and the customer asks whether you're still the best option. More and more, they ask an AI assistant.

1.3x

More likely to switch insurers after using AI tools

J.D. POWER

The clearest evidence comes from car insurance. J.D. Power's 2026 U.S. Auto Insurance Study surveyed 52,216 customers, and 32% of shoppers had used AI tools. Those who did were "more than 1.3 times as likely to switch insurers" ([J.D. Power](#)).

That's a correlation. People who already meant to leave may be the ones who turned to AI. A third of the AI users also found the AI content unhelpful. The adviser doesn't have to be good to be consulted.

Business buyers are on the move too. Google and National Research Group surveyed 2,063 US business leaders. Of those who had bought in the past six months, 58% had also switched supplier in that time ([Digital Commerce 360](#)). Read it as a direction, not a precise rate.

So the renew-or-switch moment is now a sale with an adviser in the room. The answer to "alternatives to X" is built from what others have written about you, much of it by your own customers.

Being the incumbent doesn't exempt you. In 2025 a state government agency we'd supported for more than four years put its website rebuild out to tender. We pitched as the incumbent, and lost.

The debrief was fair. Our method was well defined and our build was good value. But our case studies didn't show enough of the accessibility outcomes the panel cared about. Our team looked small for the risk. And our hosting and maintenance costs weren't set out clearly enough to compare. A newcomer would have had to answer every one of those questions. We'd assumed our history answered them for us. It didn't.

A renewal is judged like a first sale, against everyone else's evidence. We kept their campaign work, and we kept the lesson.

The loop

Kept customers produce the material the earlier stages run on. The seven stages run as a loop. Here's how it closes, one link at a time.

Referrals build Known. A happy customer puts your name in front of someone who has never heard it. At a German bank, researchers compared 5,181 referred customers with 4,633 others ([Schmitt, Skiera and Van den Bulte](#)).

Referred customers had an 18% lower risk of leaving, and that gap persisted. Their margin started about 25% higher, but that edge faded after about 29 months. Over six years they were worth about 16% more.

So a referral pays mostly through retention. It's one bank, but the direction is clear.

Your best spenders aren't your best referrers. Kumar, Petersen and Leone studied a telecom firm and a financial services firm. Customer lifetime value and customer referral value picked out different people ([Kumar, Petersen and Leone](#)). The work appeared in HBR, not a peer-reviewed journal.

Reviews build Believed. Chapter 16 set out how reviews lift purchase odds. What matters here is who writes them, and that depends on whether you ask.

Askalidis, Kim and Malthouse studied data from four retailers ([Askalidis, Kim and Malthouse](#)). Customers prompted by email gave ratings up to 0.5 stars higher than people who reviewed unprompted. Their ratings stayed stable while unprompted ones drifted down. Unless you ask, your quiet, satisfied customers are missing from the record.

The folklore says angry customers tell everyone. The unhappy do talk more, "but the difference is smaller than commonly believed" ([Anderson](#)).

Forums and review sites feed Found and Preferred. Reddit, Trustpilot and G2 all appear among the sources AI assistants cite. The vendors who count them disagree on how much (Chapter 22). When a buyer asks a machine to compare suppliers, a thread from a disappointed customer can sit next to your own pricing page.

Two companies show the loop working. In 2010 Dropbox's co-founder, Drew Houston, said a two-sided referral programme "permanently increased signups by 60%" ([Houston](#)). These are the company's own figures.

Tally shows the newer version. It grew for years on word of mouth, and now credits AI assistants with much of its discovery. It says "forums, Reddit posts, blog mentions, and authentic UGC" fed that ([Tally](#)).

No study measures the whole loop. Nobody has shown that churn or service failures led, later, to fewer AI recommendations. So the loop is my reading, not a measured effect. I act on it because each link pays for itself, as the referral sum below shows, even if the whole loop never closes.

Delivery is the marketing

Customers judge you against the promise that won the sale. Richard Oliver showed in 1980 that satisfaction depends on what a customer expected, and how far performance confirmed it ([Oliver](#)). Satisfaction then shapes whether they buy again.

The distance between the two is where the next review comes from. A promise kept is a referral in waiting. A promise broken is a forum post in waiting.

That's the idea behind the **promise ledger**: what was promised, where, how delivery is measured, the gap and the owner. Every gap in it is a future review.

Most of what it tracks sits outside marketing, in product, delivery and service. Marketing's part is to know the promises word for word, and to hear early when one slips.

Two companies that cut the wrong thing

Two well-known firms cut service at the very stage where customers decide whether to stay. Both then backed away.

In February 2024, Klarna said its OpenAI-powered assistant was handling two-thirds of its service chats and did the work of 700 agents ([Tech.eu](#)). In May 2025, chief executive Sebastian Siemiatkowski told Bloomberg that cost had been “a too predominant evaluation factor”, and that quality had suffered. “Really, investing in the quality of human support is the way of the future for us,” he said ([Entrepreneur](#)). The quotes come from secondary reports of that interview.

Commonwealth Bank of Australia cut 45 customer-service roles in 2025, citing a voice bot. On 21 August 2025 the bank reversed the cuts. It said it “should have been more thorough in our assessment of the roles required” ([ABC](#)). The reversal followed union pressure, so it shows a misjudged cut, not proven customer loss.

Both firms treated service as a cost to shrink. In the loop, service is where the next review gets written.

How to tell it is failing

- ▮ Nobody can name the person who owns the renewal moment.
- ▮ The first renewal conversation happens when the invoice goes out.
- ▮ An AI assistant, asked whether your product is worth renewing, quotes complaints you have never answered.
- ▮ Your newest public reviews are months old, or come mostly from unhappy customers.
- ▮ Sales promises things the delivery team has never heard of.

The method: the promise ledger and the renewal moment map

The promise ledger catches delivery gaps before customers write them up. The renewal moment map shows what each customer, and the machine they ask, will say before they decide.

1. **Collect the promises, word for word.** For each kind of customer, list the three to five promises that won the sale.

2. **Note where each one is made.** Ask an AI assistant what you promise too, because its version is the one many buyers hear.
3. **Give each promise a delivery measure.** Pick something you can observe: a report sent on time, a reply within a day, a result in the customer's own data.
4. **Name one owner and one early warning.** The owner is usually outside marketing. The warning is a ticket theme or a review phrase.
5. **Rate and write the gap plainly.** Use the scale below. Each gap is a future review. Either fix the delivery or change the promise.
6. **List the renewal triggers.** With a contract: the renewal date, a price rise, a failure, a rival's launch. Without one, list the repeat-purchase triggers: the season, the next job, a bad experience, a rival's offer, a new person in charge, a long gap. For each, note what the customer will ask, often of a machine, and what the machine says now.
7. **Send proof before the question.** With a contract, plan what you send at 90, 60 and 30 days. Without one, time it to the trigger: before the season, after the job, when the usual gap has passed. Send proof of value in each decision-maker's own terms.
8. **Review both every quarter.** Check them against fresh reviews and fresh AI answers. Read renewal, referral and review numbers together, allowing for a lag.

A simple keep rating helps the ledger set priorities. Higher is better, like every other scale in this book. This is my own working rule, not tested against sales.

RATING	WHAT IT MEANS	WHAT TO DO
3	Kept, with evidence you could show a buyer	Turn it into proof, and keep asking every customer for a review
2	Slipping: early warnings, no public complaints yet	Fix this quarter
1	Broken, and customers have noticed	Fix now, and answer every public mention
0	Broken, and machines repeat it	Fix, then change the promise until delivery catches up

Reviews and referrals: ask everyone, reward only referrals

Most of the review method belongs to Believed. Three points belong to this stage.

- / **Ask every customer, at a fixed point.** Pick one moment, such as the end of the first quarter, and ask everyone who reaches it, happy or not. Asking only the customers you expect to be pleased is review gating. The US Federal Trade Commission's guidance says "Don't ask for reviews only from customers you think will leave positive ones" ([FTC](#)). In the UK, the DMCC

Act's review bans took effect on 6 April 2025. The CMA's guidance treats cherry-picking positive reviews as one of them ([CMS](#)).

- ✦ **Never reward a review.** Rewards are for introductions, not opinions. Keep the referral ask in a separate message, sent at a different time.
- ✦ **Ask on the sites buyers and machines read,** not only on your own site, and answer every negative review in public.

For referrals, pick referrers by referral value, not spend, and reward both sides, as Dropbox did. Expect the payback mostly as retention. In the bank study, the margin edge faded and the lower churn stayed.

The referral sum, for finance

A CFO will ask what a referral is worth. Every input below is invented except the bank study's 16%; replace them with yours.

LINE	WORKINGS	RESULT
A new customer's lifetime margin	Your own figure	£10,000
Cost to win one through paid channels	Your own figure	£2,000
A referred customer's lifetime margin	£10,000 × 1.16 (the bank study)	£11,600
Reward paid per referral	£200 to each side	£400
Share of referrals that would have come anyway	Your estimate	50%
Reward cost per extra customer	£400 ÷ 50%	£800
Net value: paid-channel against referred	£10,000 – £2,000, against £11,600 – £800	£8,000 against £10,800
Difference per customer		£2,800

The sum is arithmetic, not evidence, and the 16% comes from one bank. The “would have come anyway” line matters most, so test it: offer the reward to half your customers for a quarter, and compare.

What keeping cannot do

Keeping customers doesn't, on its own, grow a brand. Romaniuk, Dawes and Nenycz-Thiel studied 1,093 brand-share changes across 63 UK packaged-goods categories ([Romaniuk, Dawes and Nenycz-Thiel](#)). When brands grew, penetration changed about three times as much as loyalty. When manufacturer brands declined, they lost both.

I accept that. My claim is narrower: not that retention beats acquisition, but that kept customers produce it.

Stated intentions are weak evidence too. Keiningham and colleagues found Net Promoter Score was rarely the best predictor of revenue growth among Norwegian firms ([MeasuringU](#)). Count who renewed and who referred, not who said they would. And loyal customers aren't better advocates: repeat customers' reviews didn't beat first-timers' at attracting new customers ([Shin et al.](#)).

So the loop runs on volume and freshness of honest evidence. You need many kept customers saying true things recently, not a few loyalists saying glowing ones.

The bookkeeper's promise ledger

The bookkeeper for cafés is an invented example. Here's her ledger after one quarter.

PROMISE	WHERE IT'S MADE	HOW DELIVERY IS MEASURED	RATING AND GAP	OWNER
Wage cost against takings, every Monday by nine	Home page, proposals, directory listing	Share of Mondays the report lands by 9am	2: three late Mondays in the school holidays	Her
Only cafés	Home page, partner pages	Share of clients who run cafés	3: none	Her
Numbers you can act on	Sales calls, the one café's twelve-month case	Clients who change a roster after a report	2: two clients never reply to the report	Her assistant
Someone who answers	Proposals	Replies within one working day	2: fine, except during month-end	Her assistant

Three late Mondays are three chances for a review that says "great, when it arrives". She moved the send time earlier in holiday weeks and told clients why.

Her clients have no renewal date. They stop after a bad month, or when a new manager arrives. So she sends a three-month review before the summer holidays, and a note to any new manager on day one.

Fieldline's renewal moment map

Fieldline, the invented scheduling-software firm, lost Northgate to doing nothing. Meanwhile, Harbour Plumbing, an invented customer of similar size, reached its renewal.

DAYS BEFORE RENEWAL	WHO IS ASKING	WHAT THEY ASK, OFTEN OF A MACHINE	WHAT FIELDLINE SENDS
90	Operations manager	"Is Fieldline still the best scheduling tool for plumbing contractors?"	A review of travel time in Harbour's own data, set against the 2023 pilot
60	Finance lead	"Cheaper alternatives to Fieldline"	One page in finance terms: cost, hours saved, job costs against invoices, and what switching would cost
30	IT lead	"Fieldline security issues"	The current security page and a short change log since signing

At 60 days the finance lead did ask an assistant for cheaper alternatives. But he already had the 90-day review, in his own numbers. Harbour renewed, though finance pushed the price down first. Keeping a customer isn't the same as keeping your margin.

Fieldline asks every customer for a review at the end of their first year. Harbour's operations manager wrote a long one on G2. It gave the travel-time figure and a detail Fieldline never led with: job costs now matched invoices without a weekly spreadsheet. She told the same story at a regional trade forum.

Nine months after Northgate said no, Ravi was planning next year's budget. He asked an assistant: which field-service software makes job costs match invoices? This time Fieldline was named, and the answer cited Harbour's review.

Ravi forwarded it to Dana. Dana forwarded it to the CEO, with her own line about the vans on top. Northgate asked for a second meeting, and Chapter 27 picks up from there.

The story is invented. It shows the loop as I believe it works. A kept customer, asked like everyone else, wrote the evidence the next buyer read.

By mode

The renewal moment looks different in each mode. Habitual buyers re-order on autopilot until a stock-out or a bad batch prompts a look, so consistency and availability come first. Quick-choice buyers come back only if they remember you: a good first use, then a review ask at a fixed point. Considered personal renewals, such as insurance, are often checked with AI, so send proof of value before the notice arrives. Owners have no date and wonder after a bad month. Committee members each ask their own question as the contract end nears, so map the renewal moment per person.

How Kept connects to the other stages

- It needs **Bought and Preferred**. The promise made at the sale, and the reason behind it, are what delivery is measured against.
- It feeds **Known, Found and Believed**. Referrals spread your name, and reviews and forum posts are the evidence others and machines cite.
- The builds carry it**. Part four covers the email, automation, PR and review work that runs this stage.

What not to do

- Don't cut service to fund acquisition**. Klarna and Commonwealth Bank of Australia both cut where customers decide to stay, and both reversed.
- Don't ask only the happy ones**. It skews the record, and regulators in the US and UK treat it as deceptive.
- Don't promise what delivery can't measure**. If nobody can check it, nobody will catch the gap until a customer posts it.

How you know it worked

The Selectability Score (Part five) asks the same checks at the renew-or-switch moment. Does the machine name you when a customer asks for alternatives, for what reason, and accurately? Your own repeat and renewal data is the truth.

Then read four numbers of your own, together and allowing for a lag:

- Renewal or repeat rate, by kind of customer.
- The share of new customers who came by referral.
- Fresh public reviews per month, and their average rating.
- The share of AI answers about you that cite your customers' words.

This afternoon: ask two AI assistants "alternatives to [your brand]". Note which sources they cite, and whether any are your customers' words. Then write down the three promises that won your last five sales.

This quarter: build the promise ledger and give every promise an owner. Set one fixed point for the review ask, and ask everyone. Send the proof before anyone asks.

TOOLKIT

The selection loop. Walk one real customer round it, and see where the loop leaks.

The selection loop: every promise kept recruits the next buyer




FIGURE 19.1

Read it clockwise from the next buyer: each promise kept feeds referrals and reviews, which recruit the buyer after that.

Every template for this stage is in the free Selectability Playbooks at rocmode.io/playbooks.

What we don't know

- / **Whether churn later weakens AI recommendations.** No study links a firm's churn or service failures to a later fall in how often AI cites or recommends it. That's the central link in this chapter, and it's untested.
- / **How machines weigh customer text.** How AI training and retrieval weight review and forum text isn't public. The citation shares used in this book come from vendor trackers.
- / **Business renewals and AI.** I found no B2B evidence on AI use at renewal. The J.D. Power study covers US car insurance only.

Whose evidence this is. J.D. Power sells benchmarking and data to insurers. Google sells advertising, and G2 sells visibility on its review site. Kumar consults on customer value. Keiningham worked at Ipsos Loyalty, which sells rival measures. Dropbox's and Tally's figures are their own. Peer-reviewed sources: Schmitt, Skiera and Van den Bulte; Anderson; Oliver; Romaniuk, Dawes and Nenycz-Thiel; Shin and colleagues.

IN SHORT • CHAPTER 19

WHAT'S TRUE

Today's customers write the evidence tomorrow's buyers read. Every renewal is a fresh sale, with an adviser in the room.

DO THIS QUARTER

Build the promise ledger, give each promise an owner, ask every customer for a review at one fixed point.

WATCH

Renewal or repeat rate, share of customers from referral, fresh reviews per month, and AI answers citing customers' words.

WHERE YOU ARE

01

Why it feels like this

02

What a choice is
made of

03

The way forward

04

The builds

05

Knowing it worked

04 / PART FOUR

The *builds*



PART FOUR

The *builds*

Every build is a different mix of the same seven stages.

Part three said what has to be true at each stage. But nobody builds a stage. A team builds things: a website, content, advertising, a pitch, email, events and a dozen more.

Those builds are usually run as separate crafts, each with its own team, agency, budget and dashboard. Each is judged on its own numbers: reach, traffic, rank, opens. None of those numbers says whether a buyer chose you.

Six categories, fourteen builds

There are more ways to reach a buyer than any team can run, so it helps to sort them by one question: who is speaking, and how did they reach the buyer? Six answers cover almost everything.

- ✓ **Owned:** your voice, on ground you own. The website, content and social.
- ✓ **Paid:** your voice, in space you buy. Advertising.
- ✓ **Earned:** someone else's voice, freely given. Search and AI answers, and PR and reviews.
- ✓ **Partnered:** someone else's voice or audience, by agreement. Partners and channels, and creators.
- ✓ **Direct:** a person, live. Outbound and account-based selling, the pitch, and events.
- ✓ **Relationship:** your system, talking to someone you already know. Email, automation, and customers: onboarding, loyalty, community and referral.

This extends the PESO model (paid, earned, shared, owned) that Gini Dietrich built for communications. For most firms, a person in a room and the system that remembers a customer decide as many moments as any channel, so each gets a category of its own.

That makes fourteen builds. A channel earns its own build only if it has its own shape, its own craft, and moments it leads better than anything else.

Sponsorship, affiliates, awards, account management and the rest do their work through a build that already exists, so they get a short note, not a chapter. Offer, price, product and distribution aren't builds at all: they're foundations, set before any build starts.

Ten builds get a chapter here. Outbound, events, partners and customers get shorter sections at the end of Part four.

Lay the ten chapter builds against the seven stages and the picture changes.

Ten builds, seven stages: no build covers them all,
and no stage belongs to one build.

● mainly serves ○ serves in part

	Known	Found	Answered	Believed	Preferred	Bought	Kept
Content (Ch 20)	○	○	●	●	●	○	
The website (Ch 21)		○	●	●	○	●	○
Search and AI answers (Ch 22)		●	○	●	●		
Advertising (Ch 23)	●	○			○	○	
Social (Ch 24)	●	○	○	●	○		
Creators (Ch 25)	○			●	○		
PR and reviews (Ch 26)	○	○		●	○		
The pitch (Ch 27)				●	●	●	
Email (Ch 28)			●	○	○	●	●
Automation (Ch 29)			●	○		●	●
Builds that mainly serve it	2	1	4	7	3	4	2

Ten builds × seven stages · filled = mainly serves, hollow = in part

Read across a row and you see what a build is for. Read down a column and you see something more useful: who is responsible for a stage. Believed is carried by seven builds. Known rests mostly on advertising and social. Kept rests on email and automation, and on a website most teams never think of as a retention tool.

So no build owns a stage, and no stage belongs to one team. When Believed fails, it may be the website's proof, the PR team's reviews or the pitch's case study. A team organised by craft can't see that. A team organised by stage, or by moment, can.

That's the change Part four asks for. Keep the crafts, because each needs skill. But give each stage an owner who looks across the builds, and judge every build by the stages it's meant to move, using the Selectability Score (Part five).

Brief it twice

Most marketing leaks at the brief. The IPA's BetterBriefs survey asked more than 1,700 marketers and agency staff about briefs. Eight in ten marketers said they wrote good ones, and one in ten agency staff agreed ([IPA](#)).

So use two briefs. The **campaign brief** carries one short idea, the moments and stages it must move, and a separate box of rules. The **asset brief** carries five lines for each thing you make: moment, stage, reason, proof and next step. The same five lines brief a person or an AI.

If you can only do two

A small business can't run fourteen builds. My working rule: start with the website and search and AI answers. They share the same questions and the same pages, and together they cover Found, Answered, Believed, Preferred and Bought. Add PR and reviews third, because it's the cheapest route to being believed by others. The rest can wait until those three work.

For agencies, the grid is also a menu. An agency selling one craft sells one row. The work that wins now runs down a column: an audit of one stage across every build, a source map, a body of proof.

How each chapter works

Each build chapter gives the argument: what the build is for, its anatomy by stage, the mistake most teams make and how to know it's working. The full how-to, with the briefs, templates and checklists, lives in the free Selectability Playbooks at rocmode.io/playbooks.


CHAPTER 20

Content

THIS BUILD MOVES

SEEN

SELECTED



■ MAINLY MOVES

■ MOVES IN PART

Content has three jobs: answer, prove, persuade.
Anything that does none of them is filler.

What changed

Content is what every other build holds, quotes or carries. The website holds it, search and AI answers quote it, and social, email and sales carry it. Most teams run it as a calendar and judge it by output: posts published, traffic drawn, time on page.

And output is now the cheapest thing in marketing. Graphite sampled 55,400 article URLs published from January 2020 to March 2026 ([Graphite](#)). About half of new articles are now written mainly by machines, a share that has held near 50% since early 2025. Graphite adds that these articles “largely do not appear in Google and ChatGPT”.

Volume doesn't win with the machines either. Ahrefs compared 75,000 brands against how often AI assistants mentioned them, and found “almost no relationship between content volume... and AI visibility” ([Ahrefs](#)). More content is no longer an advantage. It's the default, and it's mostly ignored.

What it is for

Content has three jobs, and each belongs to a stage. **Answer** serves Answered: does it settle the buyer's question, in their words? **Prove** serves Believed: could a buyer, or a machine, check it? **Persuade** serves Preferred: does it name the bigger problem, and give your reason?

Content feeds Found, because it's what search and AI answers point at and quote, but getting found is their job. Nor does content close the sale. It hands the buyer to the website and the pitch, with a next step towards Bought.

~50%

of new online articles are
now written mainly by
machines

GRAPHITE

Some content exists to be remembered: a distinctive series, or an idea tied to a buying situation. That's Known work, and it isn't filler. But it's judged by Known's own tests, and it isn't excused them because it's called "thought leadership".

Buyers in the thick of choosing do two things. They hunt: they arrive with a question and want it settled. And they browse: they don't yet know what to want. The answer piece is for the hunter. The case piece is for the browser. Most teams write only one kind, so their content reads like either a manual or a pitch.

The anatomy of a good one

A good piece is Chapter 15's answer page with one part added: the bigger problem, and your reason.

A good piece answers first, proves it, then gives the reason




FIGURE 20.1

The buyer's question, settled first (Answered). One question, in the buyer's words, answered in the first lines. Then go deeper.

The detail, including the comparison (Answered). Make the comparison the buyer is already making: against the named rival, and against doing nothing. That's where the next questions live.

Proof (Believed). A customer with a name, a figure with a source, the month it was true. "One named café, twelve months, wage cost down from 38% to 31% of takings" is proof (the figures are invented). "Trusted by hundreds" isn't.

What can go wrong, and who it's not for (Believed). Saying where it doesn't fit is part of being believed.

The bigger problem and your reason (Preferred). Step back from the question to the problem behind it. Name it better than the buyer can, show what leaving it alone costs in the buyer's units, and only then give your reason. "Most cafés learn their wages were too high a quarter later" names a problem the owner has lived, but never put that way. Without this part, the buyer takes your answer to someone cheaper.

One next step (Bought). One small action, and what happens after it.

Worth quoting (Found). Machines quote what's specific. My four working rules: specific beats general, yours beats borrowed, checked beats claimed, and dated beats timeless.

The mistake most teams make

They let the machine decide. Anil Doshi and Oliver Hauser ran an experiment with 293 writers of eight-sentence stories, published in *Science Advances* (Doshi and Hauser). Writers given five AI-generated ideas wrote stories rated 8.1% more novel and 9.0% more useful. But the AI-assisted stories were more like each other than stories written by people alone. It studied fiction, not marketing. Anyone reading a category's content this year will recognise it.

When every machine drafts from the same internet, the pieces converge. Your facts, cases and reason are what pull yours away. So AI is the drafter, quick at structure, summaries and forms. A person decides the question, the reason and every fact. A machine given a moment, a stage, a reason and a proof produces work that fits. A machine given a topic produces more of what everyone else has.

How you know it's working

For each piece, ask its headline question of two or three AI assistants, several times each. Is your answer used, your figure quoted, your name given?

Alongside that, track the share of live pieces with proof that is named, numbered and dated within the last year.

Retire pieces published, words written, page views and time on page as goals. Keep them as diagnostics. A piece with many views and no effect is still filler.

The bookkeeper's content plan

The bookkeeper for cafés is an invented example. Her offer is "Bookkeeping for cafés. Your wage cost against takings, every Monday by nine."

She had 31 blog posts. Her audit kept four, fixed six and cut 21: tax-deadline reminders, office news, and tips any accountant could have written. Then she planned one moment: a café owner after a bad quarter, wondering whether a bookkeeper would help. She made an answer piece on what wages should cost a café, a named café's year of Monday reports, and a case piece on why most cafés find out a quarter too late. That's three pieces, eight uses and one moment covered.

In the first month, two of her referring accountants forwarded the case piece to clients. The answer piece has done less: an assistant asked about café wages still quotes a payroll firm's figure, not hers.

The full playbook. The Content Playbook, with the content audit scorecard, the asset brief, the campaign brief and the truth index, is free at rocmode.io/playbooks.

Whose evidence this is. Graphite sells search and AI-search services. Ahrefs sells SEO tools and an AI-visibility tracker. Peer-reviewed source: Doshi and Hauser.

IN SHORT • CHAPTER 20

WHAT'S TRUE

Content has three jobs: answer, prove, persuade. Anything that does none of them is filler, and more of it is no advantage.

DO THIS QUARTER

Build a plan for your chosen moments, at most three pieces each, and audit everything live: keep, fix or cut.

WATCH

Ask assistants each piece's headline question. Is your answer used, your figure quoted, your name given?


CHAPTER 21

The website

THIS BUILD MOVES

SEEN

SELECTED



The website is your place of record: fewer visitors, bigger moments.

What changed

Most websites are still run like shop windows. The team counts who came in and how long they stayed, and judges a redesign by whether traffic went up.

But much of a buyer's research now happens elsewhere: in an AI answer, a review site, a forwarded message. Pew Research tracked 900 US adults through 68,879 Google searches in March 2025. People clicked a result in 8% of visits that showed an AI summary, against 15% without one ([Pew](#)).

Those who do arrive are further along. Microsoft Clarity studied more than 1,200 publisher sites ([Microsoft Clarity, via PPC Land](#)). Visitors sent by AI assistants signed up at 1.66%, against 0.15% for visitors from search. These are sign-ups on publisher sites, not sales.

My reading: fewer people arrive, and in considered purchases more of them are close to a decision. A site judged on traffic will look like it's failing just as it matters more.

What it is for

The website's job is to be the place of record. It's the one place where your facts are stated in full, by you, with the proof beside them. Buyers come to check it, machines read it, and the sale closes on it.

It's the one asset you own that has a job at almost every stage. Its main jobs are Answered, Believed and Bought. It helps at Found, as the largest thing a machine reads about you; at Preferred, where the reason is made against the alternatives; and at Kept, through help and the account. Only Known is mostly won elsewhere.

8% vs 15%

Google searches clicked with an AI summary, against without one

PEW RESEARCH

The visitor who has grown is the checker. They've been told about you, by a machine, a colleague or an accountant, and they've come to confirm it. They aren't looking for a story. They need proof beside the claim, and the price and terms where they expect them. Checking is where Believed and Bought are won or lost.

The checker is quick to leave. If the price is missing, they ask the machine, which answers with whatever figure it can find. A site that ignores the checker loses the buyer who was already sold.

The anatomy of a good one

The sale is made on the service or product page, not the home page. It's where a buyer lands when a machine, a search or a colleague sends them.

A good page answers the checker: the question first, the proof beside every claim




FIGURE 21.1

The buyer's question, answered in the headline (Answered). "Bookkeeping for cafés: your wage cost against takings, every Monday by nine" answers "what is this and is it for me?". "Empowering hospitality success" answers nothing.

Who it's for, and who it isn't (Preferred). To be someone's first choice, you must be wrong for someone else. Say so near the top.

What it is, and what it costs (Answered · Bought). If you can't give one price, give typical prices for common cases. TrustRadius found 81% of software

buyers want to find pricing themselves ([TrustRadius](#)). If hiding price is a deliberate choice that fits your buyers, say plainly how pricing works and what moves it.

Proof beside the claim (Believed). A number next to its source, a result next to the named customer, a promise next to the guarantee.

The reason, against the alternatives (Preferred). Name what the buyer would do instead, including doing nothing. A machine that recommended you gave a reason. The page should repeat it, or the buyer arrives and finds a stranger.

The way out, and one next step (Bought). The trial, the guarantee and the cancellation terms, stated next to the price. Then one action, the lowest-risk yes the buyer could give.

Help and account, for customers (Kept). A quiet link for people who already bought. If they can't find help, the renewal starts badly.

The order is the order of the buyer's own doubts, from "is this for me?" to "what do I do now?".

Brief every section, not just the page. Each section of a page carries one job: one stage, for one chooser at one moment. So each gets its own short brief: the job, the one claim, the proof beside it, the facts taken word for word from your one truth, the line a person should take away and the sentence a machine should be able to quote. The page brief keeps the order and the arc, so the page still reads as one story. Briefing this way gives sharper writing, from people and from AI. And because assistants tend to pull passages rather than whole pages, every section can stand on its own.

The mistake most teams make

They put the claims on the service page and the proof on a "testimonials" page three clicks away. The checker never finds it. Your own word counts for least: the checker needs someone else's word, and needs it on the spot.

The second mistake is hiding the price by default. "Contact us" sends checkers to a competitor, or to whatever figure a machine can find. Hiding the price because nobody decided isn't a strategy.

And many teams rebuild on faith. I know of no study that isolates proof beside the claim, or a single next step, and measures the sales it adds. So rebuild your most important service page from the anatomy, keep the old one running for half the traffic, and compare next steps taken. Then take the result to finance.

How you know it's working

Retire traffic as the goal. Visits, time on site and bounce rate stay as diagnostics. Instead, track the share of visitors to a service page who reach pricing, and who take the next step. And ask new customers what they checked before buying; analytics won't show it.

Fieldline's website

Fieldline is the invented scheduling-software firm from Part three. In one invented deal, at a 70-technician plumbing contractor, three people checked its site.

The head of operations asked an assistant for scheduling tools that cut drive time. Fieldline was named, with its pilot as the reason. She landed on a home page that opened with "The future of field service". The travel-time claim was four blocks down, with no figure beside it. Her finance director asked what it cost for seventy technicians. The pricing page said "Talk to sales", so he asked an assistant, which quoted an old reseller listing. The IT manager found integration logos and no limits, and wrote "unclear".

Three checkers, and the site failed each one. That contractor chose a rival. Fieldline rebuilt the three pages, and fixed two of the three checks within a quarter.

The full playbook. The Website Playbook, with the page audit scorecard, the sitemap from the map and the service page outline, is free at rocmode.io/playbooks.

Whose evidence this is. Microsoft Clarity sells website analytics. TrustRadius is a review site that sells vendors paid pricing pages.

IN SHORT • CHAPTER 21

WHAT'S TRUE

Your website is the place of record. Fewer people arrive, but those who do come to check price, proof and terms.

DO THIS QUARTER

Redraw the sitemap from your map. Rebuild your most important service page from the anatomy, and test it against the old one.


WATCH

The share of moments covered by a page, and visitors to a service page who reach pricing and take the next step.

CHAPTER 22

Search and AI *answers*

THIS BUILD MOVES



Search optimisation and answer optimisation are one job: be named, be recommended, and be right.

What changed

Search is usually run as a ranking business. An agency picks keywords, reports positions each month and counts the traffic. Lately a second agency, or a second line on the invoice, offers the same for AI answers under a new acronym.

The work is the same. Google says so: "There are no additional requirements to appear in AI Overviews or AI Mode, nor other special optimizations necessary" ([Google Search Central](#)). The entry ticket is ordinary search.

What has changed is the test. Clicks have stopped being the score. Ahrefs studied 300,000 keywords. It estimated an AI Overview cut the top page's click-through rate by 58%, and calls the result a correlation ([Ahrefs](#)). And one buyer question now becomes many: Google says its AI features may use "query fan-out", issuing related searches across subtopics. The page that wins answers the follow-ups too.

-58%

top page's click-through rate when an AI Overview appears

AHREFS

What it is for

The job has three tests. When a buyer asks without naming you, you're **named**: that's Found. You're the one **recommended**: that's Preferred. And what's said about you is **right**, your price, terms and claims repeated accurately and backed by others: that's Believed.

It also serves Answered in part, because both kinds of answer draw on pages that settle the question.

Named isn't recommended. Picture five names in an answer, with a rival called "best for teams like yours". That's a loss that looks like a win. Most AI-visibility

reports count the mention and ignore the verdict.

This build doesn't create Known. A large recent preprint found Google search interest was the strongest predictor of how often AI named a brand (Chapter 13). That interest is made by memory work. This build harvests memory; it rarely creates it. Nor does it close the sale at Bought. Where the buyer lands is the website's job.

The anatomy of a good one

A good programme is a loop. It starts with what buyers ask, ends with what the machines said, then fixes the sources behind the answer.

A good search-and-answer programme wins on other people's pages, then fixes the source, not the answer




FIGURE 22.1

The buying questions, per moment (Found). Start from your map of moments and the question bank in Chapter 14, not a keyword tool. Write the questions in buyers' own long, specific words.

Answer pages, in plain text (Answered). Each buying question gets one owning page, with the follow-ups the engine will ask on the same page. Keep it in text, not in an image or a gated PDF.

Technical access (Found). Crawlable, fast and consistent. It's a gate, not a lever. Cloudflare, which says it handles traffic for 20% of the web, began blocking AI crawlers by default in July 2025 ([Cloudflare](#)). AI firms run separate crawlers for training and for search, and blocking one doesn't block the other.

Other people's pages (Believed). Citations, reviews and lists. The machine weighs what others say about you above what you say. Each source is small and the list shifts, so rebuild it each quarter from your own questions.

Sampled presence (Found · Preferred). Ask the same unbranded questions many times on the assistants your buyers use. Record named, recommended and right, and which sources each answer cites.

Fix the source, then ask again (Believed). Never chase the answer. Fix the page it came from: correct it if it's yours, earn a place or a correction if it's someone else's.

The mistake most teams make

They buy tricks. There's no special markup: Google says you don't need "new machine readable files, AI text files, or markup". Manipulating answers is now spam. Google's spam policies cover "attempting to manipulate generative AI responses in Google Search" ([Google](#)).

Hidden instructions are treated as an attack. Microsoft's security researchers found more than 50 such prompts from 31 companies over 60 days, hidden in "Summarize with AI" buttons ([Microsoft](#)). Microsoft calls it "AI Recommendation Poisoning".

The quieter mistake is buying search and AI separately. Same questions, pages and sources. Two agencies means two reports and nobody owning the gap. And most teams keep publishing on their own site while the verdict is written on other people's pages.

How you know it's working

Report three rates per moment, per assistant, with their margins: named, recommended and right. Then ask new enquiries how they found you, and log the moment. Ask buyers, because Google counts AI Overview and AI Mode traffic inside Search Console's ordinary "Web" figures ([Google](#)). Keep rankings and traffic as diagnostics, never goals.

Fieldline's search-and-answer programme

Fieldline, the invented scheduling-software firm, had a search agency and a separate AI-visibility retainer. Both reported green, because both tracked the operations question about drive time.

Then a deal stalled. The buyer's IT manager had asked an assistant which scheduling tools connect to their accounting software. Two rivals were named. Fieldline wasn't.

The team mapped the sources behind her question. Its integration guides were PDFs behind a sign-up form. A security review had blocked "all AI bots" at the CDN, including the search crawlers. And a review site's listing showed an integration Fieldline had retired.

The fixes were dull: publish the guides as text, allow the search crawlers, correct the listing. Two months later the IT question named Fieldline in about a third of answers. The stalled deal didn't wait that long; it went to a rival.

The full playbook. The Search and AI Answers Playbook, with the search-and-answer audit scorecard, the page map and the citation source map, is free at rocmode.io/playbooks.

Whose evidence this is. Ahrefs sells SEO and AI-visibility tools. Microsoft sells security products. Google, Microsoft and Cloudflare describe their own systems. No peer-reviewed source is used in this chapter beyond the preprint in Chapter 13.

IN SHORT • CHAPTER 22

WHAT'S TRUE

Search and AI answers are one job: be named, be recommended, and be right. Found is a rate, not a rank.

DO THIS QUARTER

Give every buying question one owning page. Build the citation source map and work your top five sources.


WATCH

Named, recommended and right rates per moment, per assistant, sampled monthly and reported side by side with their margins.

CHAPTER 23

Advertising

THIS BUILD MOVES



An ad has one job: to tie your name to a moment in the buyer's memory.

What changed

Advertising is the build most firms judge by what happens next. Clicks, cost per lead, return on ad spend, sales the dashboard says the ad caused. Each one is counted within days of the ad running.

That habit rewards the wrong ads. Most people who see an ad aren't buying, and won't be for months. Dawes's rule of thumb, that about 95% of business buyers aren't buying at any one time, is in Chapter 2. An ad judged on this week's sales gets bent towards the few who were about to buy anyway.

Those are the buyers you were most likely to get without it. Switching off eBay's paid search in 68 regions made no measurable difference to sales ([Blake, Nosko and Tadelis](#)). So this chapter treats an ad as a memory, written now and collected later.

What it is for

For nearly everyone who sees it, an ad's work is to leave your name filed under their moment. That's Known.

Its second, smaller job is to put you in front of someone at the moment they look: a paid search ad, a top marketplace listing. That's Found, bought rather than earned.

I think of these as **writing** the memory and **collecting** on it. Brand advertising writes. Search ads, retargeting and offers collect. Les Binet and Peter Field studied 996 campaigns from the IPA's effectiveness databank. Their rule of thumb puts about 60% of the budget into brand building and 40% into activation

(IPA). Treat the split as a starting point. What matters is that writing gets a large share, because collecting can't work on a memory nobody wrote.

A good ad also carries your reason, which helps at Preferred. But advertising rarely builds belief. A claim isn't more true because you paid to repeat it. Answered, Bought and Kept belong to other builds.

The anatomy of a good one

A good ad has two halves. The creative half decides what gets written into memory. The media half decides how many memories, and how often.

A good ad writes one memory, then buys enough reach to make it last




FIGURE 23.1

The moment, in the first seconds (Known). Start with the situation, not the product: the hungry afternoon, the tax letter, the wage sheet. Memory is filed under these moments.

The brand, early and distinctive (Known). One study pooled 50,000 viewers' responses to 25 television ads. The later the brand appeared, the less often viewers recognised it (Gerber and colleagues). Show a distinctive asset early, then again through the ad.

A feeling (Known · Preferred). Over three years or more, 43% of emotional campaigns in the IPA databank reported very large profit effects. Only 23% of rational ones did (IPA). Pick a feeling that belongs to the moment.

One thing to remember (Preferred). One, not three. Often it's your reason, in words someone else could repeat. "Every Monday by nine" works because it's easy to keep in mind, not because no rival could do it.

43% vs 23%

emotional vs rational campaigns reporting very large profit effects

IPA

Proof, then where to find you (Believed · Found). Point at evidence others have given you, then close on a name plain enough to type or say to an assistant.

Broad reach (Known). Growth comes from people who rarely or never buy you, so reach them too, not just last week's visitors.

Continuity and share of voice (Known). Memory fades. Stay on air at a steady level all year, not in one burst.

The mistake most teams make

They trust the attributed sales. Brett Gordon and colleagues compared common measurement methods against 15 randomised experiments at Facebook ([Gordon and colleagues](#)). The common methods "often fail to accurately measure the true effect of advertising", mostly overstating it.

The ads closest to the sale look best on the dashboard, and are the most likely to be paying for buyers who were coming anyway. But don't swing to the other extreme. Andrey Simonov, Chris Nosko and Justin Rao found that where rivals bid on a brand's name and the brand stayed off, the rivals took 18% to 42% of the clicks ([Simonov, Nosko and Rao](#)). Check who else bids before you switch branded search off.

None of this means collecting ads fail. It means you can't tell from the report.

How you know it's working

Use what I call the measurement ladder. Clicks and reach sit on the diagnostic rung: they show the ads ran. Steer by the leading rung, starting with searches for your name: Binet found share of search sometimes ran a year ahead of market share ([IPA](#)). Decide budgets on the causal rung: a holdout, where you switch ads off in some regions or periods and compare.

The bookkeeper's advertising

The bookkeeper for cafés is an invented example. Her budget is small. She made one ad: Monday, 7am, a café owner opening up and wondering if last week's wages ate the takings. Her name and her one colour appeared in the first two seconds and again at the end. The line was "Your wage cost against takings, every Monday by nine".

She ran it in two of her four towns for eight weeks, and not in the other two. She counted enquiries and searches for her name in each.

Nothing moved. With four towns, a small effect could hide in the noise. But she couldn't afford to keep paying on faith, so she stopped the ads and moved the

money to reviews. She plans to retest in a year. A test that says no is worth as much as one that says yes.

The full playbook. The Advertising Playbook, with the ad audit scorecard, the ad brief, the holdout test plan and the measurement ladder, is free at rocmode.io/playbooks.

Whose evidence this is. The IPA is a trade body for agencies. Two of the Facebook measurement authors worked at Facebook, which sells ads. Peer-reviewed sources: Gerber and colleagues; Simonov, Nosko and Rao; Gordon and colleagues; Blake, Nosko and Tadelis.

IN SHORT · CHAPTER 23

WHAT'S TRUE

An ad's job is to tie your name to a moment in the buyer's memory, written now and collected later.

DO THIS QUARTER

Check who bids on your name, then run a holdout on branded search or retargeting. Move any waste into one moment ad.

WATCH

Searches for your name, moment recall among new customers, and holdout lift. Not clicks or attributed return.


CHAPTER 24

Social

THIS BUILD MOVES

SEEN

SELECTED



MAINLY MOVES

MOVES IN PART

Social is proof of life, memory among the people who follow you, and reach through people. It is no longer a free distribution channel.

What changed

Most firms run social as a free broadcast channel. Post often, grow the following, chase engagement, and report followers, likes and reach each month.

The free broadcast has largely gone, by design. In January 2018 Adam Mosseri, then Facebook's head of News Feed, announced a shift towards posts from friends and family. "Pages may see their reach, video watch time and referral traffic decrease," he wrote ([Meta](#)).

Data on what's left shows how little. Socialinsider measured 25 million posts from 130,683 Facebook business pages over 2024 and 2025 ([Socialinsider](#)). The average engagement rate was 0.15%, flat year on year. Link posts did worst, at 0.05%.

A team that still treats social as distribution is paying staff to shout into a room that has emptied.

What it is for

Social has three real jobs once the goal is being selected.

Proof of life is Believed. Buyers check a feed the way they check a website: are you real, recent, and used by people like them? A feed whose last post is from March says something, and it isn't good.

Memory among the people who follow you is Known, but only among people you already reach.

0.15%

average engagement
rate on Facebook
business pages

SOCIALINSIDER

Reach through people is the one organic route to new buyers left. A founder, an employee or a customer reaches people your company page never will, and is believed more.

Video adds Answered and Found: it shows the thing, and machines transcribe it. Social does little for Preferred, Bought or Kept. Nobody decides on a supplier in a scroll. Social's job is to have put you in the room. Reach to new buyers at scale is paid, and that's advertising (Chapter 23).

The anatomy of a good one

Every post needs the first two parts, and at least one of the next three. The feed as a whole does the rest.

Every post cues a moment and looks like you, then does one more job




FIGURE 24.1

A moment the buyer recognises (Known). "Wages ate the takings again this month" stops a café owner. "We're excited to share" stops nobody.

A distinctive look, the same every time (Known). The same colour, framing, opening line and face, so a scroller knows it's you before reading a word.

Something useful, or something that delights (Answered). Leave the viewer knowing one thing they didn't, or answer one real question. A joke that lands in your moment, in your look, is memory work too.

A real person (Believed). Anatoli Colicev and colleagues tracked 45 brands in 21 sectors, day by day (Colicev and colleagues). The brand's own posts raised awareness, but not purchase intent. What customers said raised both. Your own posts build memory. Other people's words build intent.

The reason, in passing (Preferred). "That's why we send the report every Monday by nine." Said once and lightly, across fifty posts.

Spoken and captioned (Found). Ahrefs found YouTube mentions correlated with AI visibility at about 0.737 for ChatGPT, the strongest factor in its study ([Ahrefs](#)). Its authors warn that "correlation isn't causation". Say the moment, the question and your name aloud.

The feed as a whole: proof of life (Believed). One brilliant video from last year answers worse than a plain post from last week.

The mistake most teams make

They post in the brand's voice alone, and buyers discount a brand praising itself. A café owner believes another café owner far sooner than a logo.

The other mistake is chasing the viral hit. Sharad Goel, Duncan Watts and colleagues studied a billion diffusion events on Twitter ([Goel and colleagues](#)). Moderately popular events came up about one in a thousand times. "Viral hits" came up closer to one in a million.

So post on a schedule you can keep for a year. Consistency beats volume, and both beat the viral hit. And let staff show their own work, saying who they work for. Ten staff posting the same paragraph tells the audience who wrote it.

How you know it's working

Use the measurement ladder from Chapter 23. Followers, likes and reach stay on the diagnostic rung. Steer by the feed audit score each month, and by the share of posts that feature a customer, an employee or the founder. Then ask new customers whether they saw you on social or checked your feed.

Fieldline's social

Fieldline, the invented scheduling-software firm, sells to a committee. For two years its company page posted product updates and award logos to 4,000 followers, mostly staff and partners.

The founder began posting weekly, and two engineers monthly. Customers appeared in a video series that opened the same way: a van, a map, and a dispatcher saying how long the drive was.

In one invented deal, a contractor's head of operations had seen the van videos for months, shared by a dispatcher she followed. When the problem became urgent, Fieldline was the first name she thought of. The finance lead checked the feed and saw last week's post about a customer's finance manager. The IT manager's assistant cited an engineer's video.

Fieldline still lost to a cheaper rival. But it made the shortlist, which it had never done from social before. None of it showed in the reach report.

The full playbook. The Social Playbook, with the feed audit scorecard, the monthly content rhythm and the one-question video brief, is free at rocmode.io/playbooks.

Whose evidence this is. Ahrefs sells SEO tools and an AI-visibility tracker; Socialinsider sells social analytics. Peer-reviewed sources: Colicev and colleagues; Goel and colleagues.

IN SHORT • CHAPTER 24

WHAT'S TRUE

Social is proof of life, memory among your followers, and reach through people. It's no longer free distribution.

DO THIS QUARTER

Choose three moments, fix the look on one page, film one customer video per moment, and get the founder posting weekly.


WATCH

The monthly feed audit score, the share of posts featuring real people, and searches for your name. Not likes or reach.

CHAPTER 25

Creators

THIS BUILD MOVES



A creator sells belief you can't make yourself, so you're renting it, and the rules of belief apply.

What changed

Creators are now a line in most marketing budgets. In the US, brands' direct spend with creators rose from US\$13.9 billion in 2021 to US\$29.5 billion in 2024 (IAB). The IAB, the US advertising industry body, projected US\$37 billion for 2025.

Most of that money is bought like media. Creators are chosen by follower count, handed a script and judged on reach and engagement. The creator becomes a billboard with a face.

That misses what's being bought. Buyers believe a creator because the creator has been right before, about things they care about. That belief is the creator's, not yours. You can rent it for a while. You can't own it, and a script spends it.

What it is for

A creator's real job is to make a claim believable to people who wouldn't believe it from you. That's the Believed stage, paid for.

What a brand says splits into facts and judgements. Buyers accept facts from you, but not judgements from the party who gains from them. "It's easy to use" from you is an advertisement; from someone they've watched for two years, it's a recommendation.

So a creator does three jobs, in this order. First, vouch (Believed): would the audience trust this person on this subject, unpaid? Second, show the reason (Preferred): does the post show why you, in a way someone could repeat? Third, put you in a moment (Known): do the right buyers now link your name to their situation?

Creators don't answer the buyer's detailed questions; your website does. Most creator work doesn't close the sale either. The exception is affiliate and shop selling, where the creator is paid a commission on each sale. A commission rewards closing, not belief, so watch accuracy closely, and test before you scale, because codes credit sales that may have happened anyway.

The anatomy of a good one

A partnership is a programme, not a post. Choose the audience once, get four things right in each piece, then repeat over months and review.

A good creator partnership rents belief: choose the audience, get each piece right, then repeat




FIGURE 25.1

The right audience [Known]. Start with who follows the creator, not how many. Simone Wies, Alexander Bleier and Alexander Edeling studied 802 Instagram campaigns and found engagement with sponsored posts rose with follower count, then fell (Wies, Bleier and Edeling). A larger following signals a weaker tie, so each follower is less moved. Choose on fit first and size second.

Clear disclosure, first [Believed]. Say it's paid, at the start, in words anyone understands. In Australia, the ACCC reviewed 118 influencers in 2023 and found 81% made posts that raised concerns under consumer law (ACCC). When an audience finds out later, both of you lose the belief.

The creator's own words [Believed]. Their voice, their experience, their caveats. Wies and her colleagues found that original content weakened the

81%

of influencers reviewed made posts raising consumer-law concerns

ACCC

drop-off at large follower counts.

Something true to show [Believed · Preferred]. A product to use, proof to point at, your reason. A number the creator found themselves is a fact, and it travels.

A path to check you [Answered · Bought]. The buyer will search your name or ask an assistant. What they find must match what the creator said: the same price, claim and proof.

A series, reviewed [Known]. One post is an ad. The same creator returning over six months is a relationship the audience can watch. Review after three pieces, and stop a series that isn't working.

The mistake most teams make

They write the script. The fee buys access to an audience's trust. The script turns that trust into an advertisement, so you pay the creator's rate and get a banner's effect. The creator bears the rest of the cost after you leave. That's my reading of the research, not a measured result.

So give the creator the facts that must be right, and leave the words to them. Say in the brief that they may show what didn't work, then don't penalise them for it.

Disclosure does lower persuasion. Martin Eisend and colleagues pooled 61 papers and found it "reduced brand attitudes, credibility, and source evaluation" (Eisend and colleagues). The rules require it anyway. My working rule: disclose early and plainly, and win the belief back with what you show.

How you know it's working

Keep reach, follower counts and engagement on the diagnostic rung. Judge each creator on leading measures: audience fit, from their data or a sample of 50 commenters; the share of comments that ask the buying question; new customers who name the creator when asked what was happening when they first thought of you; branded search after each post; and the share of posts whose facts matched your one truth exactly.

An example

The bookkeeper for cafés, an invented example, scored three creators. A food reviewer with 180,000 followers reached diners, who don't hire bookkeepers. A café owner with 9,000 followers posted about running costs, and 31 of 50 sampled commenters ran a café.

She paid the café owner for four posts over six months. The owner used the Monday report in her own café and posted her wage cost each month, in her

own words. In one month it went up. She said so, and said why. Each post opened with “Paid partnership” and linked to a page with the same promise. One café owner from the comments became a client and later wrote a review, which a machine can read.

The full playbook. The Creators Playbook, with the creator fit scorecard, the creator brief and the disclosure checklist, is free at rocmode.io/playbooks.

Whose evidence this is. The IAB is the US advertising industry body, and its spend figures rest on an Advertiser Perceptions survey whose sample isn’t published. The ACCC is a regulator. Peer-reviewed sources: Wies, Bleier and Edeling; Eisend and colleagues. The large academic studies measure engagement, not sales.

IN SHORT • CHAPTER 25

WHAT’S TRUE

A creator sells belief you can't make yourself.
You're renting it, so the rules of belief apply.

DO THIS QUARTER

Ask ten customers whose voices they trust, score the shortlist for fit, and brief one creator for a three-month, disclosed series.

WATCH

Audience fit, comments that ask the buying question, new customers who name the creator, and whether the facts were right.


CHAPTER 26

PR and reviews

THIS BUILD MOVES

SEEN

SELECTED



PR's job is no longer coverage. It is getting into the sources the machine quotes, with something only you have.

What changed

PR is usually judged by its clippings: how many pieces ran, how many people they reached, and what the same space would have cost as advertising. That last figure is the advertising value equivalent, or AVE. The industry's own measurement body gave up on it long ago. AMEC's 2020 Barcelona Principles are blunt: "AVEs are not the value of communication" ([AMEC](#)). In my experience, it still turns up in reports.

Reviews get run as a score to protect. Someone watches the star rating, chases a few happy customers when it dips, and leaves the bad reviews unanswered.

Neither habit asks what the buyer, or now the machine, read at the moment of checking.

What it is for

PR and reviews have one job once the goal is being selected. They put other people's words about you into the places buyers and machines check.

That makes this the main build for **Believed**. Judgements about you are only accepted from others, and this build is how you earn them.

It serves three more stages. It builds **Known** in the machine's memory: news mentions went with AI prominence in Chapter 13's exploratory study. It serves **Found**, because the sources cited for a buying question decide who gets named. And it supplies the reason at **Preferred**, in someone else's voice.

It depends on **Kept**. Reviews come from customers whose promises were kept. PR can't fix a product that lets people down; it'll only spread the news faster. And it doesn't do the work of **Answered** or **Bought**. Your facts, prices and terms belong on your own site.

The anatomy of a good one

It runs as two tracks on their own clocks. Reviews turn every week; PR turns once a quarter. For most small firms, reviews are the Monday task, so start there.

Two tracks on two clocks, both feeding the sources the machine cites




FIGURE 26.1

Reviews, asked at the right moment [Believed]. Ask when the job is done, at the first win. Never gate: the US Federal Trade Commission says, "Don't ask for reviews only from customers you think will leave positive ones" (FTC).

A public reply to every review [Believed · Kept]. The next buyer reads it. Davide Proserpio and Georgios Zervas found that when hotels began responding, their ratings rose by 0.12 stars on average, and they received 12% more reviews (Proserpio and Zervas). Hotels got fewer but longer negative reviews, once unhappy guests knew someone would read them.

The questions buyers ask, and the sources cited [Found]. Ask your buying questions in the assistants your buyers use, many times each, and read the citations. Those are your target list. A trade newsletter may matter more than a national paper.

+12%

more reviews once hotels began replying to reviews

PROSERPIO AND ZERVAS

Something only you have [Preferred]. Proprietary data, a finding or a clear point of view. Cision surveyed 1,899 journalists, who named original research, expert access and embargoed information as what they value ([Cision](#)). The finding should carry your reason, so it travels whenever the number is quoted.

The right outlets [Believed]. Take the finding to the sources your map turned up, starting with those cited most.

Ask again [Found]. Rerun the same questions each quarter and check whether you now appear.

The mistake most teams make

They pitch announcements to big lists: a new hire, a new office, an award. Nobody else needs to quote any of it. The same Cision survey found 72% of journalists say fewer than a quarter of pitches they receive are relevant.

The second mistake is on the reviews side: leaving bad reviews unanswered, or filtering them. That's now law, not just manners. The FTC's rule on fake reviews took effect on 21 October 2024 ([FTC](#)), and in the UK hiding negative reviews is banned too ([CMA](#)). Answer the bad ones first, because the next buyer reads them first. Own what was wrong and say what changed.

And don't write your own Wikipedia entry. Wikipedia is among the sources assistants cite most, and paid advocates are "very strongly discouraged" from editing their own articles ([Wikipedia](#)). Earn the independent coverage its editors can cite, and let them write it.

How you know it's working

Retire coverage counts, AVE and the star rating as goals. Steer by the share of your buying questions where you appear in at least one cited source; the number of top-cited sources you're in, quarter on quarter; whether your finding is quoted with your reason attached; reviews in the last 90 days; and the share of reviews answered, with the time taken on the negative ones.

An example

Fieldline, the invented scheduling-software firm, found a G2 review saying the travel savings "never showed up after the pilot". Nobody had replied. The customer team answered that week: the customer hadn't switched on route grouping, here were the drive-time figures since the fix, and here was a named contact. A month later, the finance question returned the review and the reply together.

The PR track went less well. The trade magazine ran Fieldline's drive-time finding as a feature, but two quarters on the assistants still cited an industry

survey instead. The weekly track is faster. The quarterly one may yet fail.

The full playbook. The PR and Reviews Playbook, with the source map and PR target list, the data-story pitch and the review response guide, is free at rocmode.io/playbooks.

Whose evidence this is. Cision sells PR and media software. AMEC is the measurement industry's own body. The FTC and CMA are regulators. Peer-reviewed: Proserpio and Zervas, whose study used hotels, so other categories are untested.

IN SHORT · CHAPTER 26

WHAT'S TRUE

PR and reviews put other people's words about you where buyers and machines check. Coverage counts aren't the job.

DO THIS QUARTER

Remove any review gating, answer every review, build the source map, and publish one finding only you could produce.


WATCH

The share of buying questions where you appear in a cited source, recent reviews, and how fast you answer bad ones.

CHAPTER 27

The *pitch*

THIS BUILD MOVES



A pitch is written for the meeting you're not in.

What changed

Most pitch decks are written for the meeting the seller attends. They open with the company's story, a wall of client logos and a tour of features. The presenter fills the gaps and handles questions as they come.

Then the meeting ends and the real decision starts. The person who liked you, your champion, has to make your case to finance, to IT and to the boss. You're not there. Neither is the voice-over that made slide nine make sense.

And now there's a fourth reader. Buyers check what a rep says against AI, and what AI says against a rep. So the better question isn't whether the meeting went well. It's whether the champion could make your case without you, and whether it would survive a finance lead checking it with an assistant.

What it is for

A pitch gives people who never met you a reason to say yes, and nothing to fear. I call the place where that happens **the second meeting**: the one your champion runs, without you, a week after yours.

It serves three stages directly. **Preferred**, because it carries your reason into a room you're not in. **Believed**, because it brings proof from someone like them. **Bought**, above all, because it names the cost of doing nothing and takes the risk out of acting.

It doesn't serve Known or Found. A pitch only reaches buyers who already have you on the list. It's mostly a confirming job, too. Buyers bought from their early favourite 77% of the time in 6sense's survey. So the pitch rarely wins a deal from nowhere. It confirms a lead or loses it.

77%

of buyers bought from their early favourite

6SENSE

Every member of the committee can say no. Three CEB researchers, writing in HBR, described groups “all with different roles and priorities, and all with veto power” (HBR). The champion has to teach your case to them, and the pitch is the material they teach from.

The anatomy of a good one

A good pitch is a case a stranger can make. The order holds whether you use seven slides, a memo or seven minutes of talk.

A good pitch is a case the champion can make without you, and it starts with the cost of waiting




FIGURE 27.1

Their moment, in their words [Known · Answered]. “Our technicians spend half the day driving” beats any headline you could write.

The cost of doing nothing [Bought]. Doing nothing is often your real rival. Put a figure on carrying on as before, in their numbers, worked out with the champion. Use a figure finance would sign; an inflated one dies in the second meeting.

What good looks like [Preferred]. Teach before you sell. CEB’s researchers found the best suppliers created “shared language and perspectives around problems and solutions” (HBR).

Your reason, which a rival wouldn’t say [Preferred]. One reason that passes the opposite test (Chapter 11), and who it doesn’t suit.

Proof from someone like them [Believed]. One customer of their size, with a number and a name they could call. Use exactly the figure your website uses.

The risk removed [Bought]. A small first step, a way out and someone to call, on the slide, not only in the contract.

A plain recommendation [Bought]. Name one option and say why, then a next step with a date and an owner.

The mistake most teams make

They write the deck about themselves, for the first meeting. The founding story, the logo wall, the feature tour, three tiers with no advice. Nobody in the second meeting asks for any of it.

Material the reader doesn't need does harm, not just no good. Richard Mayer reviewed experiments on learning and found people learned more deeply when extraneous material was left out, in 10 of 11 tests ([Mayer](#)). Those are lab studies of learning, not sales. But the champion's task is a learning task. Every slide they must skip or explain is a cost. And a menu of three tiers hands the hard choice to the person who isn't in the room.

My working rule, which no sales study I know of tests: send the pitch first, and use the meeting for questions.

How you know it's working

Retire meetings booked, deck views and "great meeting". Watch the no-decision rate on deals that reached a pitch; the forward test, meaning how often the champion's forwarded version keeps your reason and numbers intact; new questions per meeting, falling over time; and machine match, the share of your claims that assistants state the same way.

An example

I learned this one the hard way. In 2023 we pitched a new website to a services firm with offices in four cities, with a polished online proposal: our process, our guarantee, our price. We lost. When I asked for candid feedback, the head of marketing gave it: "the successful proposal was more in depth and personalised to the project. It included some initial design ideas and site structure suggestions which were based on the brief. It gave me more confidence they were on the right track."

The winner hadn't out-credentialled us. It had shown its thinking about her problem, and that's something a champion can carry into a room.

The full playbook. The Pitch Playbook, with the pitch audit scorecard, the seven-slide pitch skeleton and the one-page champion summary, is free at rocmode.io/playbooks.

Whose evidence this is. 6sense sells intent data. CEB, now part of Gartner, sold sales research and advice, and Gartner sells research and advice. Academic: Mayer's review of lab experiments on learning. No study I found tests pitch structure against win rates; the anatomy is my inference.

IN SHORT • CHAPTER 27

WHAT'S TRUE

A pitch is written for the meeting you're not in. Your champion has to make the case without you.

DO THIS QUARTER

Rebuild the master pitch in the seven-part order, run the re-tell test and machine check, and send it before every meeting.

WATCH


The no-decision rate on pitched deals, and whether the champion's forwarded version keeps your reason and numbers intact.

CHAPTER 28






Email

THIS BUILD MOVES

SEEN



SELECTED



MAINLY MOVES

MOVES IN PART

Every email should answer a question the reader has this week.

What changed

Most email is sent on the sender's calendar. The newsletter goes out on the first Tuesday, the sale blast when stock needs to move, the "just checking in" because the CRM says it's been a fortnight. Then it's judged by opens, clicks and list size.

Two things have changed. The open rate stopped meaning anything: Apple's Mail Privacy Protection loads images in the background "regardless of whether you engage with the email" ([Apple](#)). And a machine now often reads your email before the person does. Gmail "synthesizes the entire conversation into a concise summary of key points" ([Google](#)), and promotes "high-stakes items — like a bill due tomorrow". Both Apple and Google promote the urgent, the personal and the expected. A monthly blast from a sender you never reply to is what they push down.

What it is for

Email's job is to reach someone who has already raised a hand, at the moment they have a question, and answer it.

It serves **Answered**, because each email exists to answer one question. It serves **Believed**, because it carries proof to a person weighing you up. It serves **Bought**, because the gone-quiet email is where the risk gets lowered. And it serves **Kept**, because the first week and the run-up to renewal are both moments email can own. It does little for Known or Found, because you can't email people who haven't given you their address.

**5.3% →
41%**

triggered emails' share
of volume, and their
share of revenue

KLAVIYO

The best evidence that timing beats volume is vendor data. Klaviyo's benchmarks, drawn from over 183,000 of its customers, show flows, the emails set off by something a customer did, made up 5.3% of email volume ([Klaviyo](#)). Those flows produced nearly 41% of email revenue. Part of the gap is who receives them, since many go to people halfway through a purchase. Even so, an email sent because something happened does far more work than one sent because it was Tuesday.

The anatomy of a good one

A good email is a reply to a question the reader hasn't quite asked yet.

A good email starts with the buyer's moment, not your calendar




FIGURE 28.1

Triggered by a moment, not a date [Answered]. An enquiry, a quote unanswered, a first login, ninety days to renewal. If you can't name the trigger, you're writing a broadcast.

A subject line that is the answer [Answered]. "Your quote: \$4,200 a year, with a 30-day exit" beats "Following up on our chat". It's also the line most likely to survive a machine's summary.

The point in the first two lines [Answered]. No "I hope this finds you well". The first two lines are what the preview shows and what a summary keeps.

One piece of proof [Believed]. A number from a customer like them. One strong piece is believed; five read as a brochure.

What lowers the risk [Bought]. State the way out near the price: the trial, the exit, the guarantee. Never under a deadline.

One next step [Bought]. One link, one reply or one booking: the smallest yes that moves them forward.

A person's name, and a reply that lands [Believed · Kept]. A "noreply" address tells the reader you want to talk, not listen. Google lists relationships among its AI Inbox's signals.

The mistake most teams make

They keep the calendar and manage to opens. Litmus estimates Apple's feature affects "roughly 55-60% of all email opens", which it calls "not considered reliable" ([Litmus](#)). So don't set targets on opens or judge subject lines by them.

The calendar send isn't worthless. Campaigns still brought in about 59% of email revenue in Klaviyo's figures, and some do a memory job for light and lapsed buyers. My working rule is "kill or justify": a calendar send stays only if it can name the memory job it does, and for whom.

And watch what a service email sells. In 2026 the ACMA penalised Lululemon \$702,900 for order and delivery emails that carried promotional content with no unsubscribe ([ACMA](#)).

How you know it's working

Measure each email by what happens after its moment: time from trigger to send; the moment's next step, such as enquiries that book a call; replies, and what they ask, since a question in a reply is a gap in the email; the no-decision rate after quotes; renewals from customers who got the first-week emails; and spam complaints, kept well under Gmail's 0.10%.

An example

Fieldline, the invented scheduling-software firm, sent one long recap to a committee after a good demo, with the subject "Great to meet you all!" The finance lead's inbox summary said only that Fieldline had thanked them. The IT lead's question, about security review, was answered in paragraph six.

The account manager rewrote it as three emails, each the answer to one person's question. The IT lead forwarded hers to the security team the same afternoon. Fieldline still lost, to a rival a fifth cheaper. But it learned why in two weeks, in the buyer's own words, rather than months of silence.

The full playbook. The Email Playbook, with the email programme audit, the moment-to-email map and three moment emails, is free at rocmode.io/playbooks.

Whose evidence this is. Klaviyo sells email and text-message automation, and its benchmarks come from its own customers, mostly online retailers. Litmus sells email testing and analytics tools. Apple and Google describe their own products, in their own words. The ACMA is a regulator. No peer-reviewed source in this chapter compares triggered and scheduled email, and nobody has yet measured what AI summaries do to marketing email.

IN SHORT • CHAPTER 28

WHAT'S TRUE

Every email should answer a question the reader has this week. Send on the buyer's moments, not your calendar.

DO THIS QUARTER

Map one email per moment, from enquiry to renewal, each written to the anatomy with a named owner. Stop reporting opens.


WATCH

Each moment's next step, the replies and what they ask, and the no-decision rate after quotes. Not open rates.

CHAPTER 29

Automation

THIS BUILD MOVES



Automate the noticing, not the talking.

What changed

Marketing automation is usually built as workflows. Someone fills in a form, and five emails follow. Points pile up against a contact until a score tips them over to sales. The work is judged by how many emails go out and how many leads cross the line.

That misses what automation is good at. A machine never forgets to look. It can see a quote viewed three times and not answered, usage halved, or a renewal sixty days away. Most businesses miss those moments, because nobody's watching.

What's new is that an AI assistant can now hold the conversation too, and many firms are handing it the front door. So the question of who answers, a machine or a person, has become the one most teams get wrong.

What it is for

Automation has one job: to notice a moment while it's live, and get the right answer to the buyer in time. The answer may come from a machine or a person.

It serves three stages most directly. **Answered**, because each moment carries a question, and automation makes sure it gets a reply. **Bought**, because many sales are lost to delay and silence, not to rivals. And **Kept**, because falling usage and an approaching renewal are the moments where customers decide whether to stay.

It touches **Believed**, because a wrong answer from your bot is your wrong answer. A Canadian tribunal held Air Canada to what its chatbot told a customer (Chapter 16).

7x

more likely to qualify a lead when contacted within an hour

OLDROYD, MCELHERAN
AND ELKINGTON (HBR)

It does little for Known, Found or Preferred. It works on people who have already raised a hand.

Speed matters because a moment is short. An audit published in the Harvard Business Review analysed 1.25 million sales leads ([Oldroyd, McElheran and Elkington](#)). Firms that tried to reach a lead within an hour were “nearly seven times as likely to qualify the lead” as firms that waited an hour more. It’s a correlation: fast firms may simply be better at selling.

The anatomy of a good one

A good automation runs left to right, from something the buyer did to a record of what happened next.

A good automation notices the moment, then decides who should answer




FIGURE 29.1

A signal: something the buyer did [Answered]. A click, a reply, a login, a quote viewed. Never trigger a person on an email open alone; Apple’s privacy feature can fire an open when nobody looked.

A rule: the moment it means [Answered]. “Pricing page twice within 48 hours, after an enquiry” means “this buyer is comparing prices”. Write the question down. If you can’t say what the buyer is asking, you’ve only got a click.

The route: machine or person? [Bought]. My working rule is to route on stakes for the buyer. Stakes are high when the money is large for them, the decision is hard to reverse, they're upset, or the answer depends on their circumstances.

A simple answer, by machine, from the one truth [Answered · Believed]. It doesn't improvise. If the answer isn't in the approved sources, it says so and offers a person.

High stakes: a named person, with context [Bought · Kept]. One name, not a shared inbox, with everything the machine knows, so the buyer never repeats themselves.

Record what happened, and learn [Kept]. Read the records monthly, and change the rules that aren't working.

The mistake most teams make

They send sales the highest lead score. Most scoring, as run, counts activity: everything a person has done, over months. A moment is a particular combination of actions, close together, that means a buyer is at a particular step. The finance lead who visits the pricing page twice in one afternoon may score 10 points. The junior analyst who has read every blog post for a year may score 140. The score sends sales to the analyst.

The second mistake is putting a machine between the customer and a person. Gartner found 64% of customers would prefer companies did not use AI for customer service, and the top concern was that it would be harder to reach a person ([Gartner](#)). That's a stated preference, but it names the fear precisely.

How you know it's working

Keep workflows, emails sent and "deflection" as diagnostics; a high deflection rate can mean a trapped customer. Judge automation by the share of each moment's signals answered within its time limit, the share that moves to the next step, how often a buyer has to repeat themselves, the assistant's accuracy against the one truth, and deals lost where a signal fired and nobody acted.

An example

Fieldline, the invented scheduling-software firm, used to run a lead score. In one deal, a junior coordinator downloaded three guides, crossed the line and got the sales call. She had no budget. Meanwhile the finance lead viewed the proposal four times in two days, going back to pricing each time. Nobody saw it, because viewing a proposal earned five points. The deal went quiet.

Fieldline rebuilt around moments. In the next deal, the finance lead's third view reached the account manager within the hour. Her first line was a question, not

a report: "Most finance leads at this stage want year one all-in. Would that help?" She never mentioned what she'd seen.

The full playbook. The Automation Playbook, with the signal-to-moment map, the handover rule sheet and the guardrails spec for a customer-facing AI assistant, is free at rocmode.io/playbooks.

Whose evidence this is. Gartner sells research to customer service leaders, and its customer figures are a stated preference from late 2023. The lead-response audit appeared in HBR, not in a peer-reviewed journal, and I found no peer-reviewed replication.

IN SHORT · CHAPTER 29

WHAT'S TRUE

Automate the noticing, not the talking. Let the machine spot the moment, and send high stakes to a named person.

DO THIS QUARTER

Build the signal-to-moment map for your most important moments, with a named person on every high-stakes route.

WATCH

The share of each moment's signals answered within its time limit, and how many then move to the next step.

INTERLUDE

The rest, *briefly*

Four builds didn't get a chapter of their own: outbound, events, partners and customers. The same questions apply to each. What's it for? Which stages does it serve? What's the method? How do you know it worked?

Outbound and account-based

What it is for. Reaching named firms before they start looking, so you're the firm they remember when they do.

Which stages. Known first, for a short list of firms you've chosen. Then Believed, if what you send is worth keeping.

Two findings pull against a cold pitch. Gartner surveyed 646 B2B buyers in August and September 2025, and 67% said they'd prefer to buy without a sales rep ([Gartner](#)). And John Dawes of the Ehrenberg-Bass Institute estimates that only about 5% of business buyers are in the market for a category at any one time ([Ehrenberg-Bass](#)). So most cold messages reach people who aren't buying, and who'd rather not be sold to when they are.

The method. Choose the firms first, from your segments, and name the moment each is most likely to be in: usually the trigger. Then send something worth keeping, about their problem rather than your service. An invented engineering consultancy writes to the owners of forty plants in its sectors with one page on the capacity most lines leave unused, and asks for nothing. When an owner replies, the pitch takes over.

How you know it worked. Replies and meetings from the named list. Then whether those firms turn up later by other routes: a search for your name, a visit to the site, a referral. The number of messages sent tells you nothing.

Events

What it is for. An event puts you in a room with buyers, and the people they trust. Buyers can judge you in person, before they have to decide anything.

Marketers rate them highly. The Content Marketing Institute surveyed 980 B2B marketers for its 2025 report. In-person events and webinars topped the channels said to produce the best results, at 52% and 51% ([CMI](#)). That's marketers' own opinion, but it's reason enough to run events properly, rather than to run more of them.

Which stages. Known and Believed, in person. Events matter most in owner-led and committee buying, where a face and a conversation carry weight.

The method. Go with moments in mind, not a stand and a bowl of mints. Before the event, list who'll be there, which moment each is in, what they need to believe and the next step. Bring proof from someone like them. Follow up within a day, with the answer to the question they asked you. At a trade forum, invented Fieldline skips the stand: a customer's operations manager gives a talk on travel time, and the team books second conversations with those who stayed to ask questions.

How you know it worked. Count the conversations that moved a moment forward: a meeting booked, a question answered, a pilot agreed. Badge scans tell you who walked past, not who moved.

Partners and channels

What it is for. A partnership borrows belief. A partner who already has the buyer's trust says you're good, and puts you where the buyer already is.

Which stages. Believed first, through the partner's word. Then Found, because partners list you in places buyers look, and machines read those places too. A partner directory is a structured page that names you and says what you do: the kind of page an AI assistant can retrieve. That's my reading, not a measured result.

The method. Choose partners whose customers share your moments, and make the listing say your reason in your one truth's words, with one piece of proof. The invented bookkeeper for cafés lists herself in her accounting software's directory of advisers. Xero's, for one, invites owners to "find experts based on your industry, location, and specific accounting needs" ([Xero](#)). So her listing says "cafés" in the first line, not "small business".

How you know it worked. Count customers who arrived through each partner, and whether they bought. Ask an AI assistant the partner's customers' question, and see whether the partner's page brings you up.

Customers: onboarding, loyalty, community and referral

What it is for. Making staying the easy, safe choice, and turning customers into the people who bring the next one.

Which stages. Kept first. Then, through what customers say and write, Believed and Known for the buyers who come next.

The method. Start with the handover. The first weeks decide whether the promise made in the pitch survives the work, so onboarding is the first moment of Kept, not an admin task. Then keep the promise. No programme repairs a

broken one. Points for their own sake are a discount with extra steps. Byron Sharp and Anne Sharp studied a large loyalty programme: only two of its six member brands showed substantial “excess loyalty”, and those gains showed up among non-members too ([Sharp and Sharp](#)).

Community does work points can't. Customers who help each other write the evidence that tomorrow's buyers read. So host the conversation, answer in it and never edit the criticism out.

How you know it worked. Renewal and repeat rates, first. Then the share of new customers who came by referral, and fresh customer-written reviews each month. Points redeemed and members signed up are diagnostics.

Whose evidence this is. The Content Marketing Institute's survey records marketers' opinions, sponsored by The MX Group. Xero describes its own directory. Gartner's figure is a stated preference from its own survey of buyers. The 95:5 rule is a rule of thumb from buying cycles, not a measurement of any one market. Peer-reviewed: Sharp and Sharp, in the *International Journal of Research in Marketing*.

And the nods

Some channels never need a build of their own, but each needs a home. Sponsorship sits with advertising. Analyst relations, awards and word of mouth sit with PR and reviews. Affiliates, marketplaces, associations and employee advocacy sit with partners and creators. Account management and customer service sit with customers. CRM, SMS and chat sit with email and automation. Packaging and physical spaces sit with the website, as the other place a buyer checks what you said. Each gets a short section in its build's playbook.

It's a lot

Read end to end, Part four is a lot: fourteen builds in six categories, all saying one thing at every moment that matters. That's what being selected takes. It's also why most brands do a little of everything and none of it well.

Nobody does all of it at once. Start with fifteen moments and one build. Pick the build that serves the stage where you lose the most buyers. Do it properly, then add the next.

The harder question is knowing where you're losing them. Part five answers it.

WHERE YOU ARE

01

Why it feels like this

02

What a choice is
made of

03

The way forward

04

The builds

05

Knowing it worked

05 / PART FIVE

Knowing it *worked*



The Selectability Score

A brand needs one number that says how likely it is to be selected, and one list that says what to fix first.

Everything else is detail beneath those two. This chapter sets out the score in full. It's the method behind RocMode's scoring, and I've written it so you could build your own.

Be clear about what the number is. It's an index of your odds at the moments that matter, not a count of sales. Nothing yet ties it to market share or revenue. That test is still to come, and this chapter says how it'll be run. The appendix turns it into a worksheet you can fill in by hand.

The unit is the moment

The score is built from the map in Chapter 6. Each moment is one kind of buyer at one step of their journey. Only moments that exist for your brand are scored. A moment you haven't defined is left out, never counted as zero.

Each moment is heard through several buyer questions: the different ways that buyer might ask at that step. Five kinds of buyer and seven journey steps make 35 moments. Three questions for each makes 105 questions, and still 35 moments. More questions make each reading steadier without adding moments.

One moment, heard three ways: more questions steady the reading, not the count




FIGURE 30.1 More questions steady each reading; they never add moments

Six checks in every moment

Each check asks one question, and each uses the tool from its stage chapter.

STAGE	THE QUESTION	HOW IT CAN BE MEASURED	THE TOOL
Known	Does the buyer, or the machine, already know you?	Two signals: whether AI assistants recognise you by name and tie you to the moment, and your share of search against named rivals	The moment memory map (Ch 13)
Found	Do you appear when they ask without naming you?	The share of answers to unbranded questions that name you	The presence sample (Ch 14)
Answered	Does what you publish answer what this buyer needs to know?	The share of their questions and objections your site answers	The question chain map (Ch 15)
Believed	Is what is said about you accurate, and backed by others?	Accuracy, outside support and consistency, in the answers that name you	The claim, proof and source inventory (Ch 16)
Preferred	When you appear, are you the one recommended, with a reason?	How strongly you are recommended, and whether your reason is the one given	The reason gap sheet (Ch 17)
Bought	Does the place they land make it easy to say yes?	Whether the page repeats the reason, shows proof and gives a clear next step. Your own sales data, where you have it	The risk-removal ladder and the decision friction audit (Ch 18)

The seventh stage, Kept, is scored differently. It comes after the formula.


The formula

Known and Found are gates. If a buyer has never heard of you and can't find you, nothing else you do is seen. The other four stages are the quality of what they find.

Gate = square root of (Known × Found)
 Quality = weighted average of Answered, Believed, Preferred and Bought
 Moment = Gate × Quality
 Score = weighted average of every moment, from 0 to 100

Work each check as a share from 0 to 1, and multiply by 100 at the end.

Each moment scores $\text{gate} \times \text{quality}$, so a gate of zero zeroes everything



Worked lines: the bookkeeper's moment A, with owner-led weights.

FIGURE 30.2 The score formula · worked for the bookkeeper's moment A

The multiplication is the point. A brand with perfect answers that nobody finds scores zero, because a gate of zero multiplies everything down to nothing.

The square root softens the gate. Take a brand with Known and quality both at 100, so only Found varies. Named in half the answers in every moment, it scores $\sqrt{(100 \times 50)} =$ about 71 in each, and about 71 overall. Now name it in every answer in half its moments, and in none of the rest. It scores 100 in half and 0 in the others: 50 overall.

So the score favours breadth. Being on more lists some of the time beats owning half of them. That matches the evidence on breadth in Chapter 13: brands linked to more buying situations win more buyers and keep more of them.

The same presence scores 71 spread evenly, 50 in lumps

20 moments each · Known and quality held at 100 · both grids carry the same amount of shading




FIGURE 30.3 The same total presence scores higher when it is spread across more moments

It's "seen isn't selected" written as arithmetic. Being seen is the gate, and it's worth nothing until it's multiplied by quality: the answers, the proof, the reason and the way to buy.

The weights

Two sets of weights turn the formula into a number. The first is between the four quality stages, and it depends on how your buyers buy. These are the starting weights for the five buying modes (Chapter 7).

MODE	ANSWERED	BELIEVED	PREFERRED	BOUGHT
Habitual	10%	20%	15%	55%
Quick choice	15%	25%	15%	45%
Considered personal	25%	30%	30%	15%
Owner-led business	20%	35%	30%	15%
Committee	25%	30%	30%	15%

In habitual and quick-choice buying, Known and Found already sit in the gate, so inside the quality score Bought carries them, and a hurried choice leaves no room to weigh reasons. In the other three modes the buyer stops to think. Answered, Believed and Preferred decide it.

The second set of weights is between moments. Moments nearer the decision count for more, and your main buyers count for more than occasional ones.

I should be plain about these numbers. They're judgement, not yet fitted to sales data. They're published, versioned and open to challenge. These are version 1.0.

How the weights will be tested

A score nobody has checked against sales is a hypothesis. Here's how I mean to check it.

Score the same brands every quarter, moment by moment, and set each score beside that brand's share of new customers or sales, in the same category and period. If the score has value, brands that score higher should tend to hold more share, and a rise in the score should come before a rise in share, not after it.

If it doesn't, the weights are wrong, and possibly the formula too. Every change gets a new version number and resets the baseline.

This needs data no single firm holds. If you'll share your scores beside your sales or share, I'd like to hear from you at rocmode.io/research. RocMode will also run this test on its own customers' scores, and the findings will be published, whatever they show.

Kept is a moment

A first-time buyer can't have been kept. So Kept isn't a seventh weight.

It's the same six checks asked at a different moment: renew or switch. "Is it worth renewing?" Those moments sit on the same map and are scored the same way. If you hold your own renewal or repeat-purchase data, that's the truth, and the score explains it.

Four kinds of moment

Two of the measures sort every moment into one of four kinds.

	THE MACHINE PICKS YOU	THE MACHINE DOES NOT
Your site answers the buyer	Won	Hidden
Your site does not	Borrowed	Empty

Hidden is the common one: you have the answer and nobody is sent to it. Borrowed is the fragile one: you're recommended on the strength of what others wrote.

The verdict

SCORE	WHAT IT MEANS
0 to 24	Rarely selected
25 to 49	Sometimes selected
50 to 74	Often selected
75 and over	Usually selected

The rules that keep it honest

These matter more than the formula. They're what stop the score becoming the next visibility number.

- Questions that name your brand don't count toward Found. Appearing when you were asked for proves nothing.
- Every score carries a range. AI answers vary from one asking to the next, so one number alone would claim a certainty nobody has.
- Small samples are pulled toward your average, so three lucky answers can't make a moment look won.
- A stage that wasn't measured is shown as not measured and left out. It's never guessed.

- ▮ A change to the formula or the questions resets the baseline. A change of method is never reported as progress.
- ▮ Only measured outcomes move the score. Publishing a page, finishing a task or switching on a tool does nothing by itself.

That last rule is the one I'd fight for. A score that rises because you used the vendor's product is an advertisement.

One list

The score is out of 100, so the gap to 100 is what you're losing. Split that loss by stage and by moment, and the largest share is your biggest lever.

The list is the fixes, ranked by how much each should lift the score for the effort it takes. Every fix names its moment and its stage, and carries a predicted range.

Then comes the step most marketing skips. After a fix ships, ask again and record what moved. Predicted against measured is how the list earns trust.

Who owns the choice

Every team owns a number, and nobody owns the choice. The score gives the choice an owner.

This is my opinion, not a finding. One person should own the score and the list, across marketing and sales, with the standing to say no to work that moves no moment. In a small firm it's the owner. In a larger one it's often a commercial director or a head of growth. They don't own the channels; they decide which fix goes first, and which team does it. Believed belongs to everyone, which is why it needs an owner most. A contradiction can start in any team's page.

Beside your other numbers

The score doesn't replace the measures a finance director already trusts: attribution, mix modelling and experiments. eBay's paid-search test in Chapter 3 was an experiment, and that's the most trustworthy answer to "did this work?".

All three look backwards, at sales. The score looks forwards, at the choice. So use them together: the score says which fix to try, and an experiment proves whether it paid.

Where AI's influence shows up

Don't judge AI by the visits it sends. Most buyers who act on an AI answer arrive another way: they search your name, or type your address (Chapter 1). Their

visits land in your reports as branded search and direct traffic.

So read three numbers together each quarter: the share of AI answers that name you at your moments, searches for your name, and direct visits. If your name starts appearing in answers and the other two rise a month or two later, that's AI at work. It isn't proof. A launch or a campaign moves the same numbers, so note what else changed. But it's a truer read than the AI line in your analytics.

What it changes

Budget follows the biggest lever. Money tends to drift toward whatever's easiest to count. One score gives it somewhere better to go.

NerdWallet, the American comparison site, made a version of this move in public. Search traffic to some of its pages was falling. Its chief executive, Tim Chen, told investors that marketing decisions would be "based on internal rate of return ... rather than solely on in-quarter profitability" ([NerdWallet](#)).

And there's a measure that costs nothing. Ask every new customer how they found you, and include "an AI assistant" as an answer. Tally's onboarding survey found far more arrivals from AI than its analytics showed (Chapter 17).

What the score doesn't know yet

- ▮ The weights are judgement. They need calibrating against real sales, mode by mode.
- ▮ Nobody has yet shown that the score moves with share. Until that test is run, read it as an index of your odds, not a forecast of sales.
- ▮ Found and Preferred come from the same sample of answers, and those answers aren't fully independent. The real margin is wider than the arithmetic says.
- ▮ It can't see a sale. Only your own data can, and where you have it, it should outrank everything above.

Whose evidence this is. The method and the weights are mine, and RocMode, which I founded, sells software that runs them. Tally sells online forms, and its numbers are its own. NerdWallet's words come from its earnings call, as reported. eBay's test was run by eBay's own researchers. No peer-reviewed study has tested the score.

The loop

A score is only worth taking if it changes what you do next. The whole method is a loop, and it has six steps.

1. **Measure.** Ask the same questions, of the same assistants, the same number of times. Score each moment.
2. **Diagnose.** Find the stage where each moment is lost, and why. Usually something is missing: an answer you never published, a claim nobody else backs, a reason nobody can repeat.
3. **Write down the truth.** Record what's true about you for that moment, once, in one place: the facts, the reason, the price, the limits (Chapter 10). Every channel draws from that record.
4. **Prove it.** Put evidence beside every claim that needs it: a named customer, a figure with a source, an outside review (Chapter 16).
5. **Publish it.** Put the answer and the proof where buyers and machines read: your site, the review platforms, the sources assistants cite (Chapters 15 and 22).
6. **Measure again.** Use the same sample, unchanged. Record what you predicted beside what moved, and drop the fixes that did nothing.

Most teams run the first step and stop. Or they jump from the second to making more content. The work that moves the score is in steps three to five: fixing the source, not the symptom. A dashboard can tell you you're not believed. Only a truth you've written down, proved and published changes it.

In a quarter, it looks like this.

Each quarter ends by recording what moved, then asks again




FIGURE 30.4 The quarterly loop · a change to the questions or the formula starts a new loop with a new baseline

This afternoon, this quarter

This afternoon: take the five questions your buyers ask most. Put each to three AI assistants five times, without naming your brand. Count how often you're

named, and how often you're the one recommended. That's a rough reading of Found and Preferred, for nothing.

This quarter: start here. Score five moments by hand, using the worksheet in the appendix. Fix the biggest lever, then ask again.

Where RocMode comes in

I said at the start that I build software on this thinking. Here's why.

The seven stages are what being selected takes, and the loop is how you improve them. Most teams hold the pieces in seven places, or not at all. The facts sit in a brand book, the proof in a sales deck, the reviews on a platform nobody owns, and the AI answers go unread. So nobody can see that a moment was lost because one claim had no proof behind it.

That's why we built RocMode. It doesn't only measure selectability. It's built to run the whole loop in one place.

- / **It measures.** It samples AI answers to your unbranded questions, scores each check in each moment, and runs the same sample again each period, so you see what moved.
- / **It diagnoses.** It traces a lost moment to its cause: the question you haven't answered, the claim with no proof, the contradiction between two of your own pages, or the reason the machine gives instead of yours.
- / **It holds the truth.** At its centre is BrandGraph, one governed record of what's true about your brand: each fact, claim and reason, the evidence behind it, and where it's published. A claim without a source is flagged before it goes out. When the price changes, it changes once.
- / **It shows its reasoning.** Every item on the list carries its chain: which answers, which sources, which missing proof. You can check the working, not just the number.
- / **It closes the loop.** Pages and content are drafted from that one record for you to approve, and the next sample shows whether the machine's choice changed.

There are two ways to start, and they suit different firms.

- / **Start here: do it yourself, for five moments.** The appendix has the worksheet. An editable version, with every tool from the stage chapters, is free in the Selectability Playbooks at rocmode.io/playbooks. Five moments by hand take an afternoon to a day each quarter, and that's the right place to start.
- / **Use RocMode when you want fifteen or more, every quarter.** At that size the asking, counting and fixing outgrow a spreadsheet.

Be clear about what it can't do. No software earns your reviews, wins you a reference customer or makes the pitch. It tells you which of those to do first, and proves whether they worked.

Today it scores five of the stages, Found through Bought, from AI answers and from your own site. Known and the renewal moments are in the method above and not yet in the product.

The formula is published, so you can check our working or do it yourself. Agencies and consultants can run the method for their clients too, and RocMode is built for agencies running it across many clients. The method belongs to whoever uses it.

Seen isn't selected. Now there's a way to tell the difference, and to close it.

IN SHORT · CHAPTER 30

WHAT'S TRUE

A brand needs one number for how likely it is to be selected, and one list of what to fix first.

DO THIS QUARTER

Score five moments by hand with the worksheet, fix the biggest lever, then ask the same questions again.

WATCH

Each moment's score and its range, quarter on quarter, and predicted against measured lift for every fix you ship.

APPENDIX

The self-score *worksheet*

This worksheet scores your Selectability by hand. It follows Chapter 30 step by step.

Be realistic about the time. A full sample is 45 answers per moment, so fifteen moments is 675 answers: several days by hand. For a first pass, use two assistants and three runs of each phrasing. That's 18 answers per moment, 270 for fifteen moments, and roughly a day. Five moments at the first-pass size take an afternoon.

The rubrics are my own working rules, drawn from the stage chapters. They're not tested against sales. A hand score is a rough reading, and its value is in repeating it the same way each quarter.

The ten steps

1. **Pick a mode.** Choose the one from Chapter 7 that fits most of your sales. It sets the weights for the four quality checks.

MODE	ANSWERED	BELIEVED	PREFERRED	BOUGHT
Habitual	10%	20%	15%	55%
Quick choice	15%	25%	15%	45%
Considered personal	25%	30%	30%	15%
Owner-led business	20%	35%	30%	15%
Committee	25%	30%	30%	15%

1. **List up to 15 moments.** Each is one kind of buyer at one step, from your map in Chapter 12. Write each in the buyer's own words. Renew-or-switch moments count too, and are scored the same way.
2. **Write three unbranded questions per moment.** Three different phrasings of what that buyer would ask. Make them long and specific, with the buyer's situation in them. None may name you.
3. **Run a presence sample.** Put each phrasing to three assistants, five times each, in fresh sessions. That's 45 answers per moment. For a first pass, use two assistants and three runs: 18 answers. For each answer, note whether you were named, whether you were the one recommended, and the reason given. With 45 answers, the margin on Found is roughly 15 points either way; with 18, about 23. Both assume each answer is independent. Repeated phrasings to the same assistants aren't, so treat the real margin as wider.

4. **Rate the six checks from 0 to 100.** Use the rubric table below, one moment at a time. If fewer than 10 answers in a moment name you, mark that moment's Preferred as unreliable. Score it, but don't rank a fix on it alone.
5. **Work out the gate and the quality.** Gate = $\sqrt{\text{Known} \times \text{Found}}$. Quality = the weighted average of Answered, Believed, Preferred and Bought, using your mode's weights.
6. **Score the moment.** Moment = Gate $\times$ Quality. If you work in points out of 100, divide the product by 100.
7. **Average the moments.** A plain average is fine for a first score. Later, weight moments nearer the decision, and your main buyers, more heavily.
8. **Mark the band.** 0 to 24 is rarely selected, 25 to 49 sometimes, 50 to 74 often, and 75 and over usually. Write the margin beside it.
9. **List the three biggest gaps.** Look at the gate first, because a low gate multiplies everything else down. Then, for each quality check, multiply its weight by its shortfall from 100. The largest products, in the moments nearest the decision, are your list.

The rubrics

CHECK	THE TOOL	HOW TO RATE IT FROM 0 TO 100
Known	The moment memory map (Ch 13)	Rate buyer memory and machine memory each: strong 100, weak 50, absent 0. Take the average. Buyer memory comes from ten customer interviews; machine memory from needs-based prompts and asking the machine about you by name
Found	The presence sample (Ch 14)	The share of your 45 unbranded answers that name you, as a percentage
Answered	The question chain map (Ch 15)	Score each link in the buyer's chain: answered in public 1, partly 0.5, avoided 0. Divide by the number of links, then multiply by 100
Believed	The claim, proof and source inventory (Ch 16)	Score each claim about you: 0 if only you say it, 1 if one independent source backs it, 2 if two or more of different kinds do, 3 if assistants repeat it and cite them. Take 1 off for any live contradiction, but never go below 0. Average the scores, divide by 3, then multiply by 100
Preferred	The reason gap sheet (Ch 17)	Average two shares. First: of the answers that name you, the share where you are the one recommended. Second: the reason-match rate. Ask recent buyers why they chose you, and score each answer 0 (no reason), 1 (a generic reason), 2 (your reason, roughly) or 3 (your reason, in words you would use). The rate is the share scoring 2 or 3
Bought	The risk-removal ladder and decision friction audit (Ch 18)	Eight yes-or-no checks, 12.5 points each. The landing page repeats the reason, shows proof and gives a clear next step. There is a small first step, a way out, a plain recommendation, proof from someone like them, and someone to call. Where you have sales data for the moment, it outranks this

CHECK	THE TOOL	HOW TO RATE IT FROM 0 TO 100
Kept	The promise ledger and renewal moment map (Ch 19)	Not a check of its own. Score your renew-or-switch moments with the six checks above. Your renewal and repeat data outranks the result

The bookkeeper's worksheet

The bookkeeper for cafés is an invented example. Her offer is "Bookkeeping for cafés. Your wage cost against takings, every Monday by nine." All the figures below are invented to show the method. They describe her before the fixes in the stage chapters.

Her mode is owner-led business, so her weights are 20% Answered, 35% Believed, 30% Preferred and 15% Bought. She listed nine moments. Three are shown here.

MOMENT, IN THE BUYER'S WORDS	KNOWN	FOUND	ANSWERED	BELIEVED	PREFERRED	BOUGHT
A. "Wages ate the takings again this month" (comparing)	75	40	38	33	35	50
B. "Our accountant says the books are a mess before tax time" (learning)	50	10	25	33	20	50
C. "My accountant only sends numbers at year end" (buying)	50	25	38	33	40	50

Moment A's Known is strong buyer memory and weak machine memory, from her moment memory map: $(100 + 50) \div 2 = 75$. Its Answered is two links answered, two partly and four avoided, from her question chain: $(2 + 1) \div 8 = 38$. Believed is 33 everywhere. Each of her claims is backed by exactly one independent source, a single client's review, and no assistant repeats it. That scores 1 out of 3 for every claim. Bought is four of the eight checks.

Moment A's Preferred shows how that check works. Of her 45 answers, 18 name her (Found 40). In 6 of those 18 she's the one recommended: 33.3%. She asked eight recent clients why they chose her, and three gave her reason, cafés and the Monday number, in words close to hers: 37.5%. The average of 33.3 and 37.5 is 35.

Moment B is different. Only about four or five answers name her there, fewer than 10, so she marks its Preferred as unreliable.

MOMENT	GATE	QUALITY	MOMENT SCORE
A	$\sqrt{(75 \times 40)} = 55$	37	$55 \times 37 \div 100 = 20$
B	$\sqrt{(50 \times 10)} = 22$	30	$22 \times 30 \div 100 = 7$
C	$\sqrt{(50 \times 25)} = 35$	39	$35 \times 39 \div 100 = 14$
Average of the three			14: rarely selected

Fourteen sounds bleak, but it's useful. Found alone carries a margin of about 15 points, so she reads it as "rarely selected", not as 14 exactly.

Her three biggest gaps fall out of the arithmetic.

RANK	MOMENT	STAGE	WHAT IS WRONG	THE FIX
1	B	Found (the gate)	The machine never names her at tax time	Write with the two accountants who refer her (Ch 13)
2	All	Believed	One client review backs her claims, and no assistant repeats them	Publish the twelve-month case, and ask two clients to mention the Monday number in reviews (Ch 16, Ch 17)
3	A and C	Preferred	When she is named, she is rarely the one recommended, and her reason is seldom given	Give reviewers and partners her reason in words they can repeat: "only cafés, the wage number every Monday" (Ch 17)

The ranking follows step 10. Moment B's gate is the lowest, so it comes first. Then weight $\times$ shortfall: Believed is $0.35 \times 67 =$ about 23 in every moment. Preferred is $0.30 \times 65 =$ about 20 in A and $0.30 \times 60 = 18$ in C; B's is set aside as unreliable. Answered is only $0.20 \times 62 =$ about 12 in A and C, so its fixes wait.

Rather than telling her to post more, the score told her that her weakest moment is held back by the gate, and every moment by belief. Both are fixed by other people writing about her, not by her writing more.

Next quarter she asks the same 27 questions, the same 15 times each. If moment B's Found has moved, the fix worked. If not, she tries the next one.

Blank templates

Template 1: Moments and questions.

#	BUYER	STEP	MOMENT, IN THE BUYER'S WORDS	PHRASING 1	PHRASING 2	PHRASING 3
---	-------	------	------------------------------	------------	------------	------------

Template 2: Presence sample. One row per moment, per assistant. Five runs of each of the three phrasings gives 15 answers per row; three runs, for a first pass, gives 9.

MOMENT	ASSISTANT	ANSWERS	NAMED	RECOMMENDED	REASON GIVEN	WHO WAS NAMED INSTEAD
		15				

Template 3: Ratings and scores.

MOMENT	KNOWN	FOUND	GATE	ANSWERED	BELIEVED	PREFERRED	BOUGHT	QUALITY	MOMENT SCORE
--------	-------	-------	------	----------	----------	-----------	--------	---------	--------------

Template 4: The verdict.

DATE	MODE	WEIGHTS (A / BE / P / B0)	MOMENTS SCORED	AVERAGE SCORE	MARGIN	BAND	METHOD	VERSION
------	------	---------------------------	----------------	---------------	--------	------	--------	---------

Template 5: The three biggest gaps.

RANK	MOMENT	STAGE	WHAT IS WRONG	THE FIX	OWNER	PREDICTED LIFT	MEASURED LIFT
------	--------	-------	---------------	---------	-------	----------------	---------------

Download the editable worksheet free at rocmode.io/playbooks, or have RocMode run it for you.

AFTERWORD

Why I wrote *this*

I studied maths at university and planned to become an actuary. Three years in, with top grades, I realised I was heading for a life in statistics, which I found the duller corner of the subject.

Around the same time, in 2000, I started a web design agency with my best friend. It lasted two years. I spent two more designing marketing and sales processes, and somewhere along the way I fell for marketing.

It unsettled me. I'd spent my life on problems with right and wrong answers, and here was a field that seemed to have neither.

I decided that couldn't be true. Marketing had to be a formula, just a very complicated one, and I've spent my career since trying to find it.

That search ran through agencies: Cru Digital, whose start you met in Chapter 13, then a merger that became JuniorCru, then CXA, the Customer Experience Agency. In between, from June 2017 to July 2019, I was client-side, as Head of Digital at the bank then called CUA. By the JuniorCru years I saw customer experience as the frame for all of marketing: one journey, end to end, and I began designing, in my head, the software to run it. At CXA, with Cos Luccitti, I built the 6D model of customer experience that much of this book stands on.

Through all of it I held one belief. The answer to marketing is the right message, in front of the right person, at the right time. Seen through the buyer's journey, that means many messages, for many people, at many stages. In the age of AI it also means the right answer to the right question.

Then came 2025. I watched AI answers take the searches our content for a lender used to win, and for a while its leads went with them. When I asked an assistant for the top digital agencies in Brisbane, it didn't name mine.

For the best part of a year I'd been working sixteen-hour days, often more. I knew I was close to burning out. I didn't know how close.

On 8 October 2025, while driving, I micro-napped at the wheel and hit a bus head-on. There was a moment, and I remember it clearly, when I thought: "This is it. This is the end." Then I went blank.

I came to a few minutes later. I'd hit the bus at an angle, scraped along nearly half of its side and spun several times. My arm was broken in three places and my shoulder was fractured. There was a small cut on my stomach where a piece of metal had pushed in, no doubt centimetres from being the end of me. Apart from me, nobody was hurt beyond a bruised shin.

I was incredibly lucky to survive, let alone walk away relatively unscathed. Even so, my shoulder took more than nine months of rehab.

In the days after, I reconsidered pretty much everything. I couldn't keep selling time for money. That had to end. I couldn't keep risking my life, or the life of anyone in my car, by living like that. And with my arm months away from healing, I couldn't have kept working that way even if I'd wanted to.

So I made myself a promise. In the next twelve months I'd build the software I'd been designing in my head for years, and I'd write the book I'd always wanted to write: something close to the answer to marketing. Twelve months later, here we are.

I started on the software in hospital, typing with my good left arm and asking AI to help. It was slow going, and AI wasn't up to the job back then. So I made a bigger call. I handed off all the day-to-day work I'd been doing in the agency, mostly design and code, and within days I went from sixteen-hour days to about two hours of work a day.

It was incredibly liberating. Then, within a couple of months, I was bored. I had the energy for a solid eight-hour day again, and only two hours of work to fill it.

On 2 December 2025, a couple of months after the accident, Cos Luccitti, now RocMode's chief executive, came to me with an idea that was very close to my own. We called it BrandGraph at first. It became RocMode, and I've worked on almost nothing else since. You're holding the other half of the promise.

I haven't solved marketing. But I think I know why it looked unsolvable. You can't optimise anything until you've decided what you're optimising for, and marketing never agreed.

Awareness, leads, traffic, brand: every team picked its own. This book argues there's one answer: being selected. Marketing is the work of raising the odds that a buyer picks your brand over the alternatives, at every step of their journey.

You can't force that choice, so there'll never be a clean solution. But the odds can now be estimated, and soon tested and improved. AI lets us do that at a speed and scale no marketing team has had before.

None of this replaces craft. The maths can show where a choice is being lost. It takes imagination to win it back.

And yes, I see the irony. I walked away from statistics, and the answer turned out to be statistics.

So I should be plain about my interest: I build software based on this thinking. That's why the argument in this book had to stand without it.

One idea from my notebooks didn't fit the chapters, so it lives here. I think of brands as living things. They grow, adapt, defend themselves and pursue goals.

Most do it by accident.

The brand I want to help build does it on purpose. It senses how buyers and machines see it, moment by moment, without waiting for a quarterly report. It protects what's true about it, and catches a contradiction or a broken promise before it spreads. And it expresses one truth in many forms, each fitted to the moment, without ever becoming a different story.

It's no accident that the word is selection. Living things are selected by their environment. Brands are selected by buyers, and now by the machines that advise them.

That's the long-term vision, and it's a vision, not a product. The score in Part five is the first part of it: a way for a brand to sense where it stands. The rest is work for the years ahead.

AFTERWORD

What we don't *know* yet

This book leans on evidence that's young, and much of it comes from firms with something to sell. These gaps matter most.

- ✓ **Habitual buying.** No study I found shows AI changing which brand people buy on habit.
- ✓ **Memory and machines.** No published study tracks a brand from building its name to being recommended more by AI. The evidence is correlational, or uses invented brands. Nor has anyone shown that moment-based memory measures predict sales better than prompted awareness.
- ✓ **How many runs are enough.** No independent study says how many AI answers make a stable rate. The margins in Chapter 14 and the appendix are arithmetic, not evidence.
- ✓ **Which answers decide.** No study measures which questions settle a purchase. Nobody has tested whether answering price, risk or comparison questions changes how often AI recommends you.
- ✓ **Consistency and the one truth.** No study links consistency across sources to AI recommendation or to sales. The best papers suggest machines pick one version rather than flag a conflict. A structured brand record improves the answers you generate yourself, but nobody has shown that outside assistants reward it. Studies of what AI cites disagree widely, from about 1% to 44% for brands' own sites.
- ✓ **Reasons.** No study links a reason that customers repeat to a higher win rate. Nor do we know whether assistants repeat a brand's own reason, or whether the reason a machine gives is the one that decided.
- ✓ **The last step.** The size of "no decision" in business buying comes from one proprietary dataset. Most guarantee evidence comes from lab studies of intention, not from real purchases.
- ✓ **Keeping and the machine.** No study links a firm's churn or service failures to a later fall in AI recommendation. I found no business-to-business evidence on AI use at renewal.
- ✓ **Challengers and small firms.** I found no verified case of a challenger displacing a bigger rival in AI answers through better evidence. Tally comes closest, and its numbers are its own. I found no data on how owners choose a bookkeeper, an agency or a lawyer, with or without AI.
- ✓ **Cause and effect.** No randomised experiment yet links an AI recommendation to an actual purchase. No vendor or buyer has described, on the record, a large deal won or lost because of what an AI said.

- ▮ **The score itself.** Which stages carry each mode is my inference. The weights in Chapter 30 are judgement, and not yet fitted to sales.
- ▮ **The score against share.** Nobody, including me, has yet shown that a higher Selectability Score goes with more market share or more sales. Chapter 30 sets out the test, and we're running it.
- ▮ **How independent AI answers are.** The margins in this book assume each answer is a separate draw. Answers to the same phrasings, from the same assistants, aren't. Nobody has measured by how much that widens the real margin.
- ▮ **Most of the world.** The company stories here come from North America, Europe, China and Australia. I found nothing verified elsewhere.

Most of the surveys cited were run by vendors, and each chapter's note says whose evidence it is. Almost everything known about how marketers are faring comes from surveys of opinion, not measured results. The evidence was reviewed in October 2026, and it'll date quickly.

Every gap on this list can be closed, and most need data no single firm holds. If you score your moments, share what you find. Send your scores beside your sales through rocmode.io/research, and the findings will be published. The worksheet and every tool in this book are free at rocmode.io/playbooks.

CODA

Every brand is seen by someone.

Very few are selected on purpose.

*The difference lives in moments: a buyer, a question,
a choice.*

*You can now estimate which moments you win, and
which you lose.*

*Start with five of them this quarter. The tools are
free at rocmode.io/playbooks.*

Then fix the biggest gap, and ask again.

Seen isn't selected. Go and be selected.

AFTERWORD

Sources

Every source is linked where it's used. Here they are in one place, chapter by chapter. Each chapter's "Whose evidence this is" note says who ran each study and what they sell.

1. SO MUCH, SO LITTLE

TechCrunch techcrunch.com

Google blog.google

Pew Research Center pewresearch.org

SE Ranking seranking.com

Semrush semrush.com

TechCrunch techcrunch.com

Marketing Week marketingweek.com

CMO Survey cmosurvey.org

Graphite graphite.io

Constant Contact news.constantcontact.com

2. MARKETING'S ONE JOB

Ehrenberg-Bass marketingscience.info

Kantar kantarnewzealand.com

3. THE VISIBILITY TRAP

Ehrenberg-Bass marketingscience.info

Berkeley Haas newsroom.haas.berkeley.edu

Blake, Nosko and Tadelis nber.org

Gartner gartner.com

Skift skift.com

Tripadvisor sec.gov

PhocusWire phocuswire.com

4. THE ANATOMY OF A CHOICE

Vakratsas and Ambler journals.sagepub.com

5. THE CHOICE STARTS WITHOUT YOU

6sense 6sense.com

Google, via Marketing Dive marketingdive.com

McKinsey mckinsey.com

Think with Google thinkwithgoogle.com

G2 sell.g2.com

6. MOMENTS, NOT FUNNELS

Ehrenberg-Bass marketingscience.info

Google adwords.googleblog.com

Google business.google.com

Semrush semrush.com

7. FIVE WAYS PEOPLE BUY

Ehrenberg-Bass marketingscience.info

G2 images.g2crowd.com

Capterra capterra.com

G2 sell.g2.com

8. WHY PEOPLE CHOOSE

Zajonc doi.org

IPA ipa.co.uk

Salganik, Dodds and Watts princeton.edu

Scheibehenne and colleagues archive-ouverte.unige.ch

Chernev, Böckenholt and Goodman
kellogg.northwestern.edu

9. THE NEW CHOOSER

Ipsos ipsos.com

BCG bcg.com

Marketing Science pubsonline.informs.org

Decrypt decrypt.co

Digital Commerce 360 digitalcommerce360.com

Similarweb aisearch.similarweb.com

Iannelli and Ai arxiv.org

TrustRadius prnewswire.com

China Media Project chinamediaproject.org

10. ONE TRUTH, MANY ANSWERS

Aggarwal and colleagues arxiv.org

Sequeda and colleagues arxiv.org

Search Engine Journal searchenginejournal.com

Allouah and colleagues arxiv.org

11. A REASON TO CHOOSE YOU

Romaniuk, Sharp and Ehrenberg doi.org

Ehrenberg-Bass marketingscience.info

US Census Bureau census.gov

Roger Martin bridgespan.org

Wise wise.com

Wise plc investigate.co.uk

12. KNOW YOUR MOMENTS

Webster and Wind, 1972 [jstor.org](https://www.jstor.org)
 Bonoma, 1982 hbr.org
 Gartner gartner.com
 Oxford University Press global.oup.com
 Bettencourt and Ulwick, 2008 hbr.org
 The Re-Wired Group therewiredgroup.com
 Guest, Bunce and Johnson, 2006
journals.sagepub.com
 Strategyn strategyn.com
 Retail TouchPoints retailtouchpoints.com

13. BE KNOWN

Ehrenberg-Bass marketingscience.info
 Romaniuk business.linkedin.com
 Ludwichowska and colleagues
digital.library.adelaide.edu.au
 Hoyer and Brown ideas.repec.org
 Chu and Hou arxiv.org
 Malthouse and colleagues arxiv.org
 Ahrefs ahrefs.com
 B&T bandt.com.au
 Romaniuk ideas.repec.org
 IPA ipa.co.uk
 AdNews adnews.com.au
 21 Hats 21hats.substack.com

14. BE FOUND

SparkToro sparktoro.com
 Ahrefs ahrefs.com
 Profound tryprofound.com
 Google blog.google
 BrightEdge brightedge.com
 MarketBeat transcript marketbeat.com
 Zatuchin arxiv.org
 Ahrefs ahrefs.com
 Backlinko backlinko.com
 Ranqo arxiv.org
 IPA ipa.co.uk

15. BE ANSWERED

NYT reprint, SupplyChainBrain supplychainbrain.com
 River Pools riverpoolsandspas.com
 Heidi Cohen heidicohen.com
 MarTech martech.org
 INFORMS informs.org
 Eisend sciencedirect.com
 Ein-Gar, Shiv and Tormala academic.oup.com
 Crowley and Hoyer ideas.repec.org
 TrustRadius solutions.trustradius.com
 Gartner, September 2019 gartner.com
 Google blog.google
 Alphabet blog.google
 Semrush semrush.com
 WebFX webfx.com
 Gartner, 2026 gartner.com
 Ellison and Ellison nber.org
 HockeyStack hockeystack.com

16. BE BELIEVED

Nielsen nielsen.com
 McKinsey mckinsey.com
 Yext yext.com
 FTC ftc.gov
 CAP Code asa.org.uk
 Hou and colleagues arxiv.org
 Wu, Wu and Zou arxiv.org
 Moffatt v. Air Canada mondaq.com
 Ziegler, Diehl and Ruther journals.sagepub.com
 Eisend and Küster link.springer.com
 Spiegel spiegel.medill.northwestern.edu
 Luca hbs.edu
 Profound tryprofound.com
 Semrush semrush.com
 Domino's sec.gov
 FTC ftc.gov
 CBC cbc.ca

17. BE PREFERRED

Shafir, Simonson and Tversky sciencedirect.com
 Simonson gsb.stanford.edu
 Nisbett and Wilson www3.nd.edu
 Iannelli and Ai arxiv.org
 B&T bandt.com.au
 CEB in HBR hbr.org
 HBR hbr.org
 Reichheld hbr.org
 Berger and Schwartz faculty.wharton.upenn.edu
 Tally, 22 September 2026 blog.tally.so
 Tally, June 2025 blog.tally.so
 Romaniuk, Sharp and Ehrenberg academia.edu
 Rockland Dutton survey thefiduciaryinstitute.org

18. BE BOUGHT

Dixon customerthink.com
 Dixon and McKenna's deck typo9demo.canto.com
 Brown, Imai, Vieider and Camerer aeaweb.org
 Gal and Rucker papers.ssrn.com
 Samuelson and Zeckhauser link.springer.com
 Ritov and Baron cris.huji.ac.il
 Harinck, Van Dijk, Van Beest and Mersmann psychologicalscience.org
 Barton, Zlatevska and Oppewal ideas.repec.org
 Mathur and colleagues arxiv.org
 HWL Ebsworth hwlebsworth.com.au
 CNBC cnbc.com
 G2 prnewswire.com
 Baymard baymard.com
 Yoganarasimhan, Barzegary and Pani faculty.washington.edu
 Suwelack, Hogreve and Hoyer researchgate.net
 AMA summary of Janakiraman, Syrdal and Freling ama.org

19. BE KEPT

J.D. Power jdpower.com
 Digital Commerce 360 digitalcommerce360.com
 Schmitt, Skiera and Van den Bulte faculty.wharton.upenn.edu
 Kumar, Petersen and Leone hbr.org
 Askalidis, Kim and Malthouse arxiv.org
 Anderson deepblue.lib.umich.edu
 Houston slideshare.net
 Tally blog.tally.so
 Oliver doi.org
 Tech.eu tech.eu
 Entrepreneur entrepreneur.com
 ABC abc.net.au
 FTC ftc.gov
 CMS cms.law
 Romaniuk, Dawes and Nenycz-Thiel ideas.repec.org
 MeasuringU measuringu.com
 Shin et al. research.polyu.edu.hk

20. CONTENT

Graphite graphite.io
 Ahrefs ahrefs.com
 Doshi and Hauser science.org

21. THE WEBSITE

Pew pewresearch.org
 Microsoft Clarity, via PPC Land ppc.land
 TrustRadius solutions.trustradius.com

22. SEARCH AND AI ANSWERS

Google Search Central developers.google.com
 Ahrefs ahrefs.com
 Cloudflare cloudflare.com
 Google developers.google.com
 Microsoft microsoft.com

23. ADVERTISING

Blake, Nosko and Tadelis nber.org
 IPA ipa.co.uk
 Gerber and colleagues actacommerci.co.za
 Gordon and colleagues kellogg.northwestern.edu
 Simonov, Nosko and Rao ideas.repec.org
 IPA ipa.co.uk

24. SOCIAL

Meta about.fb.com
 Socialinsider socialinsider.io
 Colicev and colleagues journals.sagepub.com
 Ahrefs ahrefs.com
 Goel and colleagues sharad.com

25. CREATORS

IAB iab.com
 Wies, Bleier and Edeling journals.sagepub.com
 ACCC acc.gov.au
 Eisend and colleagues doi.org

26. PR AND REVIEWS

AMEC amecorg.com
 FTC ftc.gov
 Proserpio and Zervas open.bu.edu
 Cision newswire.ca
 FTC ftc.gov
 CMA gov.uk
 Wikipedia en.wikipedia.org

27. THE PITCH

HBR hbr.org
 Mayer cambridge.org

28. EMAIL

Apple apple.com
 Google blog.google
 Klaviyo klaviyo.com
 Litmus litmus.com
 ACMA acma.gov.au

29. AUTOMATION

Oldroyd, McElheran and Elkington thedenmangroupselling.wordpress.com
 Gartner gartner.com

THE REST, BRIEFLY

Gartner gartner.com
 Ehrenberg-Bass marketingscience.info
 CMI contentmarketinginstitute.com
 Xero xero.com
 Sharp and Sharp doi.org

30. THE SELECTABILITY SCORE

NerdWallet aol.com