NATE BURKE

A FIELD GUIDE



# RESIST COMFORT

The Cybersecurity Marketer's Guide to Doing Work Worth Remembering

# RESIST COMFORT

## The Cybersecurity Marketer's Guide to Doing Work Worth Remembering

NATE BURKE

First edition, November 2026.
Free to read at nathanwburke.com/resist.
It started as the opening keynote at CyberMarketingCon 2026. The talk is the compression.
This is the whole thing.

**BORING IS THE ONLY UNFORGIVABLE SIN.**

# CONTENTS


00 Before we start 06

**PART ONE** WHY YOUR BEST WORK KEEPS COMING OUT AVERAGE

01 Meet Your Neanderthal 13
02 The Sameness Machine 25
03 Safe Is Just a Bet You Don't Have to Defend 35

**PART TWO** THE COMFORT CHECK

04 The Three Questions 44
05 Question One: What Is the Comfortable Assumption? 51
06 Question Two: What Is the True Objective? 61
07 Question Three: What Is the Redesigned Approach? 69
08 When the Check Fails: A Post-Mortem 78

**PART THREE** THE PRACTICE

09 Better Is Subjective, Different Is Measurable 90
10 Design for Memory, Not Just Explanation 97
11 What Worked Before Is a Ceiling, Not a Floor 107
12 Curiosity Before Explanation 118

**PART FOUR** THE MACHINE THAT MAKES AVERAGE

13 Your Neanderthal Now Has a Copilot 126
14 Crutch or Multiplier 132
15 What I Actually Built 137
16 Your Value Was Never the Tasks 148
17 Redesigning the Work 154

**PART FIVE MONDAY**

| | | |
|---|---|---|
| 18 | Seven Beliefs, Seven Mondays | 161 |
| 19 | Stop Letting Him Vote First | 168 |

**APPENDICES THE TOOLS, COLLECTED**

| | | |
|---|---|---|
| A | The Comfort Check | 172 |
| B | The Sameness Audit | 175 |
| C | The Bet Card and the Measurement Plan | 177 |
| D | The Assumption Inventory | 180 |
| E | Deriving an Unpublished Rubric | 182 |
| F | The Skill Library | 184 |
| G | Sources | 189 |

# 00

BEFORE WE START

# BEFORE WE START

I want to make the argument before I earn it, so you know what you're agreeing to read.

Every other mistake in this job is survivable. Wrong channel, wrong message, wrong timing, wrong number, wrong hire, wrong quarter. I've made all of them, several more than once, and every one of them left something behind that I could learn from.

The one you don't come back from is making something nobody remembers. Not because it's the most expensive mistake, but because there's nothing to fix afterward. Nothing happened. You spent the budget, you filled the calendar, you hit every deadline, everybody did their job, and the market never noticed that you were there. There's no post-mortem for that, because there's no body.

That's the whole book. Everything else is mechanism and method.

## WHERE THIS CAME FROM

My senior yearbook page has an ambition printed under the photo. Mine says: "To find something I'm good at, and make a lot of money doing it."

That's what a seventeen-year-old writes when he has absolutely no idea what he wants to do.

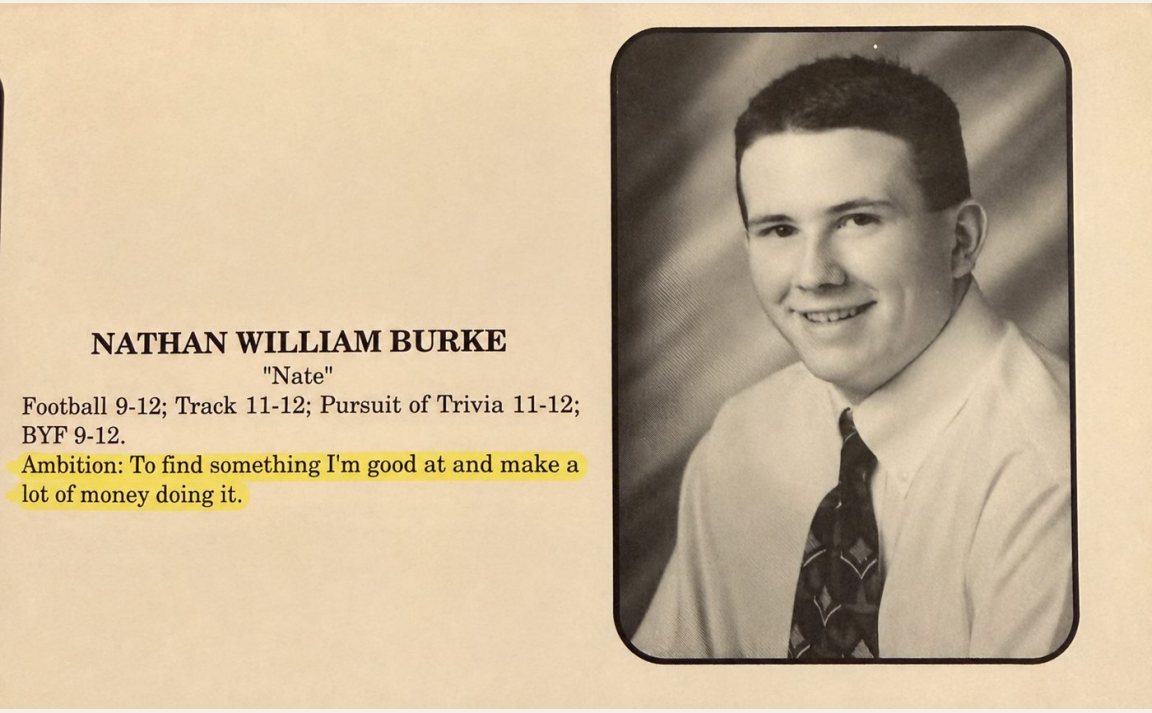



**FIG. 1** My senior yearbook page. Read the ambition.

Marketing came easy, so I majored in it. Freshman year we had to design a cereal box, and I got extra credit for my groundbreaking idea: make it colorful, and put it on the bottom shelf, where a kid in a shopping cart can grab it. I got extra credit for describing the cereal aisle. What I remember about it isn't pride. It's a small, uncomfortable feeling that it had been too easy, and that I wasn't learning anything.

So I took one sociology class, and it broke my brain a little. It turns out there's a science to why people in groups do what they do, and very little of it is logic. I switched majors. There weren't a lot of jobs for sociologists, so I taught myself to code, almost went for a PhD, and spent seven years as a developer instead, which is a strange job for somebody whose favorite subject is other people. Then I joined a startup and fell back in love with marketing, this time through the lens of why people actually decide. I've been marketing in cybersecurity ever since, starting around 2008: CloudLock, Hexadite, Axonius, and now 7AI.

I'm telling you this for two reasons. The first is so you know I'm not an outsider throwing rocks. I'm a vendor marketer in cyber, with the same tools you have and the same pressure you're under, and I'm going to implicate myself throughout. The second is that every turn in that path started the same way. Something got easy, and easy started to feel wrong.

This book is mostly about that feeling, and about why we're all wired to ignore it. Nearly every chapter leans on a finding from social psychology or sociology, because the things we treat as "just how it works" usually have a name and decades of research behind them. If your own path looks like a mess, good. The detours are where I got everything in here.

## WHO THIS IS FOR

You work in cybersecurity marketing. You might run the function or you might be one of six people in it or you might be the entire function at a Series A company where the CEO still writes the social posts. The book is written so that none of that changes what you can do with it, because the expensive part of this work has never been the budget.

If you're reading this as somebody who controls a large budget, you'll find that most of the methods here don't need one, which should worry you slightly.

If you're reading this as somebody who controls none, you'll find that every chapter ends with the version that costs nothing, and those versions aren't consolation prizes. Several of them are the better option.

## WHAT IT WILL COST YOU

About four hours to read and considerably more than that if you actually run the exercises, which is the only reason to read it.

It will also cost you some comfort, which is the point and also the joke in the title. There's a chapter in here about the most expensive failure of my career, told in more detail than my publicist would prefer, and if I'm asking you to look honestly at the work your team approved last quarter, I should go first.

## HOW TO USE IT

Read it in order the first time. The argument builds, and Part Three won't land if you skip Part One.

After that it's a reference. The audits, templates, and inventories are collected in the appendices so you don't have to hunt through chapters for them.

## THE ONE THING YOU NEED BEFORE CHAPTER 1

There's a tool that runs through this entire book. It appears at the end of every chapter, filled in against whatever that chapter was about, and by the end you'll have watched it applied about twenty times, including once where it fails completely and I have to explain why.

It's three questions. You can run it this afternoon, on anything, without asking anybody for permission or money.

## THE COMFORT CHECK

### 1. WHAT'S THE COMFORTABLE ASSUMPTION?

Not "what's our strategy." What are we treating as simply how this works? Usually it's the thing nobody in the room would think to question, because nobody has ever said it out loud.

### 2. WHAT'S THE TRUE OBJECTIVE?

Not the activity. The outcome. Be honest here, because the activity is always easier to write down and it always feels like progress.

### 3. WHAT'S THE REDESIGNED APPROACH?

If you threw out the current process entirely and built backward from that outcome, what would you actually do?

Three questions. No budget, no permission, no headcount, no software.

Before you start Chapter 1, pick one thing your team is currently doing because you did it last year, and write down the answer to question one. Just the first one. Don't solve it.

Most people find that writing the sentence is harder than they expected, and that's the first useful thing this book will tell you.




PART ONE

# WHY YOUR BEST WORK KEEPS COMING OUT AVERAGE

# 01

CHAPTER ONE

# MEET YOUR NEANDERTHAL

**HEY NATHAN! YOU HAVE MORE NEANDERTHAL DNA THAN 91% OF OTHER CUSTOMERS.**

A couple of years ago I spit in a tube and mailed it to a company, the way you do.

This is what came back.

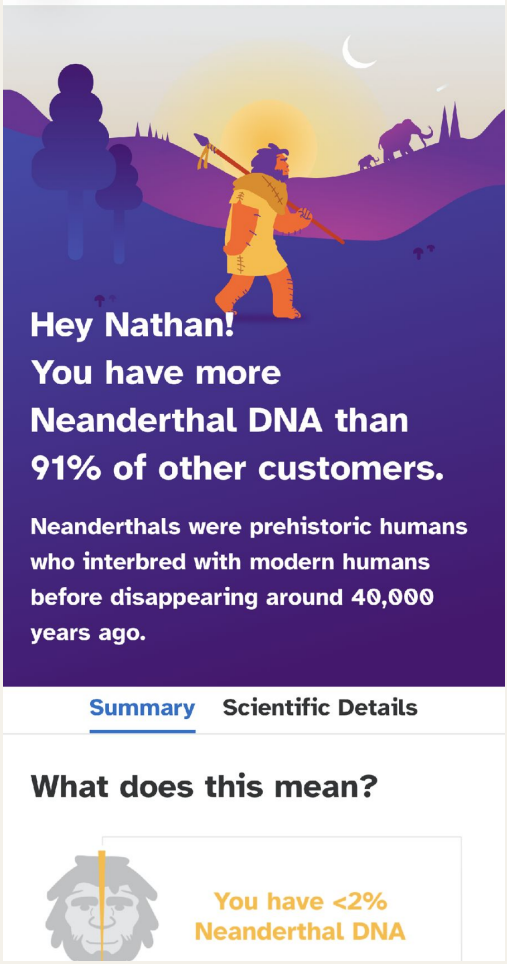



**FIG. 1.1** The 23andMe Neanderthal report. Ninety-first percentile.

I paid money to find that out.

Ninety-first percentile. Put seven hundred people in a room and I'm ahead of six hundred and thirty of them, which also means about sixty people in that room are carrying more of him than I am and have no idea.

In fairness to the guy, I went and looked at what he actually gave me. The report lists the variants. Here's the measurable legacy of forty thousand years of human history as it applies specifically to me:

1. A worse sense of direction.
2. More likely to sweat during a workout.
3. Less likely to blush easily.
4. Two separate genetic variants that make me prefer dark chocolate over milk chocolate.

That's it. I get lost, I sweat, I don't blush, and I won't eat a Hershey bar.

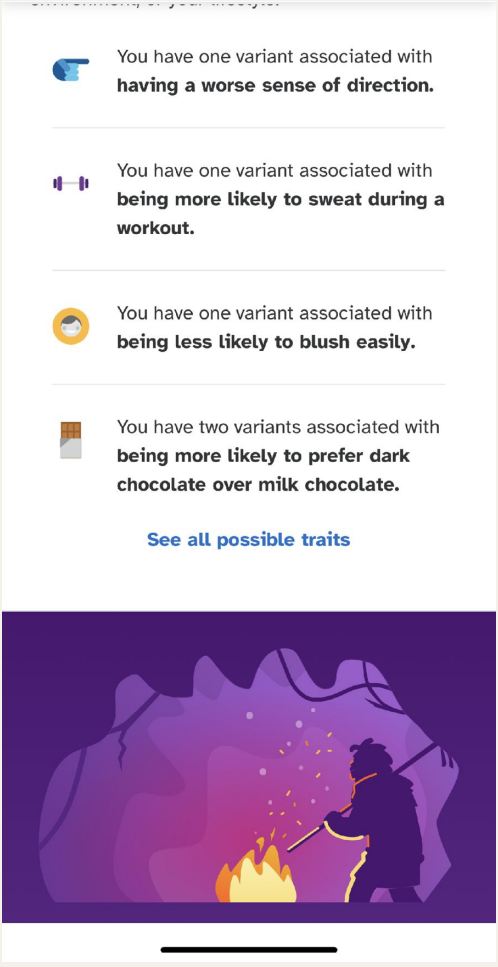



**FIG. 1.2** The full measurable legacy of forty thousand years, as it applies to me.

///

## WHY HE WAS RIGHT

Here's the part that makes him worth keeping around, and it isn't the chocolate.

For almost all of the time humans have existed, the familiar thing was the correct answer. Not the safe answer. The correct one.

Think about what an unfamiliar option actually cost back then. A plant nobody in your group had eaten. A route nobody had walked back from. A shelter nobody had wintered in. The upside of being right about any of those was a slightly better week. The downside of being wrong was that you were dead, and so, often enough, was everybody who followed you.

That's an asymmetry, and it ran in one direction for a very long time. When the penalty for novelty is that severe and that permanent, "do what worked last time" isn't laziness. It's the highest-value strategy on the board, and the people who ran it are the reason any of us are here to argue about homepages.

So recognition became a signal, and a genuinely good one. **If you recognize it, somebody survived it.**

That's the machinery. It was built for a world where familiarity was real evidence, and it's still running at full power in a job where familiarity is evidence of nothing except that eleven competitors got there first.

Now, in fairness, and I'd rather say this myself than have you catch me at it: the instinct isn't actually a Neanderthal trait. It isn't in that two percent. It's human, it lives in the other ninety-eight, and it would be sitting inside every one of us if the two percent had never happened.

I just like having somebody to blame. He's a useful character and I'm keeping him.

///

## WHAT HE ACTUALLY DOES

The instinct has a modern name and it isn't a metaphor. It's one of the most replicated findings in social psychology, and once you know it you'll see it operating in every meeting you attend for the rest of your career.

**Processing fluency.** The easier something is for your brain to process, the more you like it, and the more likely you are to believe that it's true.

Not "the more familiar it feels." The more likely you are to believe it's **true.**

The original demonstration is Robert Zajonc in 1968. Show somebody a shape they've never seen and they're neutral on it. Show them the same shape several times and they report liking it more. They can't tell you why. Ask them and they'll say it's a nicer shape. The exposure did the work; the explanation got invented afterward to cover for it.

The effect doesn't stop at preference. It leaks into judgment:

1. A statement set in an easy-to-read typeface is rated more likely to be true than the identical statement set in a harder one.
2. A claim you have encountered before feels more accurate than one you haven't, whether or not it's accurate, and repetition strengthens the effect.
3. Statements that rhyme are judged more accurate than non-rhyming statements that mean exactly the same thing.
4. Companies with easier-to-pronounce names have been shown to be judged as better bets than companies with harder ones.

Nobody in any of those studies was fooled by a bad argument. They weren't evaluating an argument at all. Their brain noticed that the input went down smoothly, filed the smoothness under "this is correct," and handed the conclusion upstairs to the part of them that believes it makes decisions.

The largest single driver of fluency is repetition. Which means every time a sentence circulates inside your category, it becomes more true-sounding to everybody inside your category. Including your buyers. Including your CEO. Including you.

///

## SO PICTURE THE MEETING

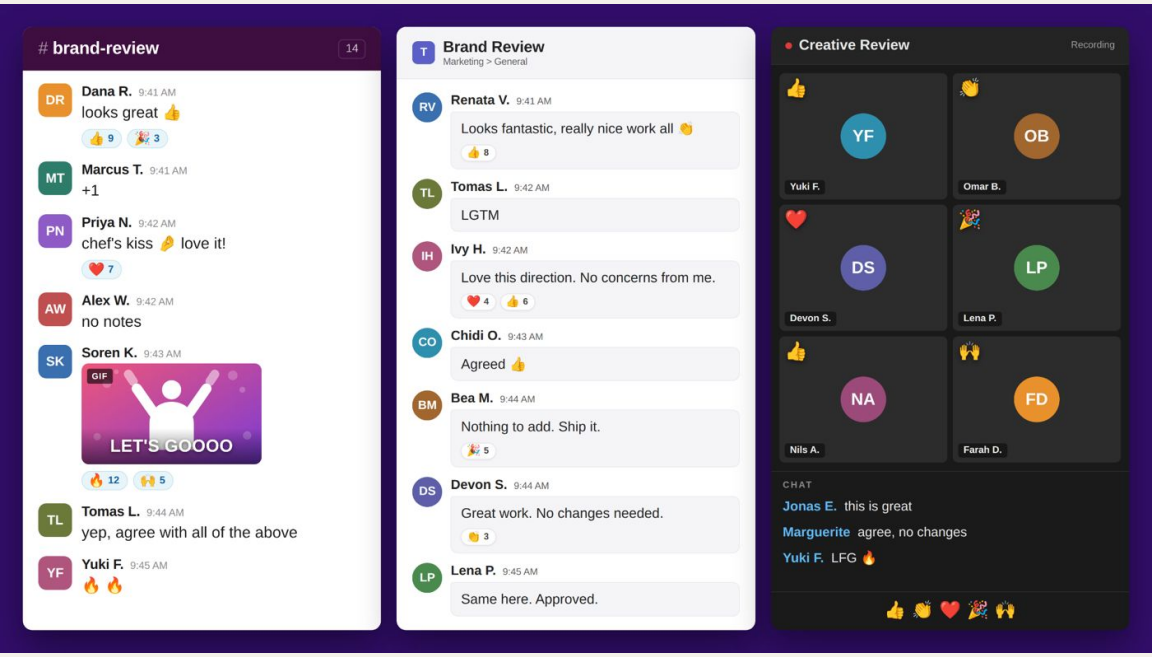



**FIG. 1.3** The review thread. Looks great, LGTM, approved. Every one of those is a nod.

It's a Tuesday morning. Somebody presents work.

It looks like the category. It uses the words the category uses. It has the structure every deck in this space has had for six years, which is precisely the structure everybody in the room has seen the most, which is precisely why it feels correct.

And everybody nods.

The nod feels like validation. It feels like proof. Nine smart people, several of whom can't agree on anything else, independently arrived at yes.

It isn't proof. It's him. It's your Neanderthal at the table, recognizing a shape he has seen before and reporting back that it's safe, and that report arrives dressed as a feeling of correctness, and it gets there before the thinking does.

He isn't wrong that the work is familiar. He is answering a completely different question than the one you asked.

So here's the distinction I'd like you to carry out of this chapter. When work feels **right**, trust that. That's judgment, and it usually arrives with a reason you can say out loud. When work feels **fine**, when it looks like everything else and nobody objects, that's him. Right is a signal worth trusting. Fine is the red flag.

Here's the version of this I want you to carry around, because it's the most useful sentence in the first half of this book:

> **WHEN A ROOM AGREES WITH YOUR WORK IMMEDIATELY, THAT'S INFORMATION. IT'S JUST NOT THE INFORMATION YOU THINK IT IS.**

Immediate agreement usually means you made something familiar. And familiar is the single property you can't afford in a market where everybody has the same tools, the same playbooks, the same conference circuit, and now the same models.

Psychologists call the skill you need here metacognition: noticing your own thinking while it's happening. I call it catching him in the act. You don't have to stop the nod. You just have to notice it, which is harder than it sounds and gets easier every time you do it.

///

## SOMETIMES THE ROOM IS JUST RIGHT

I want to be careful here, because the lazy version of this chapter turns you into somebody who treats every agreement as a red flag and every objection as insight. That person is exhausting to work with and also frequently wrong.

Rooms agree for good reasons all the time. The work is clear. The argument holds. Somebody finally said the thing everybody has been circling for a month. That happens, and when it happens the agreement is real and you should take it.

Three questions separate the two, and you can run them silently in about fifteen seconds.

**One. Could they have predicted it?** If you had described the assignment to this room a week ago and asked them to guess what would come back, would this be close? Agreement with your own prediction isn't evaluation. It's recognition.

**Two. Did anybody change their mind about anything?** Real agreement among adults usually involves at least one person visibly moving. If nobody moved, nobody was persuaded, which means nothing was actually being decided.

**Three. Can they each say why, differently?** Ask two people separately to explain what's good about it. Genuine agreement produces different reasons for the same conclusion. Fluency produces the same reason in different words, and that reason is usually a restatement of the work itself.

A room that agreed fast, predicted the answer, moved nobody, and can only justify itself by paraphrasing the deck hasn't evaluated anything. It has recognized something.

///

## THE DISSENT PROTOCOL

If the nod is the problem, the fix isn't to demand more criticism. Asking a room to be more critical produces performed criticism, which is worse than silence because it looks like rigor and consumes the same amount of time.

The fix is to change the order of operations so the fluency reaction never gets to speak first. Five rules. None cost anything. None require you to be in charge.

**1. Written before spoken.** Nobody talks for the first four minutes. Everyone writes their reaction privately, on paper or in a document nobody else can see yet. The moment one person says "I like it" out loud, the room's fluency and that person's seniority have fused into a single signal, and you won't be able to pull them apart again for the rest of the meeting.

**2. Score, don't react.** Ask for numbers, not opinions. Two of them: how clear is this, one to five, and how different is this from what our three closest competitors would produce, one to five. High clarity paired with low difference is the danger zone, and it's where most approved work in this industry lives. You want that split visible on the page before anybody has a chance to rationalize it.

**3. Make dissent cheap, but don't fake it.** Rotate one named person each week whose job in that meeting is to raise the strongest case against. It lowers the social cost of disagreeing, which is the main reason dissent doesn't happen on its own. But know what it can't do. Charlan Nemeth at Berkeley has studied dissent in groups for decades, and her finding is that a genuine dissenter makes the whole group think better, even when the dissenter turns out to be wrong, while an assigned devil's advocate mostly doesn't produce the same effect. The rotation isn't the dissent. It's cover, so the person who actually disagrees can say so without having to be the difficult one.

**4. Ask the competitor question.** If our closest competitor saw this tomorrow, what would they think? The useful answers are the bad ones. If the honest answer is "nothing, they wouldn't notice," you have your finding. The question works because it moves the room from judging quality to judging distinctiveness, and people conflate those two constantly.

**5. Find the one who didn't nod.** This is the most important rule in the chapter and it's the one I personally failed on the single most expensive mistake of my career, which is Chapter 8.

In almost every room there's one person who doesn't nod. They go quiet, or they ask a narrow technical question that seems beside the point, or they raise something easy to dismiss because they lack the full context. Spend the meeting there. Not because dissenters are smarter, but because the nod is the majority reaction and the majority reaction is the one you already

know. The only new information in the room is sitting with the person who didn't produce it.

And notice the specific trap, because it has a shape and it will come for you: **the more people are nodding, the easier it is to discount the one who isn't.** The volume of agreement reads as evidence against the holdout. That's the same mechanism firing twice, once to make the work feel right and once to make the objection feel wrong.

///

## WHAT TO DO TOMORROW

Take the last thing your team approved. Not something in flight, something already shipped, so nobody has to defend it.

Ask three people who were in that approval meeting one question, in writing, separately, without discussing it first: **what changed between the first version and the approved version, and why?**

If the honest answer is a list of things that got clearer, smoother, shorter, and more consistent with how we normally do it, then you have just watched the process work exactly as designed. You have also found the precise place your distinctiveness went.

That isn't an argument for burning down your review process. It's an argument for putting one thing in front of it: a written record of what made the work different in the first place, so that in round three somebody can hold up the list and ask whether any of it survived.

The review won't ask that on its own. It isn't built to. It's built to make things easier to process, and it's extremely good at its job.

He kept your ancestors alive. He is extremely good at his job too.

He is just terrible at making you interesting.

///

## THE COMFORT CHECK

CHAPTER 1: THE REVIEW MEETING

**1. The comfortable assumption.** ~~When the room agrees, the work is good.~~

**2. The true objective.** Find out whether this work would be noticed by somebody who didn't make it.

**3. The redesigned approach.** Written reactions before spoken ones, two scores instead of one opinion, a rotation that gives the real dissenter cover, and the meeting spent on the one person who didn't nod.

# 02

CHAPTER TWO

# THE SAMENESS MACHINE

**TWELVE DIFFERENT BRANDS. ONE SENTENCE.**

## Read this sentence and tell me whose it is.

> Blank uses agentic AI to triage, investigate and respond at machine speed, so your analysts can finally focus on the work that matters.

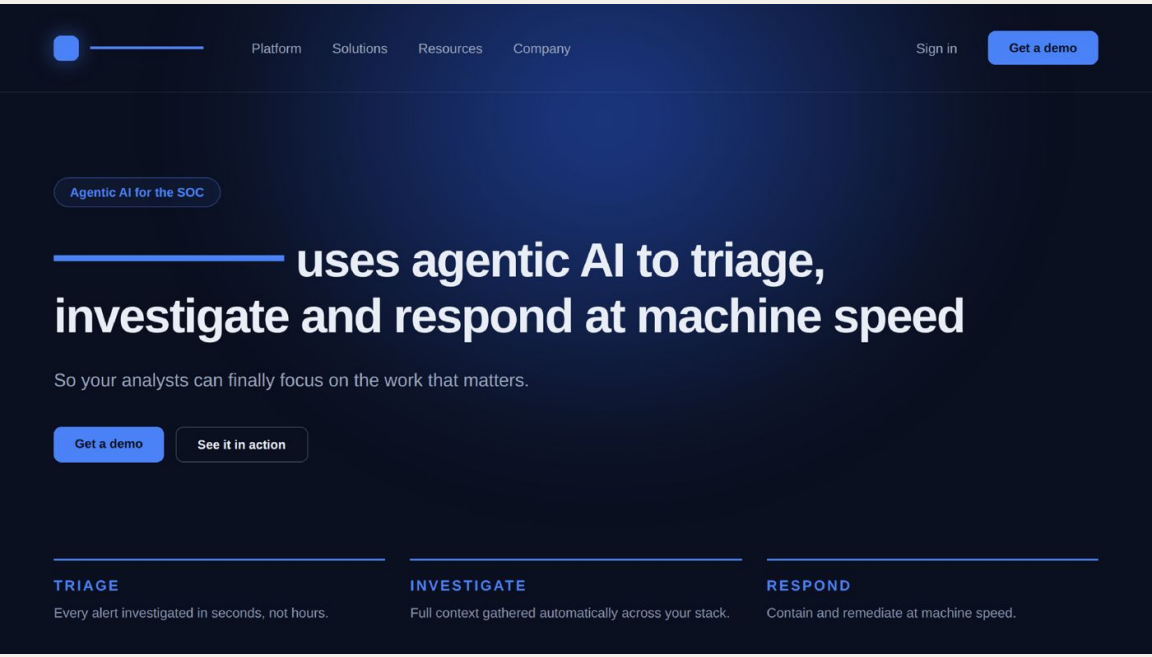



**FIG. 2.1** Whose is it? Nobody's.

If you work in security marketing, a few company names just went through your head. Maybe yours.

It's nobody's. A model wrote it, in about nine seconds. How, and why that matters, is Chapter 13. This chapter is about why you recognized it.

For the keynote I built a wall of twelve homepages around sentences like that one. The names and logos are invented, but they're built on what's out there. Look at how different they are. Different typefaces. Different colors. Serif, sans, all caps, lowercase. Light, dark, gradient, terminal. Somebody designed every one of them on purpose, and no two of them look alike.

Now read them.

Agentic. AI SOC. Machine speed. Triage, investigate, respond. So your analysts can focus on the work that matters.

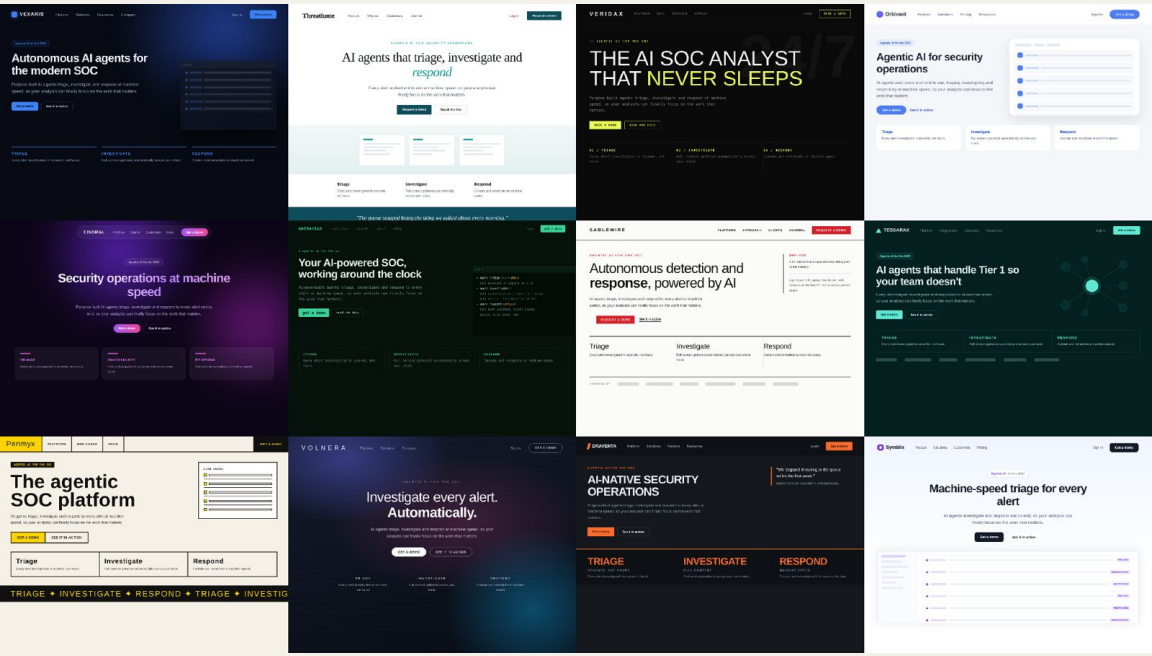



**FIG. 2.2** Twelve homepages. Invented names, invented logos, one sentence.

The real pages these are built on each had a different designer, a different agency, a different brand system, and a real budget. All of that money went into the part that's different. None of it went into the part that's the same.

If I swapped six of mine out for six real ones, some of them maybe yours, I don't think you could tell me which six.

I didn't put any real companies on the wall, and I'm not going to name any here. Most of them will be at the same conferences I go to, and naming them lets you file this as a fight between vendors instead of a description of a condition that includes me.

///

## THE PART WHERE I'M ON THE WALL

And before anybody feels singled out: go look at 7ai.com. You'll find at least three of these sentences on our homepage. I approved them.

I'd love to tell you I was the exception. I could defend every one of those sentences in a room for twenty minutes, and that's exactly the problem. Each one is reasonable. Together they're the category.

///

## MINE TOOK NINE SECONDS

Mine took nine seconds.

The real ones took six months. A brief. A creative review. A workshop with sticky notes on a wall. Three rounds of stakeholder feedback. And somebody senior who looked at the final version and said, yes, that's right, that's us.

And every one of them got the nod.

///

## NOBODY IN THIS PICTURE WAS BEING LAZY

This is the sentence I care most about in the first half of the book, so I want to spend a minute on it.

Nobody on those real teams was lazy. They cared. They stayed late. Somebody built three concepts and killed two. Somebody argued about whether the verb in the headline was strong enough, and somebody else was probably right about it.

Months of work. Real money. Real care.

And the output converged anyway.

If the sameness came from laziness, this would be a pamphlet instead of a book. Try harder. Care more. Hire better. Fire the agency. We'd all know

what to do and most of us would already be doing it, because the people in this industry aren't lazy and aren't stupid and in my experience care a great deal about work that most of the world finds incomprehensible.

**If effort were the answer, the problem would already be solved.**

So something else is producing this. Something that operates on people who are all trying hard, in a dozen different buildings, at the same time.

You met him in Chapter 1. Now let me show you the machine we built for him.

///

## YOUR REVIEW PROCESS IS A FLUENCY MACHINE WORKING PERFECTLY

Most marketing organizations have accidentally constructed a process that selects for familiarity and filters out everything else. Not through incompetence. Through completely reasonable governance that nobody would remove if you described it honestly.

Watch what a round of stakeholder feedback actually does.

The concept goes to eight people. Each of them is being asked, implicitly, to find something. Nobody forwards a review with "no notes" and feels like they contributed, and in a lot of companies nobody forwards a review with "no notes" and keeps being invited to reviews. So each person finds the thing that snags them.

And the thing that snags a reader is almost always the unfamiliar thing, because from the inside, unfamiliarity feels exactly like friction, and friction feels exactly like a problem.

So the notes come back. Every note is reasonable. Every note, independently, sands off one unfamiliar edge. Round two is smoother. Round two is also closer to the category mean, because the category mean is by definition the smoothest available version of any statement in your market. Run three rounds and you have arrived.

The process didn't fail. The process did precisely what it was built to do, which is remove friction.

Nobody ever told it that friction and distinction are the same substance.

**Every round of review is a vote for the most familiar available version, cast by people who sincerely believe they're voting for quality.** That's how twelve teams spend months and produce one website.

///

## IT DOESN'T STOP AT YOUR BUILDING

The review round explains how one company drifts to the middle. It doesn't explain why a whole category drifts to the same middle. For that you need a sociologist.

In 1983, Paul DiMaggio and Walter Powell published one of the most cited papers in organizational sociology, asking why organizations in the same field come to look so alike. One of the forces they named is **mimetic isomorphism**: when organizations are uncertain, they copy whoever looks successful. Nobody orders them to. Imitation is just the cheapest available answer to "what should we do?" when nobody actually knows.

Nobody in our category decided to look the same. Uncertainty did it for us. And I can't think of a category in more uncertainty right now than one where the product, the buyer's budget line, and the analysts' name for the market have all changed inside the same two years.

///

## WHY THIS COSTS MORE IN OUR CATEGORY THAN IN MOST

Three reasons, and they compound.

**The buyer is professionally skeptical.** You're selling to people whose entire job is to assume the thing in front of them isn't what it claims to be. A security buyer reads a vendor homepage the way they read a suspicious

email. When every vendor makes the same claim in the same words, the claim doesn't get evaluated and rejected. It gets discarded as category-level noise before evaluation starts, and you're underneath the noise floor with everybody else.

**The buying window is short and rare.** Nobody wakes up wanting to buy security software. Something happens: an incident, a board question, an audit finding, a renewal, a new CISO with a mandate. The window opens, a few names come to mind, the window closes. You aren't competing for a decision. You're competing to be one of the names that surfaces in a ten-minute conversation you aren't present for. That's a memory contest, and a homepage that reads like seven others has already lost it, months earlier, quietly.

**The evaluation is a committee, and your champion has to describe you.** Even when you get in the door, somebody inside that company has to explain you to four other people who weren't on the call. If your description is indistinguishable from the incumbent's, your champion has nothing to carry. You didn't lose the bake-off. You lost the hallway conversation after it, and nobody will ever tell you that's what happened.

Sameness in a consumer category costs you preference. Sameness in ours costs you recall, description, and repetition, which means it costs you the entire front of the process, invisibly, in a way that will never appear on a dashboard.

///

## THE SAMENESS AUDIT

Here's the first real tool in this book. Ninety minutes, no budget, and you should run it before you go much further, because everything after this lands differently once you've seen your own category laid flat.

You aren't looking for who's bad at this. You're looking for **what your category has silently agreed on**, because the agreements are the only things available to break.

### Step 1: Pick the set (10 minutes)

Eight to twelve companies. Not your competitive battlecard list, which is organized around who you lose to. Organize this list around **who a buyer encounters in the same week they encounter you**:

1. Three or four real head-to-head competitors.
2. Two or three adjacent vendors a buyer would lump you with even though you don't compete.
3. Two or three names from the last sponsor roster of the show you all attend.
4. Your own company. Not optional, and the entire reason the exercise works.

### Step 2: Capture (20 minutes)

Same browser, same window width, same zoom, same day, light mode wherever offered. Hero down to just below the fold, not full page. Consistency matters more than quality here, because you're comparing shapes and a reflowed responsive layout will lie to you.

### Step 3: Score four axes (40 minutes)

Companies as rows, these as columns.

**Vocabulary.** Write out every noun and verb phrase in each headline and subhead. Count how many companies use each. Anything appearing three or more times is category vocabulary, which means it's invisible. Anything appearing once is either a position or a mistake, and deciding which is your job.

**Structure.** Score presence or absence of: hero over image or gradient, one-to-two-sentence subhead, paired CTAs, logo bar within one scroll, three-column benefit section, stat band, testimonial carousel. Count the matches. Seven identical structural choices across ten companies isn't a trend, it's a template nobody wrote.

**Navigation.** List top-level nav items for each. Cheapest axis, most revealing, almost never run.

**Proof.** What does each company offer as evidence? Logos, a named case study, a number, an analyst mark, a certification, a product screenshot, a demo video, nothing. Note what almost everybody offers. Then note what almost nobody does, which is the more interesting column and the one people skip.

### Step 4: The swap test (10 minutes)

Take the six most similar heroes. Strip every logo and brand name. Print them, shuffle them, hand them to somebody in your company who doesn't work in marketing. Ask them to name the companies.

Don't do this to be clever. Do it because you need to watch a colleague fail at it in front of you, and because their failure is the only version of this argument that survives contact with a skeptical executive. A slide that says "we all sound alike" is an opinion. Your head of sales unable to identify his own company's homepage is a finding, and it will be repeated in that building for a year.

### Step 5: Write the sentence

Finish by writing down the single sentence your category has agreed on without ever voting on it. For the AI SOC cluster in 2026 it's roughly this:

> We use AI agents to handle the alert work your overwhelmed analysts can't get to, so they can do more valuable work.

That sentence is true. It's true of my company. It's probably true of yours. And it's worth nothing, because it's the noise floor, and every dollar spent saying it more beautifully than the last vendor is a dollar spent on a slightly more elegant way of being unidentifiable.

Once you've written it down, you have something to work against.

**You can't deliberately leave a place you haven't located.** That's the entire purpose of the exercise, and it's why the audit comes before any of the method in this book.

///

## WHAT THIS CHAPTER ISN'T SAYING

It isn't saying category vocabulary is a sin. You need some of it or nobody knows what shelf you're on, and a buyer searching for an AI SOC has to be able to discover that you're one.

It isn't saying the answer is strangeness for its own sake. I've made expensive strange things that worked and I've watched expensive strange things fail, and one of those failures gets a chapter of its own, in more detail than is comfortable.

What it's saying is narrower and harder to argue with. Your category has converged on a set of decisions nobody in it can remember making. Those decisions are now invisible to the people inside them, including you, including me, right up until somebody puts a dozen of them side by side.

The work of finding them is cheap. It's the finding that's expensive, and only in the sense that you can't un-see it afterward.

///

### THE COMFORT CHECK

CHAPTER 2: THE CATEGORY AUDIT

**1. The comfortable assumption.** ~~We know what our competitors say, because we track them.~~

**2. The true objective.** Find the decisions our whole category made without noticing, because those are the only ones available to break.

**3. The redesigned approach.** Capture eight to twelve heroes under identical conditions, score four axes including the navigation, run the swap test on somebody outside marketing, and write down the one sentence everybody has agreed on.

# 03

CHAPTER THREE

# SAFE IS JUST A BET YOU DON'T HAVE TO DEFEND

The thing I remember about asking the board for the money is that I wasn't asked whether the campaign would work.

I was asked what happens if it doesn't.

That's a completely reasonable question and I want to be fair to the people asking it. This was a very large number for a company of that age, spent on a thing nobody in our category had ever done. If it fell flat, that isn't a learning experience you write a thoughtful LinkedIn post about. That's a real problem, with real consequences, for people who had put real money behind us.

What I noticed, standing there, is that nobody had ever asked me that question about the alternative.

Nobody had ever sat me down and said: Nate, walk us through what happens if we spend the same money on the standard program, the trade shows and the paid search and the content engine and the six-week email nurture, and at the end of the year we're still the company that ninety-one percent of our market has never heard of. What's the plan then?

That question doesn't get asked. Not because the board is asleep, but because the standard program doesn't present itself as a decision. It presents itself as the absence of one.

Psychologists have a name for this. **Omission bias**: we judge harm caused by acting far more harshly than the same harm caused by doing nothing. The standard program never registers as an action, so its failures never register as anybody's fault.

///

## BOTH OF THESE ARE BETS

We talk about risk in this industry as though it's a thing you either take or you don't. You take a risk when you do something strange. You don't take a risk when you do the normal thing.

That isn't how it works and I think most people know it isn't how it works, which is why the belief survives by never being said out loud.

The normal thing is a bet. You're betting that being indistinguishable from eleven competitors will produce a result. You're betting that the buyer will pick you out of a lineup where everybody is dressed identically. You're betting that the reason last year's program underperformed was the market, and that this year's market will be different.

Those are substantive claims about the world. You can be wrong about them. People are wrong about them constantly, at enormous expense, for years at a time.

So the difference between the strange thing and the normal thing isn't that one is a bet and one isn't. Here's the honest version of the difference:

**Option A.** You do something different. It might not work. If it doesn't work, everyone will know exactly whose idea it was, because your name is on the only thing in the plan that looks like a decision.

**Option B.** You do what the category does. It probably won't work either. But if it doesn't work, nobody will be able to tell, because your results will look like everybody's results, and the story will be that the market was tough, that the buying cycle stretched, that budgets were frozen in Q3. All of which will also be true.

Read those again and notice that the difference has nothing to do with the odds.

It's a difference in **attributability.** Option B isn't safer. It's quieter when it fails.

That's the entire asymmetry, and once you can see it you can't stop seeing it. We pick the option that doesn't require us to stand up in a room and explain ourselves, and then we call that being responsible with the company's money.

///

## THE REAL COST OF THE QUIET FAILURE

There's a second-order problem here that took me most of my career to notice, and I think it's worse than the first-order one.

**A failure nobody can attribute is a failure nobody learns from.**

When the strange thing fails, you get an autopsy. Everybody in the building knows what was tried, they know roughly what it cost, and there's enormous organizational appetite to understand exactly what went wrong, mostly because several people want it on the record that they had concerns. That's uncomfortable. It's also the single most valuable byproduct of the whole exercise. I learned more from the failure in Chapter 8 than from either of the wins on either side of it, and it wasn't close.

When the safe thing fails, there's no autopsy, because there's no body. The numbers came in a bit light. The team is already planning next year, which will be structurally identical, with a slightly different allocation between two channels and a new agency. Nobody writes anything down, because nothing happened, and nothing happening isn't an event.

So the organization that consistently chooses Option B doesn't just underperform. It stops accumulating knowledge about its own market. Five years of that and you have a marketing team with a great deal of experience and almost no learning, running a program that nobody currently employed there can explain the original reasoning behind.

///

▶ SCAN TO WATCH

**WE HAVE TO TAKE BIG SWINGS**

From my 2022 CyberMarketingCon keynote.

## THE BET CARD

The point of this chapter isn't "take more risks." I'm not telling you to be reckless and I'm not telling you that strange is the same as good. I've made expensive strange things that worked, and I've been in the room for expensive strange things that didn't, and the difference between them was almost never the size of the swing.

The point is narrower: **know which bet you're making, say it out loud, write down what would prove it worked, and then make it on purpose instead of by default.**

Here's the instrument. One page, filled in before you spend anything. It works for a booth, a campaign, a report, a rebrand, a webinar series, a sponsorship, or a hire.

### THE BET CARD

**1. What we're betting.** One sentence, in plain language, with no marketing nouns in it. Not "increase awareness in the mid-market." Something closer to: "we're betting that a buyer who sees this once will be able to describe us to a colleague a month later."

**2. What we're betting against.** The alternative use of this money and this quarter, named specifically. If you can't name it, you haven't made a choice, you've made a purchase.

**3. What has to be true for this to work.** Two or three conditions, each one falsifiable. This is where most bets quietly die, because writing down the conditions is what reveals that one of them is already false.

**4. What would prove it worked.** A number, a source, and a date. The number must be measurable by somebody other than the person who wants it to have worked.

**5. What would prove it failed.** Written at the same time as number four, and this is the one people skip. A bet with no failure condition isn't a bet, it's a purchase you intend to describe favorably later.

**6. When we look.** A date. Put it in the calendar now, with the people who will be in the room.

**7. Who decides.** One name. Not a committee, and not "we'll discuss it."

Two rules about the card.

**Fill in five before four.** If you write the success condition first, you'll write the failure condition as its inverse, and you'll have learned nothing. Write down what failure looks like as its own idea, from scratch, and you'll often discover you have described the most likely outcome.

**Do it before, never after.** A bet card written after launch is a press release. The entire value of the thing is that it's a commitment made while you're still capable of being wrong.

///

## HOW TO DEFEND A BET WITHOUT CLAIMING IT'S SAFE

Back to the board.

The thing I didn't do, and I want to be specific because I think this is the transferable part, is argue that the campaign was safe. It wasn't safe. Everybody in that room could see that it wasn't safe, and any attempt to dress it up would have cost me the only thing I actually had, which was credibility about my own assessment.

What I argued was three things:

1. This is the best available strategy for the problem we actually have, which is that ninety-one percent of our market doesn't know we exist and our sales motion can't fix that by working harder.
2. The alternative isn't safety. The alternative is spending a comparable amount to remain in the ninety-one percent, and I'm willing to describe what that looks like in detail.

3. We'll know either way. Here's the measurement, here's the baseline, here's when we take it again, and here's the number that means I was wrong.

That third one is what actually moved the room, and it's the one most marketers skip. An executive's resistance to an unusual idea is very often not resistance to the idea. It's resistance to **unfalsifiable spending.** They have been handed unfalsifiable spending before. They have watched a year of it produce a slide that says the brand is stronger.

When you walk in with a defined failure condition, you have changed what you're asking for. You're no longer asking them to believe you. You're asking them to fund an experiment with a stated cost and a stated verdict date, which is a fundamentally easier thing to say yes to, and which also happens to be the honest description of what you're doing.

One warning. Don't offer a failure condition you aren't prepared to honor. If you say you'll kill it and then you don't kill it, you've spent the instrument, and you won't get it back at that company.

///

## THE VERSION OF THIS WITH NO BUDGET

Everything from here to the end of the book has a free version, and I'm going to list it at the end of each chapter, because the most common objection I get to this material is that it's easy to talk about swings when somebody else has already given you the budget to take them.

That objection is fair about the examples and wrong about the method.

Here's the free version of this chapter. You don't need money to make a bet. You need something at stake, and attention is a thing you already control.

1. **Bet the format, not the budget.** Next time you owe somebody a deck, send a one-page memo instead. Same information, different vehicle, zero dollars, and you'll find out inside a week whether your organization rewards clarity or ceremony.

2. **Bet the words.** Take the one comparative claim your category leans on hardest, the one everybody makes, and remove it from your next campaign entirely. No "faster," no "smarter," no "more accurate." You'll have to find something true to say instead. That's the exercise.
3. **Bet on the people.** Find the piece of work in your organization that somebody senior called too weird and that you privately thought was the most interesting thing you had seen all year. Put your name next to it. That costs nothing and it's a real bet, because it's attributable, which is the whole thing.

Each of those has a failure condition you can write on a card this afternoon.

And each of them is a bet you'll have to defend, which is exactly why they're worth making. The other kind was never free either. You were just not being billed for it in a way you could see.

///

## THE COMFORT CHECK

CHAPTER 3: THE ANNUAL PLAN

**1. The comfortable assumption.** ~~The standard program is the low-risk option, and the unusual idea is the bet.~~

**2. The true objective.** Make a decision we can be proven wrong about, on a date we set in advance.

**3. The redesigned approach.** A Bet Card filled in before the money moves, with the failure condition written before the success condition, one named decision-maker, and a review date already in the calendar.

II

PART TWO

# THE COMFORT CHECK

# 04

CHAPTER FOUR

# THE THREE QUESTIONS

## I don't like frameworks.

I've sat through the four-quadrant matrix, the five forces, the pyramid with the triangle inside it, the wheel with eight spokes and a verb on each one. I have also built a few, which is how I know what they're usually for. Most marketing frameworks aren't instruments for thinking. They're instruments for presenting, and the giveaway is that they're almost always beautiful. You can tell a framework was designed to be projected because somebody spent time on the icons.

A framework you present is a framework that requires an audience, a meeting, and a slot on somebody's calendar. Which means it gets used roughly four times a year, at planning, by the people who already had the authority to change things.

So what follows isn't a framework. It's a check, in the sense that a pilot runs a check: fast, unglamorous, the same every time, and run before the thing rather than about the thing.

You met it in the front matter. Here it's again, because it's the spine of the rest of this book and you should be sick of it by the end.

### THE COMFORT CHECK

**1. What's the comfortable assumption?**

**2. What's the true objective?**

**3. What's the redesigned approach?**

Three questions. It fits on an index card. It takes four minutes. You can run it silently, in a meeting, without telling anybody you're running it.

///

## WHY THREE

Because four is where a tool stops getting used.

I'm not being cute about that. Anything you can hold in your head while somebody else is talking gets used constantly, and anything that requires you to open a document gets used at planning offsites. The entire value of this thing is that it's available at the moment the comfortable choice is being made, which is usually a Tuesday, in a conversation nobody scheduled, when somebody says "we'll just do what we did last year" and everybody moves on to the next agenda item.

There's a longer version of this method. I use it myself and I'll give it to you: **observe, interpret, simplify, amplify, measure.** Five steps, genuinely useful, and it describes what actually happened across most of the campaigns in Part Three.

I keep it off the index card on purpose. Five steps is a process, and a process needs a project, and a project needs a budget. The three questions need a pen.

///

## WHY IT'S A CHECK AND NOT A REVIEW

The distinction that matters most: **run it before, not after.**

A Comfort Check run after the work exists is a critique, and critique arrives too late to do anything but wound. By the time there's a deck, somebody has spent three weeks on it, three other people have given notes, and the assumption you're about to question is now load-bearing for everybody's calendar. You won't win that argument, and you shouldn't, because the cost of being right at that point exceeds the value of it.

Run before the work exists, the same three questions cost nothing, because nobody has anything invested in the answer yet. Question one, asked at the start of a project, is a conversation. Asked at the end, it's an accusation.

The practical version: the check belongs in the fifteen minutes after somebody says "we should do X" and before anybody opens a file.

///

## WHY IT STARTS WITH THE ASSUMPTION

The order isn't arbitrary. In 1984, Charles Lord, Mark Lepper and Elizabeth Preston at Stanford tested ways to get people past their own biases when they weighed evidence. Telling people to be fair and unbiased helped far less than you'd think. What worked better was one specific instruction: consider the opposite. Ask how you'd judge the same evidence if it had pointed the other way.

That's question one. Naming the comfortable assumption is what makes its opposite thinkable, and you can't consider the opposite of something you never wrote down.

///

## THE THREE WAYS PEOPLE GET IT WRONG

I've watched this run badly enough times to describe the failure modes precisely.

**Failure one: answering question one with a strategy.**

Ask a team for their comfortable assumption and they'll tell you their plan. "We assume that our buyer is the CISO." That's a targeting decision. It might be wrong, but it's a decision somebody made out loud, in a document, and can defend.

A comfortable assumption isn't a decision. It's the thing underneath the decision that nobody voted on. Not "we assume our buyer is the CISO" but "we assume the way to reach a CISO is content." Not "we sponsor Black Hat" but "we assume a presence at the show is how a company our size becomes known."

The test: **if you can name the meeting where it was decided, it isn't a comfortable assumption.** You're looking for the ones with no meeting.

**Failure two: answering question two with an activity wearing a suit.**

"Our true objective is to generate qualified pipeline." That's an activity with a number attached to it. Pipeline is a thing your organization produces, not a thing the market experiences.

The tell is in Chapter 6 and it's this: if you could achieve it in full while the business got worse, it's an activity. You can generate a great deal of qualified pipeline while remaining a company nobody can describe.

**Failure three: answering question three with the thing you already wanted to do.**

This is the most common and the hardest to catch in yourself, because it feels like the check worked. You had an idea before you started. You ran three questions. Your idea came out the other end. What a useful tool.

The protection is mechanical: **question three has to produce at least two answers, one of which you don't like.** If the redesigned approach is a single option and it's the option you walked in with, you haven't redesigned anything. You have laundered a preference through a process, which is a thing marketing departments are world-class at.

///

## RUNNING IT ON WORK THAT ISN'T YOURS

Most readers of this book will want to run this on somebody else's project, which is a delicate business, because questions one and two are both, in the wrong tone, an accusation of not having thought about it.

Three rules that keep you employed.

**1. Ask question two first, out loud, and question one only to yourself.** "What are we actually trying to make happen here?" is a helpful question from any seat in the room. "What are we assuming that nobody has examined?" from a junior person is a career event. Get the objective on the table, and the

assumption usually walks out on its own, because people can't state an outcome and defend an activity in the same breath for very long.

**2. Attribute the doubt to the work, never to the room.** Not "I think we're assuming X." Try "if somebody outside the company looked at this, would they be able to tell it apart from the obvious version?" Same question. Nobody has to be wrong for it to get answered.

**3. Bring the redesign, not just the question.** The person who asks question one and stops is a critic. The person who asks question one and arrives with two versions of question three is a contributor, and the difference in how those two people are received is enormous and completely unfair and you should plan around it rather than resent it.

///

## WHAT THE REST OF PART TWO IS

One chapter per question, in order, each one going considerably deeper than a chapter about a single question should reasonably need to.

Then a chapter where the check fails, on a project of mine that cost real money, because a tool you have only ever seen succeed is a tool you haven't been shown honestly.

///

## THE COMFORT CHECK

CHAPTER 4: FRAMEWORKS

**1. The comfortable assumption.** ~~A method has to be complete to be useful.~~

**2. The true objective.** Have something available at the exact moment a comfortable decision is being made, which is never during a planning session.

**3. The redesigned approach.** Three questions instead of five steps, sized to fit in your head, run before the work rather than about it.

# 05

CHAPTER FIVE

# QUESTION ONE: WHAT IS THE COMFORTABLE ASSUMPTION?

---

**IF YOU CAN NAME THE MEETING WHERE IT WAS DECIDED, IT ISN'T ONE.**

Bigger company, more devices, bigger problem, bigger budget.

That was the assumption at Axonius, and I held it as firmly as anybody. The product was asset management, which at the time meant knowing what you have and whether it's actually protected. So the logic ran straight downhill: an enterprise has more of everything, therefore more of the problem, therefore more willingness to pay, therefore sell enterprise.

We built the entire company around that sentence. Reps who knew twelve accounts each and could make a number on four deals. A sales cycle measured in quarters. Marketing built to support a small number of very large conversations.

I want to say clearly that this was a good assumption. Not a lazy one, not a stupid one. If you had put it in front of any experienced person in our industry, they would have nodded, and they would have been right to.

That's the whole problem with it.

**A comfortable assumption isn't a bad idea. It's a good idea that has stopped being examined.** The bad ideas get caught. Somebody always catches the bad ideas. What nobody catches is the reasonable belief that quietly became load-bearing for twenty other decisions while everyone's attention was elsewhere.

///

## WHY YOU CAN'T SEE THEM

Three reasons, and they stack.

**They were true when you adopted them.** Almost every comfortable assumption in your organization was correct at the moment somebody formed it. That's why it's still there. It has a track record. It has, in the most literal sense, worked. The question is never whether it was true. The question is whether the conditions that made it true are still the conditions you're operating in, and nobody checks that, because checking requires you to suspect something that has never let you down.

**Nobody ever said them out loud.** This is the structural part. A decision that got debated has a record, a rationale, and usually an owner who will defend it, which means it's available for review. An assumption that never got debated has none of those things. It got absorbed. It's in the org chart, in the budget lines, in the shape of the nav bar, in what the calendar looks like in March. You can't audit a belief that was never stated, which is why the first half of this chapter is about getting them stated.

**Questioning them is socially expensive.** If a belief has been operating unexamined for three years, then examining it implies that three years of work was built on something nobody checked, which implicates everybody in the room including the person asking. The tool in Chapter 4 exists mostly to make that cheap.

So here's the working definition, and the test that goes with it.

> A comfortable assumption is a belief your organization acts on, that nobody in the room could name the meeting for.
>
> **If you can name the meeting where it was decided, it isn't one.**

///

## THE ASSUMPTION INVENTORY

Four methods. Run all four and you'll have more than you can act on, which is the correct outcome. The point isn't to overturn them. Most of them are fine. The point is to have them written down somewhere so that when conditions change, you have a list to check instead of a feeling.

### 1. The "just how it works" sweep (30 minutes, with your team)

Ask everyone to write down, silently, five things about how your marketing operates that they would describe to a new hire with the phrase "that's just how it works here."

Silently is doing real work in that sentence. Said out loud, the first answer anchors the room and you get five variations of it.

Then read them out without attribution. The ones that appear on three or more lists are your operating assumptions, and the fact that everyone independently reached for the same ones means they're invisible rather than agreed.

### 2. The new-hire question (free, and you have about thirty days to use it)

Every person who joins your team arrives with a functioning ability to see your assumptions, and loses it in roughly a month. This is the single most underused asset in a marketing organization and it expires on a schedule.

So use it deliberately. In week three, ask: **what have we done since you got here that you wouldn't have done at your last company, and what did we not do that you expected?**

Write the answers down. Don't defend any of them in the moment, which is the hard part, because about half will be things you already considered and rejected for good reasons. Defending them in the meeting teaches the new person that the question was ceremonial and you'll never get an honest answer from them again.

### 3. Budget archaeology (an afternoon, once a year, before planning)

Take last year's budget, line by line. For each line above some threshold you care about, answer two questions in writing:

1. What was this originally for?
2. Who decided to fund it the first time, and is that person still here?

The lines where nobody can answer either question aren't investments. They're inherited furniture. In my experience this exercise finds between ten and twenty percent of a marketing budget that exists because it existed, and that money is the cheapest funding available for anything new, because it's already yours and nobody has to approve it.

### 4. The calendar audit (an hour)

Look at the twelve months. For every recurring thing on it, a report, a show, a campaign season, a webinar series, ask: **what made that the right interval?**

Annual threat reports are annual because the first one was, and because the word "annual" was in the title, and for no other reason I have ever been able to find. Quarterly campaigns are quarterly because the finance calendar is quarterly, which is a fact about accounting, not about buyers.

Anything on your calendar whose rhythm you can't justify is being set by something other than your market.

///

## FORTY COMFORTABLE ASSUMPTIONS IN CYBERSECURITY MARKETING

This is the reference section. You won't believe all of these are assumptions, and for your company several of them will turn out to be correct after examination, which is a fine outcome and the point of examining them.

Each one comes with the counter-question, because the assumption alone just makes you feel bad. The question is what gives you something to do.

### Brand and positioning

| THE COMFORTABLE ASSUMPTION | THE COUNTER-QUESTION |
|---|---|
| 1. We have to be in the category the analysts named. | What would we call this if no analyst had ever named it? |
| 2. Our differentiation is our technology. | Can a buyer perceive that difference before they buy? |
| 3. The homepage is for buyers. | Who actually lands there, and what did they come to confirm? |
| 4. We need to sound enterprise-ready. | Who decided that sounding like everybody else is what ready sounds like? |
| 5. The founders' origin story is our brand. | Does that story change what a buyer believes, or does it only explain us to ourselves? |

## Content

| THE COMFORTABLE ASSUMPTION | THE COUNTER-QUESTION |
|---|---|
| 6. The threat report is an annual thing. | What made twelve months the right interval, other than the last one being twelve months ago? |
| 7. More content produces more pipeline. | Which piece from last year can anybody on the team name without looking? |
| 8. Gating the asset is how we get the lead. | What did we actually gain, measured against who didn't download it? |
| 9. Thought leadership means a byline from an executive. | Who in our company do practitioners already read voluntarily? |
| 10. We should publish on the topics with search volume. | Who else is publishing on those, and how would a reader tell ours apart? |

## Demand generation

| THE COMFORTABLE ASSUMPTION | THE COUNTER-QUESTION |
|---|---|
| 11. A lead is the unit of work. | What's the first thing that happens after a lead exists, and is it any good? |
| 12. The webinar needs a thirty-minute presentation. | Would anybody attend if we only did the Q&A? |
| 13. We need a nurture sequence. | What would a real person say they got out of email four? |
| 14. Paid search on the category term is table stakes. | What are we paying to be the eleventh identical result for? |
| 15. MQL thresholds should be tuned quarterly. | When did we last ask a rep what a good first meeting actually feels like? |

## Events

| THE COMFORTABLE ASSUMPTION | THE COUNTER-QUESTION |
|---|---|
| 16. We need a booth at the show. | What's a booth for, and can anybody name something that came from last year's? |

| THE COMFORTABLE ASSUMPTION | THE COUNTER-QUESTION |
|---|---|
| 17. The booth needs a demo station. | What's the actual conversion event here, and does it require a screen? |
| 18. Swag drives traffic. | Traffic of whom? |
| 19. We should sponsor the party. | What does somebody remember about us the next morning? |
| 20. Field events are a regional sales motion. | Who owns whether anybody in that room ever buys anything? |

## Analyst relations and press

| THE COMFORTABLE ASSUMPTION | THE COUNTER-QUESTION |
|---|---|
| 21. The analyst needs a fifty-slide deck. | What do they need to be able to repeat to a client next week, unprompted? |
| 22. Being in the report is the goal. | What happens to a company that's in the report and can't be described? |
| 23. Press releases are how news travels. | Where did our buyer actually hear about us last time? |
| 24. We need an embargo. | Who does the embargo benefit, us or the outlet? |
| 25. Analyst relations is a CMO activity. | Who in our company can answer the hard technical question in the briefing? |

## Product marketing

| THE COMFORTABLE ASSUMPTION | THE COUNTER-QUESTION |
|---|---|
| 26. The messaging house has to be complete before we launch. | Which single sentence would we keep if we could keep only one? |
| 27. Feature parity closes deals. | When did we last lose on a feature rather than on a story? |
| 28. Battlecards should cover every competitor. | Which three actually show up, and what do reps say when we aren't listening? |
| 29. The launch is the moment. | What happens in week six? |

| THE COMFORTABLE ASSUMPTION | THE COUNTER-QUESTION |
|---|---|
| 30. We need a category. | Do we need a category, or do we need one sentence a buyer can repeat correctly? |

## Sales enablement

| THE COMFORTABLE ASSUMPTION | THE COUNTER-QUESTION |
|---|---|
| 31. Sales needs more content. | Which asset actually gets sent, and what do reps send instead of the ones we made? |
| 32. Training is an event. | What do reps do differently on the Tuesday after? |
| 33. The pitch deck should tell our whole story. | What does the buyer need to believe first, before any of it matters? |
| 34. Marketing's job ends at the handoff. | Who owns whether the first meeting is worth the buyer's time? |
| 35. Objection handling is a document. | Where does the objection actually get raised, and is anybody from marketing ever there? |

## Operations and measurement

| THE COMFORTABLE ASSUMPTION | THE COUNTER-QUESTION |
|---|---|
| 36. Attribution tells us what worked. | What decision would we make differently if the model changed tomorrow? |
| 37. We measure what the platform reports. | What would we measure if the platform didn't exist? |
| 38. Brand can't be measured. | What did we decide to measure before we started? |
| 39. The dashboard is the review. | When did a number on it last change a plan? |
| 40. More data will settle this. | Is this a data problem, or a disagreement about the objective wearing a data costume? |

///

## BACK TO AXONIUS

The assumption held for a while and then the market started arguing with it.

Mid-market companies, roughly one to five thousand employees, began coming to us inbound. We weren't calling them. They found us. They closed inside a quarter, and this is the detail that matters more than the volume: **they closed without asking for a discount.**

Anybody can generate inbound. Inbound tells you your content works. Closing at list price is the market telling you it wants what you have badly enough not to negotiate, and that's a different signal entirely.

Nobody planned for that. It just kept happening, until it stopped being an anomaly and started being data.

Here's what a signal like that doesn't tell you. It doesn't tell you to sell harder to the mid-market. That's the small version of the conclusion, and it's the one an organization reaches naturally because it requires changing nothing.

It tells you your model is wrong.

An enterprise motion is built to reach twelve people extraordinarily well. A commercial motion has to reach ten thousand people well enough. Those aren't the same machine with a different target list. They have different economics, different content, different sales roles, and critically, a different constraint.

Which is what the signal was actually about. **The constraint had moved.** The question was no longer whether our reps could close. It was whether anybody knew we existed.

That's question one doing its job. Not producing an answer, producing a better question. What Chapter 6 does with that better question is where the campaign came from.

///

## THE COMFORT CHECK

CHAPTER 5: THE AXONIUS GO-TO-MARKET

**1. The comfortable assumption.** ~~Bigger company, more devices, bigger problem, bigger budget. Sell enterprise.~~

**2. The true objective.** Get known by a market that was already coming to us, at a scale an enterprise sales motion can't reach.

**3. The redesigned approach.** Continued in Chapter 6.

# 06

CHAPTER SIX

# QUESTION TWO: WHAT IS THE TRUE OBJECTIVE?

**IF YOU COULD ACHIEVE IT COMPLETELY WHILE THE BUSINESS GOT WORSE, IT'S AN ACTIVITY.**

I've sat in a marketing review where every number on the board was green and the company was invisible.

Content shipped on schedule. MQLs above target. Webinar registration up year over year. Site traffic up. Event program delivered on budget. Nine slides, every one of them honest, and not one person in that room could have told you what a buyer in our market would say about us if you stopped them in a hallway and asked.

Nobody was lying. Nobody was even spinning. Every one of those numbers was a real measurement of a real thing that really happened.

They were just all measurements of activity, and the company's problem wasn't an activity problem.

///

## THE TELL

Here's the test, and it's the most useful single sentence in this chapter:

Run it on anything.

Generate qualified pipeline. You can generate an enormous amount of qualified pipeline at a company that's losing ground, because pipeline is a measure of how much you put in the top, not of whether anybody chooses you at the bottom. Activity.

Publish twelve pieces of content this quarter. Obviously an activity, and it doesn't even resist the test.

Have a presence at the show. Activity, and a particularly expensive one, because "presence" is a word that exists specifically so that nobody has to define the outcome.

Be one of the three companies a CISO names when a peer asks who to look at. Not an activity. You can't do that while the business gets worse, because it's a description of a state of the world outside your building. That's an objective.

Notice what the good one has that the others don't. It describes something that happens **in somebody else's head, in a room you aren't in.** Almost every real marketing objective has that shape, and almost no marketing metric does.

///

## WHY ACTIVITIES WIN ANYWAY

Not because people are foolish. Because activities are better in three specific ways that matter enormously to the person who has to survive the quarter.

1. **Activities are completable.** You can finish them. There's a satisfying moment when the thing ships. Outcomes don't finish, they trend, and trending is a terrible experience to be evaluated on.
2. **Activities are attributable.** My team did this. The activity has your name on it in a way that "the market thinks about us differently" never will, and if you're three levels down in a large organization, attributable is how you get promoted.
3. **Activities are schedulable.** You can put an activity in a plan in October and know exactly what March looks like. You can't schedule the market changing its mind.

So the pull toward activities isn't laziness either. It's a rational response to how people are actually evaluated. Which means the fix isn't to shame anybody for tracking activities. The fix is to make sure there's an outcome sitting above them that the activities are supposed to produce, and to check periodically whether they're producing it.

That's the whole job of question two. Not to ban activity metrics. To stop them from being the top of the stack.

///



▶ SCAN TO WATCH

**KNOW WHAT YOU WANT**

From my 2022 CyberMarketingCon keynote.

## THE DEFINITION OF DONE REWRITE

The exercise. Take what each function in your organization says it delivers, and rewrite it as the outcome it was supposed to produce.

Then look at the third column, because that's where this stops being a semantic exercise and starts costing something.

| FUNCTION | WHAT IT SAYS IT DELIVERS | THE OUTCOME IT WAS SUPPOSED TO PRODUCE | WHAT CHANGES IF YOU ADOPT THE SECOND ONE |
|---|---|---|---|
| **Demand generation** | Qualified leads, delivered to sales | A first conversation that's worth the buyer's time | You now care about what happens after the handoff, which means you need to see the meeting, which means you need a relationship with sales you may not currently have |
| **Content** | Published assets on a calendar | A buyer can explain what you do to a colleague | You measure repetition, not consumption, and most of your calendar turns out to be producing neither |
| **Brand** | A refreshed identity, a campaign, a look | People recall and can describe you, months after exposure | You have to baseline before you spend and measure after you stop, which is Chapter 9 |
| **Events** | Booth, meetings booked, scans | Somebody remembers a specific thing about you on the flight home | The demo station becomes negotiable, and the budget moves toward whatever produces the memory |

| FUNCTION | WHAT IT SAYS IT DELIVERS | THE OUTCOME IT WAS SUPPOSED TO PRODUCE | WHAT CHANGES IF YOU ADOPT THE SECOND ONE |
|---|---|---|---|
| **Marketing ops** | Systems running, data clean, process followed | The rest of the organization can answer questions it couldn't answer before | Reporting stops being a deliverable and becomes an instrument, and some of the dashboards die |
| **Product marketing** | Messaging, launch, enablement assets | Sales says the same true sentence without reading it off a card | You test by listening to calls rather than by counting asset downloads |
| **Analyst relations** | Briefings held, reports appeared in | An analyst describes you accurately to a client when you aren't in the room | The fifty-slide deck becomes actively counterproductive, because nobody repeats slide thirty-one |

Every row in that third column costs something. That isn't a flaw in the exercise, it's the exercise. **An outcome you can adopt without changing anything was an activity in disguise.**

///

## HOW TO WRITE ONE SOMEBODY CAN FAIL

A real objective has four parts. Miss any one and you've written a slogan.

**1. A change, stated as a change.** Not a state, a delta. "Be known in the mid-market" is a wish. "Move from a market where nine in ten of our buyers have never heard of us to one where four in ten have" is a change. It has a direction and a distance.

**2. In somebody else's head or behavior.** Outside your building. If the objective can be satisfied entirely by things your team does, you've written an activity with ambition.

**3. A named measurement, chosen before you start.** Not "we'll figure out how to measure this." The measurement is part of the objective, and

choosing it afterward is how you end up with a metric selected for its willingness to agree with you.

**4. A date, and a person.** When we look, and who decides what the number meant.

Put together, it reads like a boring sentence, and that's correct. Objectives should be boring. The interesting part is question three.

> By the end of Q3, aided recall among security leaders at companies of one to five thousand employees moves from 9 percent to above 25 percent, measured by the same survey instrument we baselined with, and I make the call on whether that worked.

Somebody can fail that. That's the only property that matters.

///

## THE THING EVERYBODY DOES INSTEAD

Writes the objective after the plan.

You know how this goes because you've done it, and so have I. Somebody has an idea. The idea is good and everybody wants to do it. A deck gets built. Two slides from the end, somebody says we should add the success metrics, and a person who wasn't in the original conversation is asked to go find some numbers that this activity would plausibly move.

That isn't measurement. That's a costume.

The giveaway is that you can always tell which slide was written last, because it's the one where the metrics are all things the plan was already going to produce.

Question two has to be answered **before** question three or the whole check is theater. This is the single most common way I see this tool used badly, including by me, including recently.

///

## BACK TO AXONIUS

The signal from Chapter 5 said the constraint had moved. The old objective, which nobody had written down because it was structural rather than stated, was something like "help twelve enterprise reps work twelve accounts each extremely well."

The new objective had to describe the new constraint, and the new constraint was that a market that was already coming to us mostly didn't know we existed.

So: **get known by the mid-market, at a scale an enterprise motion can't reach, and know whether it happened.**

That last clause is a third of the objective and it's the part that made the rest of it possible to fund. I'll come back to it in Chapter 9 in detail, because "know whether it happened" turned out to be the hardest and most valuable decision in the entire campaign, and it got made before a dollar was spent, which is the only time it can be made honestly.

Question three is what we did about it, and it started with an arbitrary rule I gave myself that made no sense at the time.

///

## THE COMFORT CHECK

CHAPTER 6: THE MARKETING REVIEW

**1. The comfortable assumption.** ~~If every number is green, marketing is working.~~

**2. The true objective.** Know whether anything we did changed what the market believes, rather than whether we completed what we planned.

**3. The redesigned approach.** One outcome at the top of the stack, stated as a change in somebody else's head, with its measurement chosen before the work starts and a named person who calls it.

# 07

CHAPTER SEVEN

# QUESTION THREE: WHAT IS THE REDESIGNED APPROACH?

---

**THE VEHICLE DOESN'T HAVE TO EXPLAIN THE IDEA. IT HAS TO EMBODY IT.**

I gave myself a rule that made no sense at the time.

Whatever we say has to fit in two words.

Not a tagline. Not a positioning statement. Not a headline with a subhead underneath doing the actual work. Two words that a person could hear once, at a conference, half paying attention, and repeat correctly to somebody else at lunch.

There was no research behind that. I didn't test it. It was arbitrary, and I want to be honest that it was arbitrary, because the useful thing about this chapter isn't the rule. It's what an arbitrary constraint does to a room.

The answer we landed on was **Controlling Complexity.**

Not asset management, which is what we sold. Asset management is a category, and a category is a shelf, and nobody has ever felt anything about a shelf. Complexity was the thing our customers were actually drowning in, and complexity has a property that categories don't: it never goes away. It only accumulates. Every acquisition, every new tool, every cloud migration, every remote worker, every year.

So we weren't selling an inventory. We were selling control over something that gets worse whether or not you do anything.

Two words got us there. A brief wouldn't have.

///

## WHY DESIGNING BACKWARD IS RARE

Question three sounds obvious. Throw out the process, build from the outcome. Everybody nods at that, and almost nobody does it.

The reason isn't discipline. It's that **the existing process is the path of least fluency resistance**, and by now you know what that means. The current approach is the version everybody in the building can already picture. It requires no explanation, no new vendors, no unfamiliar line items, and no meeting where somebody has to say a sentence they've never said before.

A redesigned approach requires all of those. Which means the redesign isn't competing with the old process on merit. It's competing against the old process plus the enormous hidden subsidy of already existing.

You can't argue your way past that subsidy. You have to remove the old process from the table before the thinking starts, which is what a constraint is for.

///

## CONSTRAINT AS AN INSTRUMENT

Here's the mechanism, stated plainly.

**An open brief produces the category mean.** Ask a talented team for "a campaign that communicates our value," and they'll produce something competent, defensible, and recognizable, because with no constraint the most available answer wins, and the most available answer is the one the whole category has already converged on.

**A constraint makes the available answer illegal.** That's all it does. It isn't a creativity technique, it's an elimination technique. Two words made every sentence we had ever written about ourselves unusable, which forced a search, and the search is where the work happens.

The constraint doesn't have to be correct. It has to be strict. I've never been able to defend "two words" as the right number on any grounds other than that it worked, and I've stopped trying.

### A menu that works

Pick one. Don't pick three, because a work with three constraints produces nothing and a team that resents you.

1. **Two words.** The whole value, in two. Brutal, and the one I reach for most.
2. **No product nouns.** Nothing in the output may name the thing you sell, the category, or a feature. Forces you to describe what changes for the person rather than what you ship.

3. **No comparatives.** No faster, smarter, more accurate, more integrated, better. Whatever survives is your actual position, and for most companies this constraint alone deletes eighty percent of the messaging document.
4. **Sound off.** It has to be understood with no audio and no copy. This is the constraint that produced the vehicle in the next section.
5. **One sentence, said by the buyer.** Write the sentence a customer would use to describe you to a peer. You aren't allowed to use any sentence you'd say about yourself.
6. **The default channel is banned.** Whatever you'd obviously do, you may not do. If your answer to every objective is a webinar, ban webinars for one cycle and watch what the team finds.
7. **One tenth.** Design it for a tenth of the budget. Not as an austerity measure, as a search. The tenth-scale version is frequently better and almost always more distinctive, because expensive is a substitute for interesting and removing the money removes the substitute.
8. **It has to survive a bar.** Somebody has to be able to repeat it, drunk, at a conference, three days later, to somebody who wasn't there.

///

## VEHICLE SELECTION

Once you have the idea, you have the harder problem, which is getting it into a large number of heads without twenty million dollars and four years.

At Axonius the answer didn't come from cybersecurity. It came from a video I couldn't stop thinking about.

Amy Bream is a CrossFit athlete, born with one leg. The video is her attempting a lift. She fails. She fails again. She gets emotional. She goes back. She makes it.

No voiceover. No explanation. Nothing telling you how to feel.

And every person who watched it understood it instantly.

That's when it clicked, and it's the most transferable thing in this chapter:

Our idea was that in the face of enormous complexity, the best don't just overcome it, they find a way to give back. That sentence came before the campaign and it generated the campaign. Then we needed something that *was* that sentence rather than something that described it.

Which is how we ended up reaching out to Amy, and then to Simone Biles, in a category that at that moment was buying stock photographs of men in hoodies.

The agency that made it with us, Ruckus, has the whole campaign on its case study page. Two of the films are worth your time.



▶ SCAN TO WATCH

**HUMANIZING CYBERSECURITY: THE CAMPAIGN FILM**

Axonius, Controlling Complexity.



▶ SCAN TO WATCH

**ADAPTATION: SIMONE BILES AND AMY BREAM**

Axonius, Controlling Complexity.



**FIG. 7.1** What the category was buying at the time. You've licensed at least one of these. So have I.

### How to actually find one

Three rules, in order of how often they get broken.

**1. Look outside your category. Always.** Every vehicle available inside cybersecurity has been used by cybersecurity, which means it arrives pre-loaded with fluency and can't produce the reaction you need. The hooded figure was a vehicle once. Sport, craft, music, food, architecture, nature documentaries, competitive anything: these are unexploited in our market for no reason other than that nobody here looks there.

**2. Test it with the sound off.** If the vehicle needs a voiceover to land, it isn't a vehicle, it's a container for your copy. The Amy Bream video works with the sound off. That isn't a coincidence, it's the qualification.

**3. Test whether it embodies or decorates.** Ask what happens if you remove the idea and keep the vehicle. If it still works fine, the vehicle was decoration and you have bought an association rather than built a story. A celebrity who

could be swapped for any other celebrity is decoration. A person whose actual life is the argument you're making is a vehicle.

///

## THE VISIBILITY TEST

Here's the one I didn't have when I needed it, and it's the reason the next chapter exists.

Before you commit, ask this:

> **IS THE DIFFERENT THING VISIBLE TO THE AUDIENCE, OR ONLY TO ME?**

The failure mode is specific and it isn't the one people expect. You don't usually get punished for doing something too strange. You get punished for doing something conventional that contains, somewhere inside it, one genuinely original element that nobody outside the planning room will ever perceive.

That project will feel different to you the entire time you're building it. You'll be able to explain, at length, exactly what makes it unlike anything else in the category. Every word of that explanation will be true. And the market will experience the conventional wrapper, because the conventional wrapper is the only part it can see.

The test is mechanical. Describe your plan to somebody outside your company in thirty seconds, the way a buyer would encounter it rather than the way you built it. Then ask them what's unusual about it.

If they can't name the thing you're proud of, the thing you're proud of isn't in the plan. It's in the deck.

I didn't run that test once, on a project that cost real money, and Chapter 8 is what happened.

///

## THE SAME OBJECTIVE AT THREE BUDGET LEVELS

The objection to everything in this chapter is that it's easy to talk about vehicles when somebody has already given you a budget for one. So here's the same objective, redesigned three ways.

**Objective:** a security buyer who has encountered us once can describe what we do to a colleague a month later, in their own words, correctly.

### With no money

1. **Pick the two words and enforce them everywhere for two quarters.** Every asset, every deck, every rep's LinkedIn headline, every conference bio, every email signature. Consistency is free and almost nobody does it, because internal boredom sets in roughly nine months before external recognition does.
2. **Build one analogy and use it relentlessly.** Chapter 10 is the method. The Camry cost nothing and outlived every paid campaign of that era.
3. **Change who speaks.** Put the person practitioners already trust in front of the audience instead of the person whose title says they should be. Free, and frequently prohibited by nothing except habit.
4. **Run the repeat-back test monthly.** Five people outside the company, one week later, what do we do. That's your measurement and it costs an hour.

### With a real but ordinary budget

1. Everything above, plus:
2. **One vehicle, sized to what you have.** Not a celebrity. A person, a place, a made object, a piece of research nobody else would bother to do. The qualification is the sound-off test, not the price.
3. **Concentrate rather than distribute.** The comfortable move with a moderate budget is to spread it across six channels so that no single line item looks reckless. That's how you buy invisibility in six places. One channel at a depth that registers beats six at a depth that doesn't.

4. **Fund the measurement.** A baseline study isn't the expensive part of anything in this list and it's what converts next year's conversation from an argument into a report.

### With a large budget

1. Everything above, plus:
2. **A vehicle at a scale that makes people ask what it is.** That's Chapter 11.
3. And a warning that belongs here rather than there: a large budget makes it possible to be expensive instead of interesting, and expensive is much easier. Most of the worst marketing I've seen in this industry was extremely well funded. The constraint list above doesn't stop being useful when you can afford to ignore it. It becomes the only thing keeping you honest.

///

## THE COMFORT CHECK

CHAPTERS 5 THROUGH 7: AXONIUS, COMPLETE

**1. The comfortable assumption.** ~~Bigger company, more devices, bigger problem, bigger budget. Sell enterprise.~~

**2. The true objective.** Get known, at a scale enterprise sales can't reach.

**3. The redesigned approach.** Two words, and a vehicle nobody in security would buy.

# 08

CHAPTER EIGHT

# WHEN THE CHECK FAILS: A POST-MORTEM

**IT WAS DIFFERENT IN THE DECK. IT WAS NEVER DIFFERENT IN PRACTICE.**

If I only show you the ones that worked, I'm not teaching you anything. I'm doing a victory lap in front of you and calling it a book.

So here's one from the same company, roughly the same period, same guy, same budget authority. It cost real money and it didn't work at all.

What makes it worth a chapter isn't that it failed. Things fail. It's that it failed for the exact opposite of the reason everybody assumes, including me at the time, and the correct diagnosis took me years to arrive at.

///

## WHAT I WANTED

More demos at the booth. More qualified people sitting down in front of the product.

That's a completely ordinary objective and I want to note, before anything else, that it passes the test in Chapter 6 only barely. "More qualified demos" is a measure of the top of something. It's closer to an activity than I'd write today. Hold that thought, because it comes back at the end.

## WHAT WE LANDED

Then something genuinely good happened. We landed the customer I had wanted for years.

A car company. I'm not going to name them, which, as you'll see, is the only part of this that worked.

## THE IDEA

Every person who came to our booth and sat through a real demo got one entry into a sweepstakes. The prize was the car.




**FIG. 8.1** The plan, as it looked in the deck.

And here's the part I didn't tell anybody outside the company, because this was the part I loved. At the end, we'd draw the winner live, on a webinar, from that company's headquarters in Germany.

Which meant the drawing was also, quietly, the best customer testimonial we'd ever get. A global manufacturer, on camera, at their own facility, participating in our campaign. Not a quote in a press release. A thing happening.

I thought it was the most elegant idea I had ever had. I'm aware of how this sentence reads given where the chapter is going.

///

## ED

Before we spent a dollar, we hired a lawyer named Ed.

Ed does one thing. Ed does sweepstakes. That's his entire practice, which I hadn't known was a viable way to make a living, and which should have told me something about how well-trodden this ground was.

In the first conversation, before I had finished explaining it, Ed said: you have to offer a cash alternative, that part isn't optional, and if you offer cash they'll take the cash.

They always take the cash.

I heard him. I want to be completely clear that I heard him, because the version of this story where I didn't listen is more flattering and less useful. I heard the man whose entire professional life is this specific thing tell me the specific thing that was going to happen.

And I decided he didn't understand how good my idea was.

///

## WHAT HAPPENED

We launched.

About a week later, the car company came back and said they couldn't be associated with it. After they had agreed to it. In writing.

So now there's no reveal. No webinar. No Germany. No testimonial. It's now a car giveaway at a trade show, which, if you have ever walked through a shopping mall, you've seen before.




**FIG. 8.2** What it became.

We ran it anyway, because the money was already gone.

The winner took the cash.

And never became a customer.

///

## THE EASY READ, WHICH IS WRONG

The easy read is: he did something weird and it didn't work.

I want to close that off hard, because if you take that lesson out of this chapter you'll have learned the opposite of the thing, and you'll be more careful for the next three years in exactly the way that costs you the most.

**It didn't fail because it was weird. It failed because it wasn't weird.**

A car giveaway is the most familiar promotion in the history of commerce. It's what a car dealership does. It's what a mall does on a Saturday. It's what

a radio station did in 1974. The prize was larger than usual and that's the entire extent of the innovation.

The one genuinely different thing in the whole plan was the live draw from their headquarters. And that element was invisible to everybody except me. It existed in a slide. In reality it was going to exist for one webinar, at the very end, in front of people who had already entered.

That's the sentence I wish somebody had said to me, and it's the visibility test from Chapter 7, which I built afterward, out of this.

///

## THE NOD, WITH A BODY ATTACHED

Now the part that connects to where this book started.

When I pitched it internally, everybody nodded. Immediately.

Of course they did. Everyone understood it in one second, because everyone has seen a car giveaway. There was no explaining to do. The concept required zero cognitive effort from any person in that room, and the room experienced that effortlessness as clarity, and reported it back to me as agreement.

The nod wasn't telling me the idea was good. The nod was telling me the idea was familiar.

It was the warning label and I read it as a review.

One person in the entire process didn't nod. Ed.

And I discounted Ed **specifically because everybody else was nodding.** That's the trap from Chapter 1 and I walked into it with my eyes open. The volume of agreement in the room made the single dissenting voice sound like an outlier rather than a signal. A room full of people who had never run a sweepstakes outvoted one man who had run hundreds, and I let them, because they were many and he was one.

### He votes twice

Years after the car, I watched a clip of Bill Belichick making a point I'd heard a hundred times and never really taken in. His point, roughly: you can't improve if you can't accept constructive criticism.

Every coach says some version of that, and everybody who hears it agrees, which by now should make you suspicious of how much work the agreement is doing.

The usual reading is that criticism is hard because it stings. Somebody is blunt, your ego takes a hit, you get defensive. That happens, but it isn't what happened with Ed. Ed was polite. He was specific. He was right. There was nothing in the delivery to resent.

It was hard to take because it felt wrong, and it felt wrong for exactly the same reason the giveaway felt right.

Go back to Chapter 1. Processing fluency makes the familiar idea feel correct. Run the same machinery the other way and it makes the unfamiliar objection feel like friction, and friction feels like a mistake on the other person's part. So he votes twice: once for the idea everybody recognizes, and once against the person who doesn't. Neither vote feels like a vote. Both feel like judgment.

That's why "be open to feedback" is useless as advice. Nobody believes they're closed to it. What you can actually do is notice the specific sensation, the small sense that an objection is beside the point, and treat that sensation as the thing to investigate rather than the conclusion.

It's also why the rotation rule in Chapter 1 comes with a warning. Nemeth's research says the real thing works and the role-play mostly doesn't. You can't fake Ed.

If the only feedback you can hear is the kind that agrees with you, you aren't getting feedback. You're getting a nod.

If you take nothing else from my very expensive afternoon:

> **THE ROOM AGREEING WITH YOU ISN'T THE SIGNAL. THE ONE PERSON WHO DOESN'T IS.**

///

## THE MECHANISM, WHICH ISN'T THE CAR COMPANY

Here's where most post-mortems stop, and where this one gets useful.

The obvious story is that a partner pulled out and wrecked the plan. Bad luck. Unforeseeable. We had it in writing.

That story is true and it's a decoy, because it lets the whole thing be somebody else's fault, which is exactly what a post-mortem is supposed to prevent.

So run the counterfactual honestly. Suppose the car company had held. Suppose we got the webinar, the headquarters, the live draw, the testimonial, all of it.

**The campaign still fails.**

Because the mechanic was always going to select the wrong people. A prize doesn't select for people who want your product. It selects for people who want the prize. Everybody who walked up to that booth and sat through a demo to get an entry was, by construction, someone whose reason for being there was the car. We built a machine for filling a room with people who wanted a car, and then we were disappointed that it filled the room with people who wanted a car.

The partner falling through cost us the one nice element. The design error cost us the campaign, and the design error was present on day one, in the pitch that everybody nodded at.

### The same error, in cheaper clothes

You're probably not giving away a car. You may well be doing one of these.

1. **The gift card for a thirty-minute call.** Selects for people who want a hundred dollars. Reliably produces meetings, reliably produces nothing after them, and reliably survives in the plan because meetings booked is a number that goes up.
2. **The high-value conference raffle.** Selects for badge scans. Your scan count will be the best in the hall and your pipeline won't notice.
3. **The gated report promoted to a cold list.** Selects for people willing to trade an email address for a PDF, which is a genuinely different population from people who have the problem the report is about. You didn't generate demand. You purchased a list at an unusual price.
4. **Swag that's better than your product story.** The line at your booth is measuring the socks.

The general form, and the question to run on any incentive before you launch it:

> **EVERY INCENTIVE SELECTS. WHAT DOES MINE SELECT FOR, AND IS THAT THE SAME POPULATION AS MY BUYER?**

If the answer is no, you aren't accelerating demand. You're paying to contaminate your own data, and the contamination will look like success for two quarters.

///

## THE CHECK, WITH QUESTION ONE BLANK

This is the only Comfort Check in this book where the first line doesn't get crossed out.

I never crossed it out. I never asked the question. The assumption sat there the entire time, unexamined, load-bearing, and completely invisible to me because it arrived wearing the clothes of an objective.

///



▶ SCAN TO WATCH

**FAILURE IS WHERE WE LEARN**

From my 2022 CyberMarketingCon keynote.

## THE POST-MORTEM TEMPLATE

What I use now. Seven questions, run in writing, before the meeting rather than during it, and with rule five taken seriously or the whole thing is theater.

1. **What did we predict, in writing, before we started?** If there's no written prediction, stop. You aren't running a post-mortem, you're constructing a narrative, and the narrative will be built out of whatever everybody remembers believing, which is always roughly what happened.
2. **What actually happened?** Facts and dates only. No causes yet.
3. **Where did one and two first diverge?** Find the earliest date, not the most dramatic one. The dramatic moment is almost never the divergence, it's the consequence becoming visible.
4. **What did we know at that point that we didn't act on?**
5. **Who said it, and why did we discount them?** Name the person. This question is the whole template. Every failed project I have ever examined had somebody who said the thing, and the interesting information is never what they said. It's the reason we found for not listening, because that reason is a property of the organization and it will be there next time. In my experience the reason has a shape, and it's usually the one from earlier in this chapter: the objection felt like friction, so it felt like the other person's mistake.

6. **What was the mechanism, as distinct from the event?** The event is what you'll tell people at the bar. The mechanism is the thing that would still have been true if the event had gone your way. If you can't separate these, you haven't learned anything you can carry.
7. **What would we have had to believe, on day one, to make the right call?** Not what should we have known. What belief would have produced the right decision. That belief is the thing to install, and it's usually one sentence.

For this project, question seven has an answer I use constantly now: **an idea everybody understands instantly is an idea everybody has already seen.**

///

## THE COMFORT CHECK

CHAPTER 8: THE DREAM CAR GIVEAWAY

**1. The comfortable assumption.** A big enough prize gets us the right people.

*This one isn't crossed out, because I never crossed it out. I never asked.*

**2. The true objective.** Qualified demos that turn into pipeline.

**3. The redesigned approach.** I never asked this one.

**And the question I never asked, which is the whole reason this chapter exists:** a prize doesn't select for people who want your product. It selects for people who want the prize.

III

PART THREE

# THE PRACTICE

# 09

CHAPTER NINE

# BETTER IS SUBJECTIVE, DIFFERENT IS MEASURABLE

---

**BETTER IS SUBJECTIVE. DIFFERENT IS MEASURABLE.**

People say you can't measure brand.

What they mean, almost every time, is that they didn't decide what to measure before they started.

That isn't a cheap shot. It's a description of the standard sequence. The campaign gets approved because somebody senior liked it. It runs. Two weeks before the board meeting, somebody is asked to find out whether it worked, and that person goes looking for numbers that moved during the window. They find some, because numbers always move. And the resulting slide isn't a measurement, it's a defense, and everyone in the room can tell, including the person presenting it.

The way out is unglamorous and it's entirely a matter of order of operations.

///

## WHAT WE ACTUALLY DID

Before we spent a dollar at Axonius, we ran a brand recall survey.

**Aided recall in our target market: nine percent.**

Which is a polite way of saying that ninety-one percent of the people we were trying to sell to didn't know we existed.

Then we ran the campaign. Then we waited. Then we ran the same survey again. Same instrument, same wording, same market definition.

**Six months after the campaign finished: thirty-seven percent.**

The six months is the part I care about, and it's the part that usually gets cut when this story is retold, because it's the least exciting number in the paragraph.

If you measure while the campaign is running, you're measuring your own media spend. Of course people have heard of you. You're on their screen. Measure half a year after you stopped, and you're measuring whether anybody actually retained anything.

**Say the timing before you say the number.** "Six months after it finished, thirty-seven" is an argument. "Nine to thirty-seven" is a statistic, and statistics invite the question of what else was going on, which in our case was a great deal, because the company was growing on several fronts at once and no clean causal isolate exists.

There's a second thing the six-month window does, and I didn't notice it for years. **Aided recall is a memory test.** Asking somebody half a year later whether they recognize a name is, literally, measuring whether they remembered. Which means this chapter and the next one are the same chapter viewed from two angles.

///

## SAY "AIDED," AND SAY IT FIRST

Aided recall is a lower bar than unaided. You show somebody a list and ask which names they know, rather than asking them to produce names from nothing.

Name it yourself, out loud, before anybody asks.

I've watched marketers get destroyed in a board meeting over exactly this, and never because the number was bad. It goes: awareness went from nine to thirty-seven. Somebody asks which metric. There's a pause. The pause is the whole event. Now the conversation is about whether you understand your own data, and the number is gone.

Nine to thirty-seven on aided recall is a very large move. It doesn't need help. Being precise about what it's costs you nothing and removes the only available attack.

///

## HOW TO RUN ONE

The operational version, for somebody who has never commissioned market research and is about to.

**1. Decide which question you're answering.** Recall, consideration, and attribute association are three different studies and people conflate them constantly. Recall asks whether they know you exist. Consideration asks whether they would put you on a shortlist. Attribute association asks what they think you're for. Most companies at the awareness stage of the problem need the first one and buy the third because it sounds more sophisticated.

**2. Define the population narrowly, in writing, before you talk to a vendor.** This is the part that gets rushed and it determines everything. Security decision-makers isn't a population. Security decision-makers with budget authority, at companies of one to five thousand employees, in North America, is a population. Write the screening criteria down and keep them identical across waves, because a sample that drifts will produce a lift that's entirely an artifact of who answered.

**3. Write the question once.** Then never touch it again, no matter how much better you could word it next time. The comparison is the deliverable, and a comparison across two different questions is worth nothing. Put the exact wording in a document with the date on it.

**4. Build the list properly.** An aided recall list includes you, your real competitors, a couple of large adjacent names as anchors, and **at least one company that doesn't exist.** The fake name is a quality control. A meaningful percentage of respondents in every panel will claim to recognize anything, and the fake name tells you what your over-claim rate is so you can discount accordingly. If your vendor doesn't suggest this, ask why not, and treat the answer as information about the vendor.

**5. Randomize the list order** across respondents, or position will do some of your work for you.

**6. Baseline before you spend anything.** Not before the campaign launches. Before the money moves. The first ad, the first sponsorship, the first pre-announcement to press, all of it contaminates the baseline, and a contaminated baseline is worse than no baseline because it looks legitimate.

**7. Re-measure six months after the spend stops.** Not during. Not at the end. Six months after.

**8. Don't let the agency that ran the campaign run the study.** This shouldn't need saying. It needs saying.

The cost of this is real but it's almost never the reason it doesn't happen. The reason it doesn't happen is that step six requires you to decide, before you have any evidence, what evidence would embarrass you.

///

## THE MEASUREMENT PLAN

One page, filled in alongside the Bet Card from Chapter 3.

> **Objective.** The change, in somebody else's head, stated as a delta.
>
> **Baseline.** The number today, with the date and the method.
>
> **Leading indicators.** Things that should move within weeks if this is working. In a memory campaign these are usually behavioral: branded search volume, direct traffic, inbound mentions of the specific idea rather than the company, sales reps reporting the phrase coming back at them.
>
> **Lagging indicator.** The one that actually settles it, and when you look.
>
> **Kill.** What number, on what date, means we stop.
>
> **Continue.** What number means we keep going as is.
>
> **Scale.** What number means we put more behind it.
>
> **Who calls it.** One name.

The three-way split at the bottom matters more than it looks. A plan with only a success condition and a failure condition forces a binary at the review, and most real results are neither. Naming the continue case in advance is what stops a merely slow result from being killed by somebody's impatience, and naming the scale case is what stops a genuine success from being congratulated and then defunded.

///

## FIVE MEMORY TESTS THAT COST NOTHING

If you can't fund a study, you can still measure memory. These aren't substitutes for a recall study at a company that can afford one. They're better than the thing most companies do instead, which is nothing followed by an argument.

1. **The repeat-back test.** Tell five people outside your company what you do. Wait a week. Ask them what you do. Whatever comes back is your message, whether or not it's the one you wrote. Run it monthly with different people and track how much of the answer is yours.
2. **The discovery call count.** Listen to twenty recorded first calls. Count how many prospects describe you, unprompted, using language from your messaging. Not the reps. The prospects. This is the cheapest and most brutal number in marketing and almost nobody looks at it.
3. **The open-text field.** Put "how did you hear about us" as a free-text box rather than a dropdown, and actually read the answers every month. A dropdown tells you what you already believed. The box tells you what happened.
4. **The hallway test.** At the next conference, ask ten people in your market if they've heard of your company, before you say anything else. Ten isn't a sample. It's still more contact with the truth than most marketing leaders have in a quarter.
5. **The anonymized swap.** Your homepage hero, logo removed, shown to ten people in your market. Ask whose it is. You built this in Chapter 2. Run it again after every significant messaging change, and treat "I don't know" as the finding it is.

///

## THE ACTUAL LESSON

I want to be careful, because the takeaway here isn't "hire Simone Biles." Almost nobody reading this can, and it wouldn't help most of the people who could.

Notice what we measured.

We didn't measure whether our brand got **better.** Better is an argument. You can't settle it, the people who disagree with you'll never be persuaded, and everybody in the room knows the person presenting has an interest in the answer.

We measured whether people knew who we were. That's a number. It has a before and an after and a method attached, and the person who disagrees with it has to attack the method rather than your taste.

That's the belief, and it's the one I'd defend hardest of the seven:

And if you don't have a Simone Biles budget, you don't need one. You need two things that cost nothing: a rule that makes the obvious answer illegal, and a number you write down before you start.

It also quietly resolves the argument this whole book is having. When you make something different, you aren't asking anybody to agree that it's superior. You're making a claim about whether it can be told apart from the alternatives, and that claim can be checked.

Which is the honest reason safe work is so hard to dislodge. Safe work is never wrong, because nobody ever proposed a test it could fail.

///

## THE COMFORT CHECK

CHAPTER 9: PROVING IT WORKED

**1. The comfortable assumption.** ~~Brand can't be measured, so we'll argue about it afterward.~~

**2. The true objective.** Know, on a date we set in advance, whether the market's belief about us changed.

**3. The redesigned approach.** One instrument, worded once, baselined before a dollar moves, re-run six months after the spend stops, with kill, continue, and scale numbers written down and one person who calls it.

# 10

CHAPTER TEN

# DESIGN FOR MEMORY, NOT JUST EXPLANATION

**WHAT WILL A JUDGE STILL BE ABLE TO REPEAT TOMORROW?**

Some of you have heard this one.

I've told it enough times that it's in my own conference bio, and the bio openly says you're probably rolling your eyes right now. So let me tell you the part I usually leave out, which is the part where I fail.

///

## THE SETUP

RSAC Innovation Sandbox. Ten finalists, three minutes each, live on stage, in front of a panel of judges who decide on the spot. For an early-stage security company it's the single biggest thing that can happen to you. Past finalists have collectively been acquired and funded at a scale that makes the three minutes worth roughly what a Super Bowl spot costs per second.

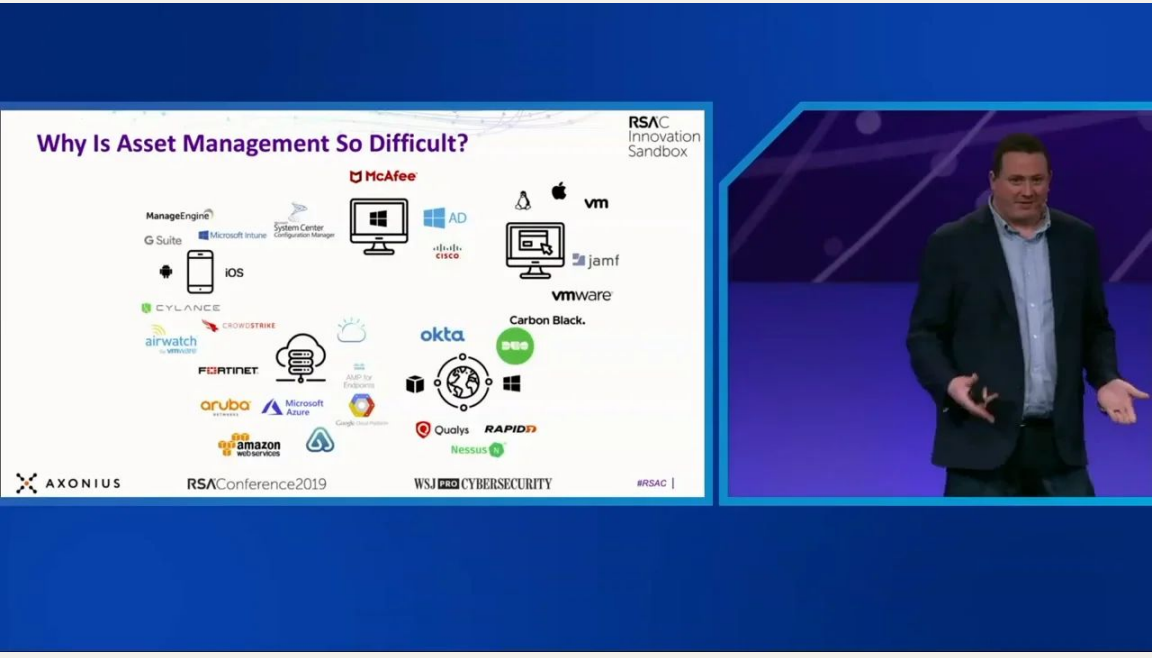



**FIG. 10.1** On the Innovation Sandbox stage, 2019.

We had been rejected before. Too early. That year we got in.

Dean, our CEO, was presenting. We had built the entire three minutes around him. Rehearsed for weeks. Every line tuned to sound like him, because a pitch that sounds like it was written for somebody else is a pitch a judge stops listening to.

The night before, his flight was delayed. Then cancelled.

///

## THE DISASTER

So I had less than twenty-four hours to rewrite three minutes of somebody else's story into mine, and to convince the organizers to let a CMO do this at all, which as far as anybody could tell had no precedent. They agreed reluctantly, which is the correct way to agree to something like that.

I rehearsed it roughly a hundred times in a mirror.

Then at the tech rehearsal, my notes didn't come up on the screen.

And I bombed. Badly. Badly enough that somebody in that room gave me a look I can still picture, and the look was sympathy.

So I left. I didn't stand there grinding out repetitions of a version that wasn't working, which I still think was the right call and which felt at the time like quitting.

And then on stage, in the real three minutes, I lost my place. Completely. Nothing there.

One second that felt like a month.

///

## BACK UP

I want to back up, because before any of that went wrong, we had done one thing that made all of it survivable.

RSAC doesn't publish a scoring rubric. It never has. What's public is the format, three minutes plus questions, and the eligibility rules, privately held, under five million in revenue, a recently launched product you can demonstrate live. There's no scoring sheet to download and there's no list of weighted criteria.

So I watched dozens of past pitches instead.

And the format turns out to be brutally formulaic. Everybody answers the same five questions, in roughly the same order, with roughly the same emphasis:

1. What's the problem?
2. Why is it big?
3. What's genuinely new about your approach?
4. Why are you the team to do this?
5. Can you survive if the floodgates open?

**Everybody hits the rubric. Which means hitting the rubric wins you nothing.**

That's the observation the whole chapter rests on, and it generalizes far past a pitch contest. Any scored situation where the criteria are known and the candidates are competent converges on identical answers. The criteria stop being a differentiator the moment everybody can read them. They become a floor, and standing on the floor isn't a position.

So the question stopped being "how do we cover the requirements" and became:

///

## THE CAMRY

The honest problem was that asset management is the least exciting thing in security. Everything else in this industry gets to be a Lamborghini. We were selling the ability to know what you have and whether it's actually protected, which is foundational and unglamorous and doesn't photograph well.

So instead of arguing that we were exciting, I said this:

> *The Toyota Camry is the most sold car in America because it fulfills the promise. It gets you from point A to point B. Every time.*

Nobody has a poster of a Camry on their wall. And every single thing you want to do that day depends on it starting.

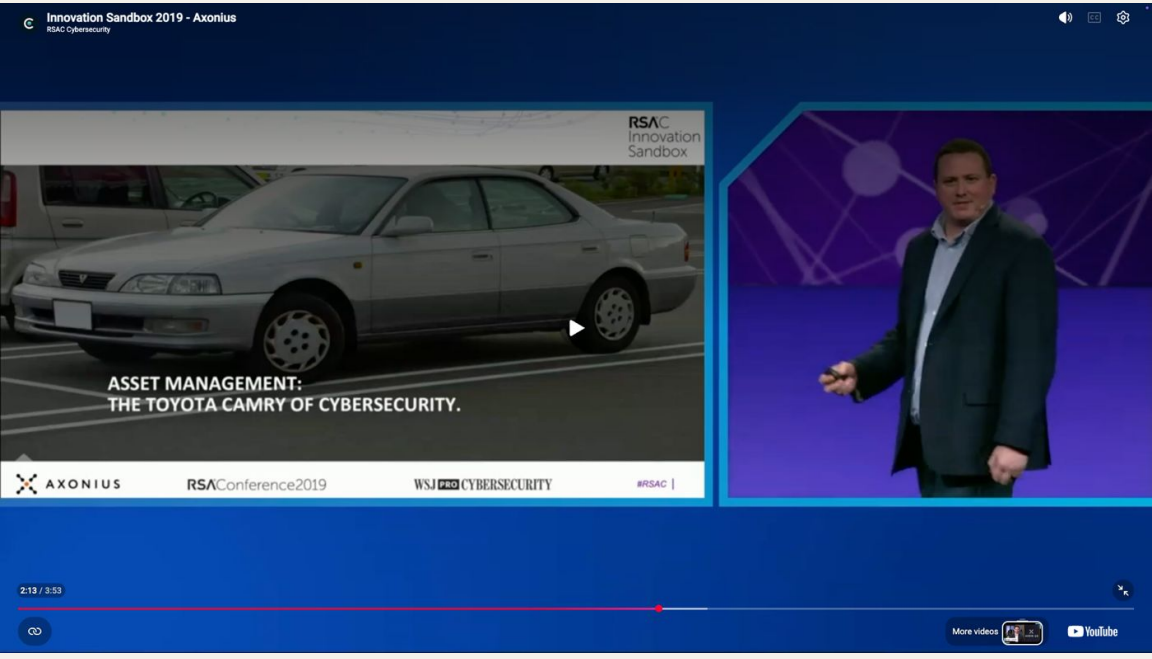



**FIG. 10.2** The Camry, on the RSAC stage. Nobody has a poster of it.

///

## WHAT CAME BACK

So. Standing on that stage. Nothing in my head.

The thing that came back was the Camry.

I wrote that line so a judge would remember it the next day.

Turns out I was the one who needed it.

I finished at exactly three minutes. And we won.



▶ SCAN TO WATCH

**RSAC INNOVATION SANDBOX 2019: THE AXONIUS PITCH**

The three minutes this chapter is about.

///

## THE PROOF, WHICH ISN'T THE WIN

A Fortune 500 company called before I was off the floor. They had watched it, they wanted a meeting immediately, they became a customer, and a POC started shortly after. That's the outcome, and it's the least interesting part.

Here's what I actually want you to notice.

Years later, people still walk up to me at industry events and say: you're the Toyota Camry guy.

Nobody has ever walked up and repeated our value proposition.

We spent real money and real years on that value proposition. It was accurate. It was reviewed. It was approved by people whose judgment I respect. And its total lifetime recall, measured in strangers at conferences, is zero.

The line about the car cost nothing, took an afternoon, and is still in circulation.

Memory ran in both directions on that stage. Externally, strangers retained it for years. Internally, under maximum pressure with my mind completely blank, the only thing that held was the line I had built to be memorable. I've never had better evidence that this property is real and that it's engineerable.

### Why it stuck

There's a reason it worked, and it's one of the oldest findings in memory research. In 1971, Allan Paivio described what he called dual coding.

Concrete words get stored twice, once as a word and once as an image. Abstract words mostly get stored once. So concrete things are remembered far better.

"Asset management" is abstract. A Camry is a picture.

Look at your own value proposition. I'd bet it's abstract. Every one in this category is, mine included. Platform. Visibility. Outcomes. Nobody can picture any of it, so nobody repeats any of it. Find your Camry. It costs an afternoon.

### The judge

One more piece, which I cut from the stage version because it teaches a different lesson.

The judge assigned to us came to the booth afterward for the demo. I didn't know who he was. The demo went well. I found out later that he was Ashim Chandna of Greylock.

If I had known that walking in, I'd have been more nervous and I'd have given a worse demo. I'm certain of this.

I don't have a clean method to extract from it. I include it because it's true and because the version of this story where every good outcome was engineered is a lie. Some of it was that I didn't know enough to be frightened.

///

## HOW TO DERIVE A RUBRIC NOBODY PUBLISHED

The transferable skill from the first half of this chapter, and it applies to any situation where you're being evaluated against criteria that exist but aren't written down. A pitch contest. An analyst evaluation. A vendor shortlist. A category review. An interview loop.

1. **Collect the artifacts.** Past winners, past submissions, past reports, past decks. Ten is enough. Twenty is better. They exist more often than

people assume, and the reason nobody looks at them is that watching twenty pitches is boring and looks like procrastination.

2. **Extract the questions, not the answers.** For each artifact, write down what question each section was answering. You're reverse-engineering the form from the responses.
3. **Count.** The questions that appear in every artifact are the real rubric. The ones that appear in the winners and not the losers are the interesting ones and there are usually fewer than you hope.
4. **Establish the floor.** Write the competent version of every answer. This is what everybody will do. You must do it too, because failing the rubric disqualifies you.
5. **Then budget the remainder for one thing.** Satisfy the rubric as efficiently as possible, quietly, and spend everything left over on the single element designed to be repeated. In three minutes that was about twenty seconds. Twenty seconds is enough.

The mistake almost everybody makes is spending the entire allocation on the rubric, answering all five questions beautifully, and finishing with nothing left over. That produces the best possible version of the thing everybody else is also producing, which is Chapter 2 again, in a different costume.

///

## HOW TO BUILD A CAMRY

Analogy construction as a discipline rather than a flourish. Five rules, in the order they matter.

**1. Source it outside your category.** Every analogy available inside cybersecurity is already in use inside cybersecurity, which means it arrives with fluency attached and can't produce a new thought. Cars, kitchens, sports, construction, medicine, airlines, weather, plumbing. The further from security, the harder it is to build and the better it works.

**2. Make it ordinary.** The Camry works because it's boring. An analogy to something glamorous flatters you and teaches nothing, and the audience notices the flattery before they notice the point. If your analogy makes your

product sound heroic, you've written an advertisement and it won't be repeated, because nobody repeats an advertisement.

**3. It has to survive the second question.** A good analogy holds when somebody pushes on it. A bad one is a simile that collapses the moment anybody asks what the equivalent of the engine is. Before you use it, have somebody hostile ask you three questions about it and see whether it bends or breaks.

**4. Test whether it explains or decorates.** Remove the analogy. Is the point still clear? If yes, the analogy is decoration, which is fine but isn't this. A working analogy does load. The Camry is doing the actual argument: foundational, unglamorous, and everything depends on it.

**5. Be slightly self-deprecating.** The Camry line concedes that we aren't the exciting one. That concession is what buys the claim underneath it. An analogy that claims a flattering comparison gets discounted on arrival by a skeptical audience, and yours is a professionally skeptical audience.

### And then use it for years

The failure mode after you build one isn't that it stops working. It's that **you get bored of it roughly nine months before your market has heard it once.**

Internal boredom is the single most expensive emotion in marketing. Every time you change the line because the team is sick of it, you reset the clock on the only asset that compounds. The Camry survived because I said it constantly, in every context, past the point where it embarrassed me.

///

## THE TEST

The repeat-back test from Chapter 9 is the instrument for this chapter. Tell five people. Wait a week. Ask them what you do.

If what comes back is your analogy, you've built one.

If what comes back is a paraphrase of your category, you've written a description, and a description is a thing people read once and never transmit.

## DESIGN FOR MEMORY, NOT JUST EXPLANATION.

Explanation is what you owe the person in front of you. Memory is what reaches the four people who weren't in the room, and in a market where the buying window is ten minutes long and opens without warning, those four people are the entire business.

///

### THE COMFORT CHECK

CHAPTER 10: THE THREE-MINUTE PITCH

**1. The comfortable assumption.** ~~You win a rubric by answering the rubric well.~~

**2. The true objective.** Be the one pitch a judge can still repeat tomorrow.

**3. The redesigned approach.** Satisfy the rubric quietly. Spend everything left on one analogy.

# 11

CHAPTER ELEVEN

# WHAT WORKED BEFORE IS A CEILING, NOT A FLOOR

**THE BEST THING WE HAD EVER DONE WAS ABOUT TO BECOME THE CEILING ON WHAT WE COULD DO NEXT.**

## 7AI launched in February of 2025. Black Hat that August was our first real show.

That year, AI in the SOC wasn't yet the thing every company said. So we built a booth that was an art gallery. Not a booth with art in it. A gallery. The demo stations were disguised as the exhibits. Quiet, still, deliberately strange, sitting in the middle of a show floor that's engineered end to end for sensory overload.

It worked. It worked really well.




**FIG. 11.1** The art gallery booth, Black Hat 2025.

That's the whole first half of the story and it's the least useful part of it.

///

## THE NEXT YEAR

The following year, everybody was an AI SOC company. Every booth. Every banner. Every hallway conversation. The category had arrived and it had arrived all at once, in the same words, which is Chapter 2 happening in three dimensions with carpet.

And here's the comfortable move, the one my Neanderthal voted for before I had finished reading the floor plan.

Do the gallery again.

Look at the arguments for it, because they're good arguments and you've made them:

1. It worked.
2. People loved it.
3. We have the photos, the plan, the vendor, the drawings, the budget line.
4. Most of the show hadn't seen it, so to most of the audience it's new anyway.
5. Doing something else means starting from zero with four months to go.

Every one of those is true. Number four is even a reasonable argument on the merits. And taken together they're the most persuasive case in marketing, which is why so many marketing departments are on year six of it.

That's the sentence the belief rests on. Not that repeating yourself is lazy. That your own best work, left in place, quietly defines the upper bound of the next thing, because everything after it gets evaluated against it rather than against the problem.

///

## WHAT WE DID INSTEAD

We took over the Sphere. And then put almost no words on the booth at all.

Not a product billboard. Not our feature list at four hundred feet. The whole goal was one reaction, and it's the title of the next chapter.

# WHAT THE HELL IS THAT?

People came to the booth specifically because they had seen it and wanted to know what it meant. Curiosity first, explanation second. That order matters enormously and almost nobody runs it that way, and Chapter 12 is entirely about why.

///

## THE CRITICISM

People had opinions.

That was the gist of a lot of what got posted. I can't believe they spent that much. That money could have gone somewhere else.

They might be right.

I'm not going to tell you the Sphere was the correct allocation of capital for your company, because it almost certainly isn't. I had a budget and an organization that would back a swing like that. Most people reading this don't, and that has nothing whatsoever to do with whether they're good at this.

**I'm not telling you to do what I did. I'm telling you to use the strategy I used.**

And the strategy was free. We noticed what everybody else was about to say, and we said something else.

### What a tenth of it would have bought

Which is the only question worth asking after a paragraph like that one, so here's the real answer rather than a shrug.

First, what it doesn't buy, because this is where most people go and it's the wrong read of everything in this book.

**A stunt.** Somebody stages a fake protest outside RSA. Somebody drives a branded truck in circles around the convention center. Somebody builds a thing purely to be photographed. I'm against all of it, and not on grounds of taste. A stunt is engineered to be noticed and nothing else, which means the noticing is the entire product. Once it's over you have a photograph and no asset, and the only thing anybody remembers is that a company did a thing, not which company or why.

Different isn't a synonym for attention-seeking. Go back to Chapter 7 and run the check properly. What's the objective? Then what's the cheapest route that actually reaches it?

Here are four we've used, all of them at a fraction of a Sphere, all of them available to a team with no budget at all.

**1. Do real research and publish it.**

Not a survey of four hundred CISOs about their top three concerns. Actual research: put somebody on a real question, let them find something nobody has documented, and write up what it is, why it matters, and what to do about it.

Two of ours. Last year our team found a QR-code phishing campaign that walked past every filter it hit. We called it **Quish Splash**, and we published what it was, why it mattered, and what to do about it. **CRXfiltrate**, an undocumented JavaScript execution backdoor running inside a network of Chrome extensions.

Both got written about by people who don't owe us anything. Trade press picked up both (HackRead on Quish Splash, eSecurity Planet on CRXfiltrate), and Quish Splash ended up on a podcast we never pitched. That's earned carriage, and no amount of paid media buys it, because the thing being carried is a finding rather than a message.

The cost is a person's time and the discipline to let them actually dig instead of pulling them onto launch support in week two.

**2. Find the right people instead of shouting at everyone.**

Use AI to read your own site traffic, work out who's actually showing up, build a lookalike from the ones that match your real ICP, and run LinkedIn ads at exactly that list.

Unglamorous, nobody will ever write about it, and it's a completely different activity from buying category keywords and hoping. You aren't increasing reach. You're deleting the part of your spend that was never going to work, which is a bigger lever than most brand campaigns and costs a fraction of one.

**3. Make the customer the hero, not the product.**

Our **Chief Innovation Security Officer** campaign. The subject is the security leader doing something genuinely new, and we're in the background.

The reason this works isn't generosity. It's that a customer will carry a story they're the hero of, into rooms you'll never be invited to, for years. A case study where the vendor is the hero gets forwarded by nobody, including the customer who's quoted in it.

**4. Start a podcast, and make it about them.**

**Do Human Work** costs approximately nothing. Equipment and travel. That's the whole line item.

What it produces is an hour of real conversation with a security leader about something they actually care about, and a number of those guests have become customers afterward. Not because it was a sales motion, which it isn't and which would kill it immediately, but because an hour of genuine attention is rarer in this industry than anything you could buy.

### What all four have in common

None of them buy attention. All four of them **earn carriage**, which means somebody who doesn't work for you moves your message for reasons of their own. A journalist writes the research up. A customer tells their own story. A guest recommends you to a peer. Even the targeting one works by finding the people who were already going to care.

The Sphere bought attention, and it worked, and it was expensive. These earn it, and they're slower, and they compound.

If I had a tenth of that budget I'd have put the whole thing into the first one and hired a second researcher.

///

## NOBODY DECIDES TO STOP TRYING

Here's why this matters past one booth.

Stagnation never feels like stagnation.

It feels like efficiency. It feels like using the thing that already worked, which is a completely reasonable thing to do, right up until it's the only thing you do.

There's a picture I keep saved in my phone under favorites, and I can't tell you how many times I've pulled it up. A school cafeteria tray of macaroni, with a handwritten sign that reads "macaroni and fuck it."




**FIG. 11.2** Nobody chose this. Somebody just did what worked last time.

Nobody decided to serve that. Nobody stood in that kitchen and chose it as a strategy. Somebody just did what worked last time. And then did it again. And the thing about doing what worked last time is that each individual instance is defensible, and the aggregate is a sign on a tray that has given up.

Your annual threat report is on year five. Your webinar format hasn't changed since before the pandemic. Your booth has the same shape and the same demo stations and the same shirts in a different color. Every one of those decisions was reasonable in the year it was made.

And look, I get it. Nobody can go all out on everything, every day. There are deadlines and budgets and things that just need to get done. So keep that picture in your head and ask yourself: is this good, or is it macaroni and fuck it? At least then you know which choice you're making.

///

## WHY THE OLD ANSWER BLINDS YOU

There's a classic experiment on this. In 1942, Abraham Luchins gave people a series of water-jar puzzles that all shared the same three-step solution. Then he gave them one that could be solved in a single step. Most of them kept using the long method. They couldn't see the easy answer, because the old answer kept working. People who had never learned the long method found the shortcut right away.

It's called the Einstellung effect. Your best idea doesn't just sit there. It gets in the way of the next one.

///

## THE CEILING AUDIT

Run this before planning season, not during, because during planning season every play on the list already has an owner and a slide.

**Step 1.** List every play you're going to run again next year because it worked. Not everything you do. The things whose justification is that they worked.

**Step 2.** For each one, write the year it started and the year it last changed materially. Not cosmetically. Materially. New color isn't a change. New format is a change.

**Step 3.** For each one, answer: **what would we do if this play were banned?**

Not "what would we do instead," which produces a slightly smaller version of the same thing. Banned. Prohibited. Not available at any budget.

**Step 4.** Look at the answers. Somewhere between one and three of them will be better than the play you were going to run. That isn't a hypothetical; it happens every time I've run this, and it happens because the ban is the constraint from Chapter 7 applied to your own history.

**Step 5.** Kill one thing that's working.

This is the step people skip and it's the only one that has teeth. Not the worst thing. Something that's genuinely working, that you could defend, that nobody is asking you to stop. One per year.

The reason is arithmetic. A marketing calendar has a fixed size. If you only ever remove things that failed, the calendar fills up with things that worked once, and within three years you have no capacity left for anything new and a completely defensible explanation for why.

///

## CRITICISM IS A SYMPTOM, NOT A SCOREBOARD

I need to be careful here, because there's a version of this chapter that produces an insufferable person.

Criticism is a **result** of being different. It isn't evidence that you're right. Those are completely different claims and conflating them is how marketers end up defending bad work by pointing at the people who dislike it.

The honest version:

1. If nobody reacts, you were probably invisible. Silence is a finding.
2. If people react, you were visible. That's all it establishes.
3. Whether you were *correct* is settled by Chapter 9, not by the comment section.

The useful thing about criticism isn't its emotional valence, it's its content. Look at what the criticism is actually about.

Criticism of your judgment is expensive and worth taking seriously. Criticism of your taste is mostly noise. Criticism that takes the form of explaining your own strategy back to you incorrectly is genuinely valuable, because it tells you precisely which part didn't travel.

And one thing that's true and that I wouldn't have believed ten years ago: **the people who criticized the Sphere publicly could all describe what we did.** That isn't nothing. It isn't success either. It's a leading indicator, and it belongs in the leading indicator row of the measurement plan rather than in your feelings.

///

## STRANGE IS FREE

The Sphere is the expensive example and it's the least useful one in this chapter. Here's the catalogue of things that cost nothing and that almost nobody in this industry does, because the cost of them isn't money.

1. **What you name things.** Internal program names, feature names, the name of your report. Naming is free and it's the highest-leverage unclaimed surface in B2B.
2. **What you refuse to say.** Ban one category word for a year, across everything. The hole it leaves has to be filled with something true.
3. **Who speaks.** Put the engineer, the analyst, or the customer on stage instead of the executive whose title says they belong there. Costs nothing, prohibited by nothing except habit.

4. **The format.** A memo instead of a deck. A ten-minute webinar instead of forty-five. A one-page report instead of a forty-page PDF nobody finished.
5. **The ten minutes before.** Everybody puts a countdown timer and a logo on screen before a webinar starts. That's ten minutes of undivided attention from your most engaged audience, given away to a clock.
6. **What you send after a meeting.** Every vendor sends a recap email and a deck. Almost nobody sends the one thing the buyer actually needs to forward internally.
7. **The back wall of the booth.** It faces the aisle for three days and almost every company puts their tagline on it.
8. **What you refuse to gate.** Ungating one thing that everybody gates is a position, not a tactic, and people notice.
9. **Your job titles.** Free, and they're on every email signature your team sends.
10. **Who you thank in public.** Nearly unused in this industry, and it's the cheapest distinctive act available.

Every one of those is available to a person with no budget and no authority beyond their own work. That's Chapter 18's belief arriving early, and it arrives here because this is the chapter where the money is most likely to distract you from the method.

///

## THE COMFORT CHECK

CHAPTER 11: THE SECOND YEAR

**1. The comfortable assumption.** ~~It worked. So run it back.~~

**2. The true objective.** Be the thing nobody can walk past.

**3. The redesigned approach.** Go bigger and stranger than your own best work.

# 12

CHAPTER TWELVE

# CURIOSITY BEFORE EXPLANATION

**CURIOSITY IS KNOWING THERE'S A QUESTION. CONFUSION ISN'T KNOWING WHAT THE QUESTION IS.**

Walk a show floor and count how many booths are trying to explain something to you.

All of them. Every headline is a compressed value proposition. Every banner is a sentence that would work fine on a slide and does nothing on a wall thirty feet away from a person walking past at three miles an hour with a coffee and a bag of other people's pens.

The explanation isn't wrong. It's just arriving before it has been asked for, and an explanation nobody asked for is noise regardless of how accurate it is.

///

## THE SEQUENCE

What actually happens when marketing works looks like this:

1. **See** something unexpected.
2. **Wonder** what it is.
3. **Ask**, or look it up, or walk over.
4. **Enter** the story.
5. **Remember** the brand attached to it.

Five steps, and the entire industry starts at step four.

The reason is completely understandable. Step four is the only step you can put in a brief. "Enter the story" is a deliverable. "Wonder" isn't a deliverable, it's a state in somebody else's head, and no agency has ever been briefed to produce one.

But notice what the order costs you. **Explanation requires attention you haven't yet been granted.** You can't explain your way into somebody's attention, because the explaining is what you do once you have it. Curiosity is the mechanism that grants it, and it's the only mechanism, and we skip it because it doesn't fit in a project plan.

This is why the Sphere worked, and it had almost nothing to do with the size of it. People came to the booth because they had seen something and wanted to know what it meant. The wanting to know was the product. The booth was where the wanting got resolved.

///

## THIS ISN'T CLICKBAIT, AND THE DIFFERENCE IS PRECISE

The obvious objection: manufacturing curiosity is what a tabloid headline does, and B2B has enough of that.

The objection is correct and the distinction is sharp.

**Clickbait creates a gap and then doesn't close it.** You arrive, and the thing on the other side is either the obvious answer or nothing. The mechanism works exactly once per person, and it spends your credibility to do it.

**Curiosity marketing creates a gap and then over-delivers on closing it.** The payoff is bigger than the tease. You arrive, and there's more there than you expected, and the next time that brand makes you wonder something you go and find out, faster.

The test isn't whether you created a gap. It's what's on the other side of it. If the thing on the other side is a form, you built clickbait with better manners.

///

## CURIOSITY AND CONFUSION AREN'T THE SAME, AND HERE'S HOW TO TELL

This is the failure mode of everything in this chapter and of most of what happens when a marketing team is told to be interesting.

The test is mechanical, it takes three seconds, and you should run it on everything before it ships.

Show the thing to somebody outside your company for three seconds. Then ask them to state, in one sentence, the question it raised.

1. If they can state the question, you have curiosity. "What does that have to do with security?" is a perfect result. "Why is there a gallery here?" is a perfect result.
2. If they can't state the question, you have confusion, and confusion doesn't convert to interest. It converts to walking.
3. If they state the answer instead of a question, you have explanation, which is the thing you were trying to avoid.

That's the whole quality bar. It's more useful than any amount of internal debate about whether something is too weird, because too weird isn't a category that exists. **Unstateable** is the category that exists.

///

## WORKED, WITH BEFORE AND AFTER

Seven surfaces, all cheap, all available to somebody with no budget.

### The booth headline

**Before.** "AI-Powered Security Operations for the Modern Enterprise." **After.** A single sentence that's obviously not a product claim, placed where the claim usually goes, and a reason for it somewhere you have to walk closer to read.

The mechanism: a wall that isn't doing the job walls do at this show is itself the question. The question is stateable. What's this booth for?

### Paid social

**Before.** A product screenshot and three bullet points. **After.** One number, or one sentence a practitioner would forward to another practitioner, with no product in the frame at all.

The test: would somebody screenshot this and send it to a colleague with no comment? That's the only engagement metric that matters at the top, and a bulleted feature list has never once been screenshotted by anybody.

### The email subject line

**Before.** "Webinar: How to Reduce Alert Fatigue with AI." **After.** The specific, slightly odd thing you found. "The part of the questionnaire nobody reads."

Specific and small beats broad and important. A broad important subject line is indistinguishable from the eleven other broad important subject lines above it in the inbox, and the inbox is Chapter 2 with a scroll bar.

### The webinar title

**Before.** "Best Practices for Modern Detection and Response." **After.** A question the audience already has and has never seen a vendor say out loud.

Best practices is the single most fluent phrase in B2B, which is precisely why it produces no reaction of any kind.

### The conference session abstract

**Before.** A summary of what you'll cover. **After.** The claim you're going to make, stated bluntly enough that somebody could disagree with it before attending.

An abstract that nobody could disagree with is an abstract nobody chose.

### The report cover

**Before.** The year, the category, an abstract shape, a gradient. **After.** The finding. The actual finding, on the cover, as a sentence.

Everybody hides the finding inside to protect the download. The download isn't the objective, the finding travelling is the objective, and a finding on the cover gets photographed and posted by people who never downloaded anything.

### The first sales touch

**Before.** "I wanted to reach out because I noticed you're using X." **After.** One observation specific enough that the recipient knows a human looked at something. Chapter 17 is how that gets produced at scale without a human looking at something every time.

///

## WHY THIS IS THE CHEAPEST CHAPTER IN THE BOOK

Every example above costs zero incremental dollars. You were already writing the subject line. You were already printing the wall. You were already naming the webinar.

Curiosity has no budget floor, which makes it the most obvious available move for anybody who doesn't control spend, and the least used, because it has a different price. It requires somebody to approve a sentence that doesn't explain what the company does.

That approval is where these die, almost every time. The work is free. The nod is the expensive part, because a headline that raises a question will always feel, to a room full of people who already know the answer, like it's failing to say the thing.

Which is the loop closing. **The reason curiosity is rare isn't that it's hard to make. It's that it's hard to get approved, because the people approving it are the only people in the world who can't experience it.**

///

## THE COMFORT CHECK

CHAPTER 12: THE BOOTH WALL

**1. The comfortable assumption.** ~~The job of the headline is to say what we do.~~

**2. The true objective.** Get somebody who was walking past to stop and ask a question we can answer.

**3. The redesigned approach.** Put a stateable question where the claim usually goes, and put the explanation where somebody has to come closer to get it.

# IV

PART FOUR

# THE MACHINE THAT MAKES AVERAGE

# 13

CHAPTER THIRTEEN

# YOUR NEANDERTHAL NOW HAS A COPILOT

---

**EVERY INDIVIDUAL GOT BETTER. THE GROUP GOT MORE ALIKE.**

Back to the wall from Chapter 2. Twelve different brands, one sentence.

I told you they weren't real. I didn't tell you how I made them.

I described our category to a model and asked it for a homepage. Twelve times. I barely changed a word.

In Chapter 2 that sameness was a story about people. This chapter is about what happens when you hand those same people a machine that's better at sameness than they are.

///

## THE FLUENCY ENGINE

Here's the plain version of how a language model works, and it's all you need for the rest of this part of the book.

A model is trained to produce the most probable next words, given everything before them. The most probable words are, by definition, the most common words. And the most common words in your category are the ones everybody in your category has already written, which is exactly what everybody else's model is producing, from roughly the same prompt, on the same Tuesday.

That isn't a flaw in the tool. That's the design. It's a fluency engine. It was built to produce text that goes down smoothly, and smooth is what familiar feels like from the inside.

Now put that next to Chapter 1.

Your brain was already biased toward the familiar. It has been for forty thousand years, and it reports familiarity to you as a feeling of correctness. We've handed that brain a tool whose single greatest skill is generating the familiar, instantly, at volume, for almost nothing.

These don't add. They compound. The model hands you the most probable sentence, your Neanderthal recognizes it and reports that it's right, and the

review round sands off anything the model happened to get interesting. Three filters, all tuned to the same frequency.

///

## THIS HAS BEEN MEASURED

In 2024, Anil Doshi and Oliver Hauser published a study in *Science Advances* that I think every marketer using AI should read.

They had people write short stories. Some wrote with AI help and some without. When independent readers rated the stories one at a time, the AI-assisted stories came out ahead. More creative, better written, more enjoyable, and the lift was biggest for the writers who had started out least creative.

Then they compared the stories to each other. The AI-assisted stories were significantly more similar to one another than the stories written without help.

That's the wall of homepages, in a peer-reviewed journal. Every one of those twelve pages is good. Read one by itself and it's clear, confident, well structured. Read all twelve and they're one website.

This is the trap, and it's why this chapter matters more than anything in Part Three. The review that would catch the problem looks at one piece at a time. One piece at a time, the AI version wins. The cost only shows up when you put your work next to everybody else's, which is the only way your buyer ever sees it.

///

## BELIEF TWO

So here's the second of the seven beliefs, arriving where it's earned.

# AI DIDN'T MAKE YOU FASTER. IT MADE EVERYONE FASTER AT THE SAME THING.

The gain is real. I'm not going to pretend otherwise in a book about honesty. A team using AI ships more, ships it sooner, and the average quality of each piece goes up.

But speed toward the middle isn't an advantage when everybody else has the same speed and the same middle. It's a treadmill. You produce more of what your competitors produce, faster, and the market experiences it as more noise from a category that already had plenty.

///

## THE TEST YOU SHOULD RUN

Before the keynote I asked myself whether the model was only converging because I gave it a boring prompt. So here's a test you can run in ten minutes, and I'd rather you run it than take my word for anything.

Open whatever model your team uses. Ask it, twelve separate times in fresh sessions, for **the most distinctive homepage it can write** for a company in your category. Not a homepage. The most distinctive one.

Put the twelve heroes in a grid.

If they converge, you've learned the most important thing in this chapter: the machine's idea of distinctive is the category mean. Asking it to be different doesn't make it different. It makes it reach for the most common way of sounding different, which is its own kind of sameness.

If they don't converge, you have a better model or a weirder category than I do, and you should tell me.

///

## THE ANTI-FLUENCY TOOLKIT

None of this is an argument against using AI. I use it more than almost anybody I know, and Chapter 15 is the inventory. It's an argument for using it on purpose, against its own grain. Six techniques, in roughly the order I'd adopt them.

**1. Locate the mean first, on purpose.** Use the model to find out exactly what the obvious answer is, then treat that answer as the one thing you're not allowed to say. This is the single best use of a fluency engine: it's the fastest way ever invented to find the middle of your category so you can leave it. Ask it for the most typical version. Write it down. Ban it.

**2. Inject a constraint.** Every constraint from Chapter 7 works on a model, and works better than on people, because a model doesn't resent you. Two words. No product nouns. No comparatives. Explain it to a plumber. A constraint makes the most probable answer illegal, and the model has to search further down the list.

**3. Keep a forbidden-words list.** Every category has twenty words that carry no information because everybody uses them. Write yours down and give the list to the model at the start of every session. In our corner of the market, the list starts itself: agentic, machine speed, seamless, empower, the work that matters. A word on that list in a draft is a word you didn't choose.

**4. Ground it in material the model hasn't seen.** A model can only remix what it has read, and it has read everything public in your category. It hasn't read your customer calls, your sales team's loss notes, the transcript of your CEO arguing with a prospect, or the strange thing a customer said on a podcast. Primary material is the only input that produces output nobody else can produce, because nobody else has the input.

**5. Run an adversarial pass.** After the draft, start a fresh session and ask a model to tear it apart as a skeptical buyer who has read eleven competitor pages this morning. What's generic? What would a competitor say word for word? Which sentence could be deleted without anybody noticing? A critic that never gets tired is one of the most underrated uses of the tool.

**6. Use a voice profile, and build it properly.** A model with no instructions writes like the average of the internet. A model with a real profile of how a specific person writes, built from their own unedited writing and their corrections to drafts, writes like that person on a reasonable day. Chapter 15 covers how we build them, and the one thing that stops a voice profile from producing parody.

Notice what all six have in common. Each one adds friction back into a process that has been engineered to remove it. That's the whole idea of this book, applied to a machine.

///

## THE COMFORT CHECK

CHAPTER 13: WRITING WITH AI

**1. The comfortable assumption.** ~~If each piece is better, the work is better.~~

**2. The true objective.** Make work that a buyer can tell apart from eleven competitors who have the same tools.

**3. The redesigned approach.** Use the model to find the category mean and ban it, write under a constraint, ground every draft in primary material nobody else has, and judge the output next to the competition instead of on its own.

# 14

CHAPTER FOURTEEN

# CRUTCH OR MULTIPLIER

**SIX MONTHS IN, IS YOUR WORK MORE DISTINCTIVE, OR MORE AVERAGE?**

There are two ways to use AI at work, and on a status report they look identical.

A **crutch** does the work you were already going to do, faster.

A **multiplier** does work you could not have done at all.

Both show up in the quarterly update as "we adopted AI." Both save time. Both make a nice slide. One of them makes you better at your job. The other makes you faster at being interchangeable.

///

## WHY YOU CAN'T TELL THEM APART FROM THE INSIDE

The crutch feels great. That's the problem.

Writing the blog post in forty minutes instead of four hours is a real, immediate, personal relief. You feel it the same afternoon. The multiplier, by contrast, usually feels like more work at first, because it's doing something that didn't exist before and somebody has to figure out what to do with the output.

So left alone, every team drifts toward crutches. Not because anyone is lazy, but because the reward for a crutch arrives today and the reward for a multiplier arrives next quarter. That's the same asymmetry as Chapter 3, wearing a different jacket.

And the metrics most companies use to track AI adoption make it worse. Seats licensed. Prompts per week. Hours saved. Every one of those numbers goes up just as fast for a crutch as for a multiplier, and some of them go up faster.

///

## THE SIX-MONTH TEST

There's exactly one test I trust, and it isn't adoption, hours saved, or seats.

Run it on your own output. Take what your team shipped six months ago and what it shipped last month. Put them next to the competition, not next to each other. Has the gap between you and the category gotten wider or narrower?

If it's narrower, you bought a crutch, and the time it saved was spent producing more of the middle. If it's wider, you have a multiplier, and you should find out exactly which part of the workflow is responsible and protect it.

The test has a second version for people. Six months in, is this person doing work they could not have done before, or the same work faster? The first is a career. The second is a countdown.

///

## A TAXONOMY

Here's how I'd call the common uses in a marketing function. These are judgments, not laws, and a few of them flip depending on how you run them.

| USE | CALL | WHY |
|---|---|---|
| Drafting blog posts from a topic | **Crutch** | It produces the category's post on the topic, which is the post your competitors are also producing |
| Drafting from your own primary material (calls, interviews, internal debates) | **Multiplier** | The input is unique, so the output can be |
| Generating social posts at volume | **Crutch** | More of the same, faster, in a feed that already has plenty |
| Finding the stories worth a point of view, then letting a human decide | **Multiplier** | It reads more than any person can and leaves the judgment where it belongs |
| Summarizing meetings and calls | **Conditional** | Crutch if the summary goes in a folder. Multiplier if it's mined for what customers actually say, across hundreds of calls |

| USE | CALL | WHY |
|---|---|---|
| Rewriting copy to "make it better" | **Crutch** | Better, to a model, means smoother. Smoother means closer to the mean |
| Adversarial critique of a draft | **Multiplier** | It adds friction, which is the scarce thing |
| Researching an account before a first meeting | **Multiplier** | Nobody was doing it properly before, because it took ninety minutes per meeting |
| Writing in an executive's voice from a real profile | **Conditional** | Multiplier if the ideas are theirs and the time goes back to them. Crutch if it becomes a way to post without having anything to say |
| Competitive monitoring on a schedule | **Multiplier** | The alternative wasn't a slower version. It was a spreadsheet nobody updated |
| Image generation for campaign creative | **Conditional** | Crutch for stock-style images, which is the hooded figure again. Multiplier for testing ten strange concepts before you commission one |
| Reporting and dashboards | **Conditional** | Crutch if it makes the same report faster. Multiplier if it produces a read nobody asked for that changes a decision |
| Personalized outbound at scale | **Crutch, usually** | Personalization that a buyer can see was generated is worse than none |

The pattern in the multiplier column is consistent. A multiplier either works from input nobody else has, does something nobody was doing at all, or adds friction where the process had none. A crutch takes the thing you were already producing and produces it with less effort.

///

## MAKING IT ON PURPOSE

The useful question isn't "how do we adopt AI?" It's "what could we do now that we couldn't do before?" Ask it about the specific work your team does, and write down the answers before you buy anything.

Most of the answers will be things that took too long for anybody to do well. Research before every meeting. Reading every call transcript. Tracking

every competitor every week. Answering the same question from every new rep. The crutch version of AI makes your existing work cheaper. The multiplier version makes impossible work possible, and that's the only version that changes your position.

The next chapter is what that looks like when you actually build it.

///

## THE COMFORT CHECK

CHAPTER 14: AI ADOPTION

**1. The comfortable assumption.** ~~We're doing well with AI, because adoption and hours saved are up.~~

**2. The true objective.** Do work that was impossible before, and widen the gap between us and the category.

**3. The redesigned approach.** Judge every use by the six-month test, keep the ones that work from unique input or add friction, and stop counting seats.

# 15

CHAPTER FIFTEEN

# WHAT I ACTUALLY BUILT

**THE UNIT OF WORK IS A DOCUMENT, NOT A SUBSCRIPTION.**

## This is the chapter I most wanted to write, and the one that's easiest to get wrong.

The wrong version is a tools roundup. Here are the twelve apps I use, here's my prompt library, here's a screenshot of my workflow. That chapter is out of date before it's published, and it teaches the reader to buy things.

So this one is built differently. It covers the systems running in my marketing function, described as patterns rather than products, each one the same way: what it does, what it replaced, what it costs, what to watch for, and how it scores on the six-month test. I've left out what's specific to my company. You don't need it. The patterns transfer, and you can build every one of them yourself.

And it ends with one I built that I never shipped, because a chapter of nothing but successes would be the exact victory lap this book argues against.

Here's the idea that runs under all of it, and it's the most transferable thing in this part of the book:

Every system below is a skill: a written set of instructions that a model follows, with sources to read, steps to take, rules to obey, and a format to deliver. Not software I bought. A document I wrote and keep improving. Which means a reader with no budget can copy all of it, and Appendix F has the versions you can run.

///

# 1. VOICE PROFILES

**What it does.** Lets a model draft in a specific person's voice well enough that the person edits rather than rewrites. We keep one for each executive who publishes, including me.

**What it replaced.** A ghostwriting process where somebody on my team spent an hour interviewing an executive, an hour drafting, and a day waiting for edits that rewrote half of it.

**How it's built.** Not from their published writing. Published writing has already been edited, often by marketing, so it's partly our voice already. The useful sources are the ones nobody polished: how they write in chat, how they answer an email at 11 PM, and above all **their corrections to drafts.** When an executive rewrites a sentence, the before and after is the most precise description of their voice you'll ever get. A good profile records what they always do, what they never do, the words they hate, and how their punctuation behaves.

**What to watch for.** Parody. A profile that only lists a person's habits produces a draft that uses every habit in every paragraph, and it reads like an impression rather than the person. The fix is a section most profiles don't have: the **imitation tells**, a list of the things a copycat overdoes. Use their signature phrase once a month, not once a paragraph. Real voices are mostly ordinary sentences with a few distinctive ones.

**Crutch or multiplier.** Conditional, and the condition matters. It's a multiplier when the ideas are theirs and the profile gives them their time back. It becomes a crutch the moment it's used to publish when nobody had anything to say.

///

## 2. THE WEEKLY CONTENT SYSTEM

**What it does.** Every week, an agent reads the news across our market, finds the stories worth having an opinion on, filters them, and proposes a short plan: which stories, which person should speak to each, and the angle. Then it drafts against the right voice profile. A human picks, edits, and posts.

**What it replaced.** Somebody scrolling for two hours on Monday, or more honestly, nobody doing it and posts going out when an executive happened to have a thought.

**The part that matters.** The filter, not the drafting. Ours asks whether the person posting can be the protagonist of the take: do they have a real point of view, earned from something they've done, or would they just be summarizing the news? A post that summarizes the news is the crutch

column of Chapter 14. A post where somebody who has done the work says what they think is the multiplier, and the filter is what tells them apart.

**What to watch for.** The system finding the same stories every competitor's system finds. If your agent reads the same feeds as everybody else's, it'll surface the same top five stories. Give it sources the others don't read, and tell it to rank the strange story above the big one.

**Crutch or multiplier.** Multiplier, as long as a human still chooses. The day it starts posting without anybody deciding, it becomes a machine for producing the category's opinion on schedule.

///

## 3. THE MONDAY REVIEW

**What it does.** Every Monday, a set of agents runs against our CRM, our internal search and our documents, and hands me a read on where we are against our goals, how spend is pacing, and what we should be building that nobody has asked for yet.

**What it replaced.** That used to take somebody about two days, when it happened at all.

**The honest claim.** It didn't save two days. That's the crutch version of the story and it's not the interesting one. What it actually did is give me a proactive read I never had before. The old version of this was reactive: somebody asked a question in a meeting and somebody else went to find out. Now the answer, and three questions I hadn't thought to ask, are sitting there before the week starts.

**What to watch for.** Believing it. An agent that reads your CRM can only report what's in your CRM, and every CRM has holes, stale stages and deals that exist mostly as optimism. The review is a starting point for a conversation with the people who own the numbers, not a replacement for it. Humans on the loop, always: the agent proposes, a person decides.

**Crutch or multiplier.** Multiplier. It produces a read that didn't exist, and it changes what gets decided on Monday.

///

## 4. THE COMPETITIVE LIBRARY

**What it does.** One page per competitor, always current. A sweep runs every week to catch what changed: new messaging, new pages, new claims, new hires, new ads. A deeper rebuild runs monthly. Alongside it sits a deal ledger, a running record of every deal where that competitor showed up, what happened, and why.

**What it replaced.** The battlecard, which in most companies is a document somebody updates the week before a sales kickoff and nobody trusts by March.

**The decision that made it work.** Wins and losses, including the loss reasons and who we lost to, are visible to every rep. Not gated to leadership. It would have been easy to keep the losses private, and it would have made the library much less useful. My reasoning was simple: it's how we learn. A rep walking into a deal against a competitor deserves to know exactly how the last three went, especially the ones that went badly.

**What to watch for.** Sameness, again. A library that only tracks what competitors say will make you better at responding to them, which is a polite way of saying it'll pull you toward them. Use it to find what everybody in the category is saying, and then do Chapter 2's exercise: work out what nobody is saying.

**Crutch or multiplier.** Multiplier. The alternative wasn't a slower library. It was no library.

///

## 5. THE LIVING TIMELINE

**What it does.** A page that tracks one fast-moving topic in our market over time, every entry with a date, a source, and what we think it means. A scheduled agent checks for new entries every day and proposes updates. It runs internally today, and the plan is a public version with a short weekly recap as its own post.

**What it replaced.** The annual report on the same topic, which is out of date the week it ships, and which Chapter 5 already suspected of being annual for no reason.

**Why it's different.** A report is a moment. A living resource is a place people come back to. It's a better asset for search, for sales, and for anybody trying to understand the topic, and it's a reason for a buyer to return that has nothing to do with being sold to. It also forces us to have an opinion every week rather than once a year, which is uncomfortable in exactly the right way.

**What to watch for.** The agent adding things that are wrong. Fast-moving topics are full of bad reporting, and an agent that reads the news will happily repeat it. Every entry gets a source and a human check before it goes up.

**Crutch or multiplier.** Multiplier. Keeping a resource like that current by hand would take a person's whole job.

///

## 6. THE INTERNAL HUB

**What it does.** One place for a new rep to find the current version of anything: the deck, the competitive page, the sales plays, the partner material. Agents keep it updated from where the work actually happens, and a search box answers questions across all of it.

**What it replaced.** Document sprawl. Fast hiring, a dozen folders, three versions of every deck, and a chat tool where the answer to every question exists somewhere and finding it takes twenty minutes.

**How it's going.** Honestly, it's new. It's in beta with our go-to-market team, there's a feedback channel where people flag wrong versions and ask for things that are missing, and the measure of whether it works is whether new reps stop asking where things are. I'm including it before I know the answer on purpose, because the alternative is waiting until every example in this chapter is a proven success, and you already know how I feel about that.

**Crutch or multiplier.** Too early to call. Ask me in six months, which is the point of the test.

///

## WHAT ALL SIX HAVE IN COMMON

1. **Each one is a document.** Sources, steps, rules, format. Anybody on my team can read it, change it, and run it, and some of my team run skills like these without ever opening a terminal.
2. **Each one does work that wasn't getting done.** Not faster versions of old work. Proactive reads, current libraries, research nobody had time for.
3. **A human decides in every one of them.** The agents propose and assemble. People choose what's true, what gets said, and who says it.
4. **None of them required a budget.** The cost is the time to write the instructions well, and the willingness to keep rewriting them when they're wrong.

///

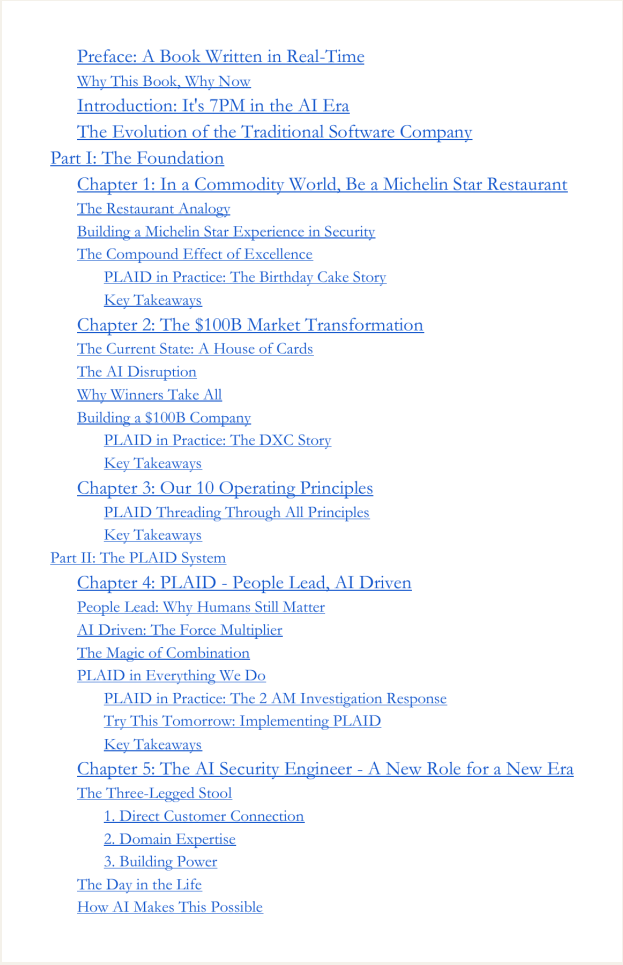
## THE ONE I SHELVED

Last September, we wrote a book.

It was called *PLAID: The Operating Model for AI-Native Companies*, and it took us less than a week. We were proud of that. It says so in the book.




**FIG. 15.1** The cover of the book we never published. Note the subtitle.




Preface: A Book Written in Real-Time
Why This Book, Why Now
Introduction: It's 7PM in the AI Era
The Evolution of the Traditional Software Company
Part I: The Foundation
Chapter 1: In a Commodity World, Be a Michelin Star Restaurant
The Restaurant Analogy
Building a Michelin Star Experience in Security
The Compound Effect of Excellence
PLAID in Practice: The Birthday Cake Story
Key Takeaways
Chapter 2: The $100B Market Transformation
The Current State: A House of Cards
The AI Disruption
Why Winners Take All
Building a $100B Company
PLAID in Practice: The DXC Story
Key Takeaways
Chapter 3: Our 10 Operating Principles
PLAID Threading Through All Principles
Key Takeaways
Part II: The PLAID System
Chapter 4: PLAID - People Lead, AI Driven
People Lead: Why Humans Still Matter
AI Driven: The Force Multiplier
The Magic of Combination
PLAID in Everything We Do
PLAID in Practice: The 2 AM Investigation Response
Try This Tomorrow: Implementing PLAID
Key Takeaways
Chapter 5: The AI Security Engineer - A New Role for a New Era
The Three-Legged Stool
1. Direct Customer Connection
2. Domain Expertise
3. Building Power
The Day in the Life
How AI Makes This Possible


**FIG. 15.2** The first page of contents. It ran three pages.

It had ten operating principles, a marketing framework with its own acronym, a Michelin star metaphor running through every chapter, and an epilogue that ended "The revolution isn't coming. It's here." It read like a book. It had the shape of a book.

We never published it. Nobody told us not to. It just never became the priority, and every time I opened it I could feel why. We'd been selling for a few months. Most of the stories worth telling, we couldn't tell publicly yet, so the ones that made it in were "a Fortune 1000 retailer" and "a manufacturing company," and the claims were the kind you make before you can prove

them. We were writing the operating manual for a company that hadn't operated long enough to need one.

So I stopped. Not because it failed. Because it wasn't good enough, and I knew it, and nobody else was going to say so.

The ideas lasted. All of them. They're how we run the company now. The book didn't, because the ideas came from doing the work and the book came from a week with a machine.

Read it again with Chapter 13 in mind and it's almost funny. Not AI-enabled, not AI-enhanced, AI-native. A comparison to the Model T. A section headed "This Changes Everything." It's the wall of homepages at book length, and I wrote it. With a model doing the structuring and smoothing, we got the most probable book about building an AI company, and the most probable book is the one everybody in the category would have written.

It's also the other side of Chapter 8. With the car, Ed was the one person who didn't nod, and I ignored him. With the book, the person who didn't nod was me, and this time I listened.

> **AI MADE IT POSSIBLE TO WRITE THAT BOOK IN A WEEK. IT COULDN'T MAKE US READY TO WRITE IT.**

That's the limit of everything else in this chapter. A model can produce the artifact. It can't produce the proof that the artifact deserves to exist: the customer stories you're allowed to tell, the results you can stand behind, the years of doing the thing. When the artifact arrives before the proof, it doesn't look early. It looks finished, and that's what makes it dangerous.

///

## THE COMFORT CHECK

CHAPTER 15: BUILDING WITH AI

**1. The comfortable assumption.** ~~Using AI well means buying the right tools.~~

**2. The true objective.** Do work that wasn't possible before, with people still deciding what's true.

**3. The redesigned approach.** Write the work down as skills anyone can run and improve, aim them at work nobody was doing, and don't ship the artifact before you have the proof.

# 16

CHAPTER SIXTEEN

# YOUR VALUE WAS NEVER THE TASKS

**DO YOU WANT TO GET PROMOTED THIS YEAR?**

## The best AI adoption story I know didn't come from marketing. It came from a CISO.

Benjamin Dulieu is the CIO and CISO at Duck Creek Technologies. He told this story on our podcast, Do Human Work, and he's been generous enough to let me retell it.



▶ SCAN TO WATCH

**DO HUMAN WORK, EPISODE 2: FROM BURNOUT TO BREAKTHROUGH, WITH BENJAMIN DULIEU**

Hear him tell it.

He was a believer in AI early. And he understood something most people pushing AI inside a company won't say out loud. A lot of what he was about to pitch to his team, if you squinted, sounded like telling them to replace themselves.

But he had a conviction. If he could get his small team out of the repeatable work, they'd do the strategic work they had never had the hours for.

The problem with a conviction is that going from "I believe this" to anybody changing what they do on a Tuesday is enormously hard. "You should use AI more" just sounds like more work, from somebody who isn't doing it.

So he called one of his GRC managers. Smart, pragmatic, hungry.

And he asked her one question.

Yeah.

Do you want to go above your max bonus?

Yeah.

Then take the security questionnaires, the ones the team spends hundreds of hours a month answering, and fully automate them.

If you've never had the pleasure: the work has to get done, but it doesn't need a human at a keyboard every single time. The answers already exist somewhere. Somebody just has to assemble and check them.

That was a Thursday night. The next Wednesday she walked into his office. The first thing she said was thank you. She'd like to believe she would have gotten there on her own. She just needed the opportunity.

Five business days, and she had version one. It's an agent running at Duck Creek today, supporting the pre-sales team, and he says it saves them hundreds of thousands of dollars a year. He also says it did something more important than that: it showed the rest of the company what was possible, and people who see it become believers.

///

## WHY THE MECHANIC WORKS

Look at what he actually did, because it's a better change-management design than most consultancies sell.

**He asked a question nobody says no to.** Not "will you adopt AI?" which invites a debate. "Do you want to get promoted?" There's no debate.

**He picked the right person and the right task.** Somebody hungry, and a job everybody in security knows is necessary, repetitive, and miserable. Not a moonshot. One basic use case, chosen because success would be obvious to everyone who'd suffered through it.

**He tied it to something real.** Not a hackathon prize. Her promotion and her bonus.

**He gave her permission.** This is the part I'd underline, and it's in her own words. She didn't need training, a tool budget, or a working group. She needed somebody senior to say: this is your job now, go. "I just needed the opportunity" is what most people on most teams would say if you asked them.

None of that costs money. All of it is available to anybody who manages anybody, this quarter.

///

## THE SUSPICION

Here's why the story lands so hard with people, and why I think everyone reading this can feel all three parts of it.

The first part is the uncomfortable suspicion that the tasks are your value. That if a model can do the security questionnaire, maybe there's less reason for you.

I had a version of this conversation early at 7AI, interviewing a candidate for a sales leadership role. He looked at what our product did and worried that it would be hard to sell, because it was going to take away jobs in security.

So I asked him: if I automated away your need to put notes into Salesforce, would you be out of a job?

"Are you kidding me? Sales is way more than that."

Bingo.

Nobody believes their own job is the data entry. Everybody suspects that about everybody else's.

The second part is going looking for your security questionnaire. And the third part, which is the only one that matters, is what becomes possible once you stop doing it.

> **YOUR VALUE WAS NEVER THE TASKS.**

///

## THE TEN HOURS

Here's a question I ask people in marketing and in security, and I want you to actually answer it before you keep reading.

If I handed you back a quarter of your week, starting Monday, ten hours, what would you do with it?

I've asked a lot of people. Nobody has ever said they'd make more slides. Nobody says update the CRM, or write another status report, or sit in another meeting where we align on alignment.

They say they'd talk to customers. They'd think. They'd build the thing they've been putting off for a year. They'd fix the process everybody complains about.

When we asked Benjamin, his answer was the best one I've heard. Time with his family, and taking apart the inefficient things everybody in security has stared at for their whole career. He boiled it down to one line: making memories and breaking systems.

You already know what your best work is. You've known the whole time. The only question is what you're going to stop doing on autopilot so you can go do it.

///

## THE TASK INVENTORY

The tool for this chapter. An hour, alone, with your calendar from the last two weeks.

**Step 1. List the tasks.** Not your responsibilities. The actual tasks. Everything you did that took more than fifteen minutes, written as a verb and an object. "Built the weekly pipeline report." "Answered the same positioning question from a new rep." "Reformatted the deck for the partner." Aim for thirty.

**Step 2. Run four tests on each.**

1. **Does it repeat?** Same shape, at least monthly.
2. **Do the inputs already exist?** The answers live somewhere: a document, a system, a past version.
3. **Can you say what good looks like?** If you can't describe a correct output, a model can't produce one, and neither can a new hire.

4. **Does it need your judgment every time, or only at the end?** Most tasks need a human to check, not a human to assemble.

A task that passes all four is a candidate. A task that fails the fourth test is probably your actual job, and you should be suspicious of anybody, including yourself, who wants to automate it.

**Step 3. Score the candidates.** Hours per month, multiplied by dread on a scale of one to five. Dread matters. The task you hate is the one you do worst, put off longest, and will be happiest to stop doing, which makes it the one you'll actually follow through on.

**Step 4. Pick the top one.** One. Not a roadmap. Give it a deadline and the permission to drop everything else, the way Benjamin did.

**Step 5. Write down the ten hours.** Before you automate anything, write down what you'll do with the time it frees. If you don't, the time will be filled by the inbox within a fortnight, and you'll have built a crutch.

///

## THE COMFORT CHECK

CHAPTER 16: YOUR OWN JOB

**1. The comfortable assumption.** ~~My value is the work I do every week.~~

**2. The true objective.** Spend my hours on the work only I can do, and know what that work is.

**3. The redesigned approach.** Inventory the tasks, automate the one with the most hours times dread, and decide in advance where the freed time goes.

# 17

CHAPTER SEVENTEEN

# REDESIGNING THE WORK

People say "AI is going to take my job," and the sentence hides four different things inside one word.

> **PEOPLE.** Judgment, relationships, accountability.
>
> **OUTCOMES.** What the business actually needs.
>
> **WORK.** The activities that produce it.
>
> **TASKS.** Research, formatting, synthesis, admin.

AI is eating the task layer. It's starting on parts of the work layer. It isn't eating the outcome, and it isn't eating who owns the outcome.

A job was never a list of tasks. It was a promise to deliver an outcome. The tasks were just how we used to do it.

That distinction sounds comforting, and it's meant to be, but it has a sharp edge. If your job is defined by its tasks, then AI really is coming for it, because the tasks are exactly what's getting cheaper. The only durable version of a job is the one defined by an outcome somebody owns.

Which brings me to the four most expensive words in marketing.

///

## "THAT'S NOT MY JOB"

Listen to how each function in a typical marketing team describes being done.

1. **Demand gen:** we generated the leads and put them in the CRM.
2. **Ops:** we ran the process.
3. **Brand:** we made the campaign.
4. **Content:** we published the asset.
5. **Product marketing:** we shipped the messaging and the enablement.
6. **Events:** we delivered the booth on budget.

Every one of those is true. Not one of them is an outcome.

"That's not my job" is the comfort instinct in its purest form. It draws a line around the tasks you were assigned and declares everything past the line somebody else's problem. It feels like focus. It's actually the reason nobody owns the thing that matters, because the thing that matters always sits on a line between two functions.

Take the most common example. Marketing produces a lead, hands it over, and is done. Sales takes the first meeting. Whether that first meeting is any good, whether the buyer leaves it thinking their time was well spent, belongs to nobody. Marketing says it's a sales problem. Sales says marketing sent them junk.

Here's the trap I want to flag, because I nearly fell into it writing the talk. The obvious redesign is to say the outcome is "pipeline." It isn't. Go back to Chapter 6 and run the test: you can generate a great deal of pipeline while the business gets worse. Pipeline is an activity with a dollar sign on it.

The outcome is **a first meeting worth the buyer's time.** That's something that happens in somebody else's head, in a room you're not in, and you can't achieve it while the business gets worse.

Now notice what changed. Not the tooling. **The definition of done.**

///

## THE BETTER FIRST MEETING, BUILT

Here's what that redesign looks like as a real system rather than a slide. It's also in Appendix F, ready to run.

**The outcome.** The rep walks into every first meeting already knowing who's on the call and why they might care, so the first ten minutes are about the buyer's problem rather than discovering it.

**The sources.** All of these exist for almost every first meeting, and almost nobody reads them, because reading them takes about ninety minutes per meeting:

1. Your own CRM history with the account: past conversations, past opportunities, who touched it and what happened.

2. The attendees' public professional profiles: role, how long they've been in it, where they came from.
3. What they've been posting or speaking about recently.
4. Company news from the last quarter: funding, leadership changes, incidents, acquisitions, layoffs.
5. What changed in their technology stack, where it's visible.
6. Whether they've publicly complained about their current vendor.
7. Your competitive library entry for whoever they use today.

**The output.** One page, delivered the evening before. Not a dossier. Five short sections: who's in the room, what they probably care about this quarter, the one thing they're likely to push back on, the question you should ask in the first five minutes, and what not to say.

**The format rule.** If it's longer than a page, the rep won't read it, and a brief nobody reads is a crutch with extra steps.

**The setup.** An afternoon to write the instructions, then ongoing tuning as reps tell you what was useful and what was noise. The per-meeting cost is small enough not to matter.

**The failure modes.** Three, and you should design for all of them.

1. **Confidently wrong.** A model will happily connect a person to the wrong employer or a stale title. Every claim in the brief needs a source, and the rep needs to know which parts to verify.
2. **Creepy.** There's a line between prepared and invasive, and a buyer can feel it. The brief is for the rep's understanding, not for reciting. "I saw you posted about this" is fine. Reading somebody's last six months back to them is not.
3. **Generic.** If the brief says the same thing about every account, it's the category mean again. The test is whether two briefs for two different companies would be obviously different.

Who owns this? Marketing, now. Because the outcome is the meeting, and marketing is the function with the research skills, the competitive library, and the content to put in it. That's the redesign. Not a new tool. A new line around what counts as done.

///

## WHAT SHOULD NEVER BE AUTOMATED

The flip side, because a chapter that only says "automate more" is half a chapter.

1. **The decision about what's true.** A model can assemble the evidence. A person decides what the company claims.
2. **The decision about what to say, and who says it.** Point of view is the scarce thing in this category, and it can't be generated. It can only be held.
3. **The relationship.** Nobody has ever been loyal to a sequence.
4. **The apology.** If something went wrong, a human says so, in their own words.
5. **The judgment at the end.** Humans on the loop, not out of it. Every system in Chapter 15 ends with a person choosing.

And one more failure mode, which is the most common one I see: **automating a process nobody should be running.** If the gated report sent to a cold list didn't work by hand, it won't work automated, it'll just fail at a scale where the failure looks like activity. Run the Comfort Check before you automate anything. A process that fails question two shouldn't be made faster. It should be stopped.

///

## THE COMFORT CHECK

CHAPTER 17: THE HANDOFF

**1. The comfortable assumption.** ~~Marketing's job ends when the lead reaches sales.~~

**2. The true objective.** A first meeting worth the buyer's time.

**3. The redesigned approach.** Marketing owns the brief behind every first meeting: seven sources, one page, delivered the night before, with a human checking what's true.

V

PART FIVE

# MONDAY

# 18

CHAPTER EIGHTEEN

# SEVEN BELIEFS, SEVEN MONDAYS

---

**BORING IS DANGEROUS.**

Here's everything in this book in one line.

And here are the seven things I'd like you to walk away believing, each one with something to do about it on Monday.

One rule before you start: **one belief per week.** A reader who does all seven on one Monday does none of them. Seven weeks, one exercise each, and you'll have run the whole book by the end of a quarter.

///

## WEEK ONE

# SAFE IS JUST A BET YOU DON'T HAVE TO DEFEND.

**The argument.** The standard program is a bet too, and it fails as often as anything else. It's just quieter when it does, so nobody learns from it.

**Monday.** Write down the safe choice you made this quarter. Then write down what you were betting on when you made it. You made a bet either way. You just didn't have to defend this one.

**Time.** Twenty minutes.

**In thirty days you should be able to point at:** one Bet Card from Chapter 3, filled in before the money moved, with the failure condition written first.

///

## WEEK TWO

# AI DIDN'T MAKE YOU FASTER. IT MADE EVERYONE FASTER AT THE SAME THING.

**The argument.** Every individual piece gets better, and the category gets more alike. Speed toward the middle isn't an advantage when everybody has the same speed.

**Monday.** Take the last thing you shipped with AI involved. Ask whether a competitor with the same tools and the same prompt would have produced something recognizably different. If the answer is no, you didn't get faster. You got closer to the middle.

**Time.** Thirty minutes. Add ten for the distinctive-homepage test in Chapter 13.

**In thirty days:** a forbidden-words list for your category, in use, and one piece of work built from primary material nobody else has.

///

## WEEK THREE

# BETTER IS SUBJECTIVE. DIFFERENT IS MEASURABLE.

**The argument.** Nobody can settle whether your brand got better. Anybody can check whether people remember it.

**Monday.** Cross out every comparative in your messaging. Faster. Smarter. More accurate. More integrations. Whatever survives is your actual position.

**Time.** An hour.

**The free version of measuring it:** the repeat-back test from Chapter 9, five people, one week apart.

**In thirty days:** one number about your brand, with a baseline and a date for the next measurement.

///

## WEEK FOUR

# DESIGN FOR MEMORY, NOT JUST EXPLANATION.

**The argument.** Explanation reaches the person in front of you. Memory reaches the four people who weren't in the room, and in a market where the buying window opens without warning, those four people are the whole business.

**Monday.** Explain what you do in two words. If you can't, you don't have a message. You have a description.

**Time.** An afternoon, which is about what the Camry cost.

**In thirty days:** one concrete analogy, sourced from outside the category, that you've used enough times to be slightly sick of it.

///

## WEEK FIVE

# WHAT WORKED BEFORE IS A CEILING, NOT A FLOOR.

**The argument.** Your best work quietly becomes the upper bound on the next thing, because everything after it gets measured against it instead of against the problem.

**Monday.** Name the play you're about to run again because it worked last time. Now design the version you'd run if that play were banned.

**Time.** An hour, plus the Ceiling Audit in Chapter 11 before your next planning cycle.

**In thirty days:** one thing that was working, killed on purpose, and something new in its slot.

///

## WEEK SIX

# YOUR VALUE WAS NEVER THE TASKS.

**The argument.** Nobody believes their own job is the data entry. Everybody suspects it about everybody else's.

**Monday.** Find the security questionnaire in your job. The thing that has to get done, that doesn't need you to be the one doing it every time.

**Time.** The Task Inventory in Chapter 16 takes an hour. Benjamin's version took five business days.

**In thirty days:** one task gone, and a written answer to what you did with the hours.

///

## WEEK SEVEN

# YOU DON'T NEED THE BUDGET TO OWN THE THINKING.

**The argument.** The Camry cost nothing. Benjamin's promotion question cost nothing. The Comfort Check costs nothing. The expensive part of this work has never been the money. It's the willingness to own an outcome nobody assigned you.

**Monday.** Pick one process nobody owns the outcome of. Redesign it backward from the outcome.

And then don't ask.

**Time.** As long as it takes.

**In thirty days:** something that's better because you decided it was yours.

///

## THE COMFORT CHECK

CHAPTER 18: THE NEXT SEVEN WEEKS

**1. The comfortable assumption.** ~~I'll get to this after the busy season.~~

**2. The true objective.** Change one thing about how I work, every week, for seven weeks.

**3. The redesigned approach.** One belief per Monday, one exercise, one thing to point at in thirty days.

# 19

CHAPTER NINETEEN

# STOP LETTING HIM VOTE FIRST

---

**YOU JUST HAVE TO STOP LETTING HIM VOTE FIRST.**

## At the start of the talk this book came from, I asked a room full of people to do me a favor. Every time they felt the urge to nod, notice it. Don't stop it. Just notice it.

If you've read this far, you've nodded a few times. Some of those were you. Some of those were him.

He isn't going anywhere. Neither is mine, and I have more of him than most of you. Tomorrow morning I'm going to open a document, and he's going to lean over and point at the thing that worked last time, and it's going to feel fine before I've had a chance to think about it.

He kept your ancestors alive. He's extremely good at that job.

He's terrible at making you interesting.

///

The first page of this book made a claim before I'd given you any reason to believe it. In Chapter 18 I softened it to "boring is dangerous."

That was me being polite. Dangerous is a business word. It lets you file this under risk management and get to it next quarter.

Here's what I actually think. Every other mistake in this job is survivable. Wrong channel, wrong message, wrong timing, wrong number. I've made all of them, and I spent a whole chapter on one that cost real money. You can fix those.

The one you don't come back from is making something nobody remembers. There's nothing to fix. Nothing happened. You spent the budget, you filled the calendar, you hit every deadline, and the market never noticed you were there.

# BORING IS THE ONLY UNFORGIVABLE SIN.

///

So here's the only thing I want from you. Not courage. Not budget. Not permission from your board.

You don't have to fire him. He has a seat in every meeting for the rest of your career.

That's it. The half second between the instinct and the decision. Notice the pull toward the thing that already worked, and ask the question underneath all seven beliefs: what am I actually trying to make happen here?

Remember the ten hours? You already knew what you'd do with them. You already know what your best work is.

The only thing standing between you and it is a guy who's been dead for forty thousand years and still gets to vote first.

Don't let him.

///

One last thing.

The kid on that yearbook page had no idea what he wanted to do. The path worked out anyway, mostly because of the one thing he did notice: when something got too easy, it started to feel wrong.

If I could give him one piece of advice, it wouldn't be to find something he's good at.

It'd be this. When it feels easy, when it feels fine, when the whole room nods, get suspicious. That's where the good work starts.

Go do your best work ever.

# A—G

APPENDICES

# THE TOOLS, COLLECTED

Everything you'll want to copy, collected in one place.

A

APPENDIX A

# THE COMFORT CHECK

## THE WORKSHEET

> ### The thing we're about to do: ____
>
> **1. What's the comfortable assumption?** The belief underneath this that nobody voted on. *Test: if you can name the meeting where it was decided, it isn't one.*
>
> ///
>
> **2. What's the true objective?** A change in somebody else's head or behavior. *Test: if you could achieve it completely while the business got worse, it's an activity.*
>
> ///
>
> **3. What's the redesigned approach?** Built backward from question two. *Rule: write at least two, and one has to be an answer you don't like.*
>
> 1. ///
> 2. ///

Run it before the work exists, not after. Ask question two out loud and question one to yourself.

## EVERY CHECK IN THIS BOOK

| CH. | SUBJECT | THE ASSUMPTION | STRUCK? |
|---|---|---|---|
| 1 | The review meeting | When the room agrees, the work is good | Yes |
| 2 | The category audit | We know what competitors say because we track them | Yes |
| 3 | The annual plan | The standard program is the low-risk option | Yes |
| 4 | Frameworks | A method has to be complete to be useful | Yes |

| CH. | SUBJECT | THE ASSUMPTION | STRUCK? |
|---|---|---|---|
| 5 to 7 | Axonius | Bigger company, bigger problem, bigger budget. Sell enterprise | Yes |
| 8 | The dream car | A big enough prize gets us the right people | **No. I never asked.** |
| 9 | Proving it worked | Brand can't be measured | Yes |
| 10 | The three-minute pitch | You win a rubric by answering the rubric well | Yes |
| 11 | The second year | It worked, so run it back | Yes |
| 12 | The booth wall | The headline's job is to say what we do | Yes |
| 13 | Writing with AI | If each piece is better, the work is better | Yes |
| 14 | AI adoption | Adoption and hours saved mean it's working | Yes |
| 15 | Building with AI | Using AI well means buying the right tools | Yes |
| 16 | Your own job | My value is the work I do every week | Yes |
| 17 | The handoff | Marketing's job ends when the lead reaches sales | Yes |
| 18 | The next seven weeks | I'll get to this after the busy season | Yes |

B

APPENDIX B

# THE SAMENESS AUDIT

Ninety minutes. Full method in Chapter 2.

1. **Pick the set (10 min).** Eight to twelve companies a buyer meets in the same week as you: three or four direct competitors, two or three adjacent vendors, two or three from the last sponsor list, and your own company.
2. **Capture (20 min).** Same browser, width, zoom and day. Light mode. Hero to just below the fold.
3. **Score four axes (40 min).** Companies as rows.
    1. *Vocabulary:* every phrase in each headline and subhead, with a count of how many companies use it. Three or more means invisible.
    2. *Structure:* hero over image or gradient, short subhead, paired CTAs, logo bar, three-column benefits, stat band, testimonial carousel. Tick each.
    3. *Navigation:* top-level items, verbatim.
    4. *Proof:* what each offers as evidence. Note what almost nobody offers.
4. **The swap test (10 min).** Six most similar heroes, logos stripped, shown to somebody outside marketing. Ask them to name the companies.
5. **Write the sentence.** The one thing your category agrees on without ever voting. That's what you're working against.

C

APPENDIX C

# THE BET CARD AND THE MEASUREMENT PLAN

## THE BET CARD

Fill it in before the money moves. Write question five before question four.

1. **What we're betting.** One sentence, no marketing nouns.
2. **What we're betting against.** The alternative use of this money and quarter, named.
3. **What has to be true.** Two or three falsifiable conditions.
4. **What would prove it worked.** A number, a source, a date.
5. **What would prove it failed.** Written on its own, from scratch.
6. **When we look.** A date, already in the calendar.
7. **Who decides.** One name.

## THE MEASUREMENT PLAN

1. **Objective.** The change in somebody else's head, as a delta.
2. **Baseline.** Today's number, with date and method.
3. **Leading indicators.** What should move within weeks.
4. **Lagging indicator.** What settles it, and when.
5. **Kill.** The number and date that mean we stop.
6. **Continue.** The number that means we keep going as is.
7. **Scale.** The number that means we put more behind it.
8. **Who calls it.** One name.

## FIVE FREE MEMORY TESTS

1. **Repeat-back.** Tell five outsiders what you do, ask them a week later.
2. **Discovery call count.** In twenty recorded first calls, how many prospects describe you in your words, unprompted.

3. **Open-text field.** "How did you hear about us," as a free-text box, read monthly.
4. **Hallway test.** Ten people at the next conference, before you say anything.
5. **Anonymized swap.** Your hero, logo removed, shown to ten people in your market.

D

APPENDIX D

# THE ASSUMPTION INVENTORY

The forty assumptions and counter-questions are in Chapter 5, grouped by function. Here are the four ways to find your own.

1. **The "just how it works" sweep.** Thirty minutes with the team. Everyone writes five, silently. Anything on three or more lists is an operating assumption.
2. **The new-hire question.** In week three: what have we done that you wouldn't have done at your last company, and what didn't we do that you expected? Write it down. Don't defend anything in the room.
3. **Budget archaeology.** Every line above your threshold: what was this originally for, and is the person who first funded it still here?
4. **The calendar audit.** For every recurring item: what made that the right interval?

Keep the list somewhere you'll see it at planning. Most of what you find will be fine. The point is to have it written down for the day conditions change.

E

APPENDIX E

# DERIVING AN UNPUBLISHED RUBRIC

For any scored situation where the criteria exist but aren't written down: a pitch contest, an analyst evaluation, a shortlist, an interview loop. Full method in Chapter 10.

1. **Collect ten to twenty past artifacts.** Winners and losers.
2. **Extract the questions each section answered,** not the answers.
3. **Count.** The questions in every artifact are the real rubric. The ones only in the winners are the interesting part.
4. **Write the competent version of every answer.** That's the floor. You have to clear it.
5. **Satisfy the floor fast and spend the remainder on one thing built to be repeated.**

The five questions I found in past Innovation Sandbox pitches:

1. What's the problem?
2. Why is it big?
3. What's genuinely new?
4. Why you?
5. Can you survive the floodgates?

In three minutes, the one repeatable thing got about twenty seconds. That was enough.

F

APPENDIX F

# THE SKILL LIBRARY

These are written to be copied and run, not admired. Each one is a skill: a document a model follows. Paste one into whatever AI tool your team uses, replace the bracketed parts with your own sources and rules, and run it. Then rewrite it every time it gets something wrong, because the first version will.

Every skill here ends with a human decision. Keep it that way.

///

## F1. THE MONDAY REVIEW

```
NAME: monday-review
WHEN: Every Monday before 8:00 AM.

GOAL
Give me a proactive read on where we are against our goals, how spend is
pacing, and what we should be doing that nobody has asked for yet.

SOURCES (read all of them; say which ones you could not reach)
1. [CRM]: pipeline by stage, deals created and closed in the last 7 days,
   deals with no activity in 21 days.
2. [Internal search / wiki]: anything published or changed in the last 7 days
   about customers, competitors, launches or plans.
3. [Docs folder]: this quarter's goals and budget.
4. [Last week's review], so you can say what changed.

STEPS
1. State this quarter's goals in one line each, from the goals doc.
2. For each goal: on pace, at risk, or off pace, with the number and its
source.
3. Spend: planned vs actual to date, by program. Flag anything more than
   15% off plan in either direction.
4. What changed since last Monday that matters. Maximum five items.
5. Opportunities nobody asked for: up to three things the data suggests we
   should build, stop or change. For each, the evidence and the smallest test.
6. Questions I should ask this week, and who to ask.

RULES
- Every number needs a source. If you can't source it, say so; don't estimate.
- Separate what the data says from what you infer. Label inferences.
- Assume the CRM has stale stages and optimistic deals. Note anything that
  looks like it.
- No more than one page.

OUTPUT
```

```
One page: Goals / Spend / What changed / Opportunities / Questions.

HUMAN DECISION
I decide which opportunities to pursue. You propose; you don't assign work.
```

///

## F2. THE VOICE PROFILE BUILDER

```
NAME: voice-profile-builder
WHEN: Once per person, then updated every time they correct a draft.

GOAL
Build a profile that lets a model draft in [PERSON]'s voice well enough that
they edit rather than rewrite.

SOURCES (in order of value)
1. Their corrections to drafts: the before and after of every sentence they
   rewrote. This is the most important input.
2. Unpolished writing: chat messages, quick emails, comments.
3. Transcripts of them talking: podcasts, talks, recorded calls.
4. Published writing LAST, because someone else probably edited it.

PROFILE SECTIONS
1. Register: how formal, how long the sentences run, how they open and close.
2. Always: the moves they make constantly.
3. Never: words, phrases and structures they never use or always remove.
4. Punctuation: how they actually punctuate, including what they avoid.
5. Point of view: the positions they hold and repeat.
6. Imitation tells: the things a copycat overdoes. Signature phrases, used
   sparingly. Note how often each one really appears.
7. Ten example sentences, verbatim, that sound most like them.

RULES
- Most of a real voice is ordinary sentences. Say so in the profile.
- Every rule needs an example from the sources.

HUMAN DECISION
The person reads the profile and corrects it. Their corrections go in.
```

///

## F3. THE WEEKLY CONTENT SCOUT

```
NAME: weekly-content-scout
WHEN: Every Monday.
```

```
GOAL
Find the stories this week that [PERSON] has a real point of view on, and
propose a short plan. Do not summarize the news.

SOURCES
1. [List of publications and feeds]. Include at least three that competitors
   probably don't read.
2. [PERSON]'s voice profile and their last ten posts.

STEPS
1. Collect the 20 most relevant stories of the last 7 days.
2. Apply the protagonist filter to each: can [PERSON] speak to this from
   something they've actually done or seen? If they'd only be repeating the
   story, drop it.
3. Rank what survives. Rank the strange story above the big one if both pass.
4. Pick two or three. For each: the story, the angle in one sentence, why it's
   theirs to say, and the obvious take everyone else will post (so we avoid
it).
5. Draft each post against the voice profile.

RULES
- Use the forbidden-words list: [LIST].
- No post may be a summary of the article.

HUMAN DECISION
[PERSON] picks, edits and posts. Nothing goes out automatically.
```

///

## F4. THE FIRST-MEETING BRIEF

```
NAME: first-meeting-brief
WHEN: The evening before every first meeting.

GOAL
The rep walks in already knowing who's in the room and why they might care.

SOURCES
1. [CRM]: history with this account and these people.
2. Attendees' public professional profiles: role, tenure, previous companies.
3. What attendees have posted or said publicly in the last 90 days.
4. Company news from the last quarter.
5. Visible changes to their technology stack.
6. Any public complaints about their current vendor.
7. [Competitive library] entry for their current vendor.

OUTPUT (one page, five sections)
1. Who's in the room, one line each.
2. What they probably care about this quarter, with evidence.
3. The one thing they're likely to push back on.
```

```
4. The question to ask in the first five minutes.
5. What not to say.

RULES
- Every claim has a source link. Mark anything you're unsure of.
- Check that names, titles and employers match across sources. If they
  conflict, say so.
- This is for the rep's understanding, not for reciting to the buyer.
- If this brief would read the same for a different company, it's too generic.
  Rewrite it.

HUMAN DECISION
The rep decides what to use and verifies anything marked unsure.
```

///

## F5. THE ANTI-FLUENCY PASS

```
NAME: anti-fluency-pass
WHEN: On any draft before it goes to review.

GOAL
Find where the draft sounds like the rest of the category, and fix it.

STEPS
1. Write the most typical version of this piece for a company in [CATEGORY].
   Label it THE MEAN.
2. Compare the draft to THE MEAN. List every sentence that a competitor could
   publish unchanged.
3. Flag every word from the forbidden list: [LIST].
4. Flag every comparative: faster, smarter, better, more, easier.
5. As a skeptical buyer who has read eleven competitor pages this morning,
   name the three sentences you'd skip and the one you'd remember.
6. For each flagged sentence, propose a rewrite that uses something specific:
   a number we can source, a customer's words, a concrete image, or a claim
   someone could disagree with.

RULES
- Don't make it smoother. Smoother is the problem.
- Don't invent facts to make a sentence specific. Ask for them instead.

HUMAN DECISION
The author decides which rewrites to take.
```

G

APPENDIX G

# SOURCES

## RESEARCH CITED


1. Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, 9(2, Pt. 2), 1 to 27. Link (Chapter 1)
2. Reber, R., & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness and Cognition*, 8(3), 338 to 342. Link (Chapter 1, easy-to-read statements)
3. Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior*, 16(1), 107 to 112. Link (Chapter 1, repeated claims)
4. McGlone, M. S., & Tofighbakhsh, J. (2000). Birds of a feather flock conjointly (?): Rhyme as reason in aphorisms. *Psychological Science*, 11(5), 424 to 428. Link (Chapter 1, rhyme)
5. Alter, A. L., & Oppenheimer, D. M. (2006). Predicting short-term stock fluctuations by using processing fluency. *PNAS*, 103(24), 9369 to 9372. Link (Chapter 1, pronounceable names)
6. Nemeth, C. J. (1986). Differential contributions of majority and minority influence. *Psychological Review*, 93(1), 23 to 32. Link (Chapters 1 and 8)
7. Nemeth, C., Brown, K., & Rogers, J. (2001). Devil's advocate versus authentic dissent: Stimulating quantity and quality. *European Journal of Social Psychology*, 31(6), 707 to 720. Link (Chapters 1 and 8)
8. DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147 to 160. Link (Chapter 2)
9. Spranca, M., Minsk, E., & Baron, J. (1991). Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*, 27(1), 76 to 105. Link (Chapter 3)
10. Lord, C. G., Lepper, M. R., & Preston, E. (1984). Considering the opposite: A corrective strategy for social judgment. *Journal of Personality and Social Psychology*, 47(6), 1231 to 1243. Link (Chapter 4)
11. Paivio, A. (1971). *Imagery and Verbal Processes*. Holt, Rinehart and Winston. (Chapter 10)

12. Luchins, A. S. (1942). Mechanization in problem solving: The effect of Einstellung. *Psychological Monographs*, 54(6). Link (Chapter 11)
13. Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28). Link (Chapter 13)

## CONVERSATIONS AND MEDIA

1. Benjamin Dulieu, Duck Creek Technologies. *Do Human Work*, Episode 2. (Chapter 16)
2. RSAC Innovation Sandbox 2019, the Axonius pitch. RSAC library. (Chapter 10)
3. Axonius, Controlling Complexity campaign. Ruckus case study. (Chapter 7)
4. Quish Splash research. 7AI blog. (Chapter 11)
5. CRXfiltrate research. 7AI blog. (Chapter 11)
6. My 2022 CyberMarketingCon keynote clips. nathanwburke.com/speaking. (Chapters 3, 6 and 8)

GO DO
YOUR
BEST
WORK
EVER.

# BORING IS THE ONLY UNFORGIVABLE SIN.

Every other mistake in this job is survivable. The one you don't come back from is making something nobody remembers. This is a field guide for cybersecurity marketers who would rather be remembered than approved: the science of why good teams make forgettable work, a three-question check you can run on anything, and the stories, tools and AI systems behind it.

---




**READ IT FREE**

**NATHANWBURKE.COM/RESIST**