



# The Fox Advantage

Paul Dervan

# Contents

## \part\_01 - what just collapsed

- 01 The marketing department autopsy report

---
- 02 The robots aren't coming. They've already moved in

---
- 03 The algorithm will see you now

---
- 04 Everything. Everywhere. All at once

---
- 05 Average is the new invisible

---

## \part\_02 - better together

- 06 Drivers Wanted

---
- 07 Robots don't have skin in the game

---
- 08 Kill bugs fast

---
- 09 Critical thinking has never been more critical

---
- 10 AI has turned us all into little Hemingways

---
- 11 The friction is the point

---
- 12 A very tidy wrong answer

---

# 01

## The marketing department autopsy report

I was in a planning meeting recently that took three hours. Three hours to agree not more than a headline. By the time we'd finished, a competitor had already shipped something similar and moved on to the next thing.

This is what we built. We built marketing like a civil service. Layered. Titled. Ritualised. The strategist thinks. The creative director adds vision. The copywriter adds words. The designer adds polish. The social team schedules it. The compliance team deletes half of it.

Three weeks later, something stumbles out the door. It's late. It's beige. It cost a fortune. Meanwhile, a 23-year-old with Claude and Midjourney shipped three campaigns before we agreed if the headline should say "empowering" or "enabling".

We didn't just build a slow system. We built one optimised for a world that no longer exists. Like a loyalty card for a DVD rental shop.

Let's not call this a disruption. Let's call it what it is. Small teams with tools are now moving faster than big teams with process. While some are still rewriting the brief, others have already tested five ideas and moved on. The specialist stack wasn't streamlined. It was sidelined.

We used to pay for execution. Design. Craft. Copy. Now most of that is cheaper, faster and often, if I'm being honest, often better when done by machine. If your value is tied to what you produce, you've probably already been replaced. You just haven't been told yet.

We were trained to be hedgehogs. Pick a lane. Stay there. Climb the ladder. But the foxes, the curious, cross-trained, judgment-led ones, are moving faster, breaking rules, and skipping the whole approval chain.

AI makes average execution free. So the value has moved. You're not paid for your polish anymore. You're paid for your taste. Not whether you can write the ad, but whether you can tell the difference between an ad that works and one that just fills space.

Your job now is to spot the tension before the brief. To choose the 10x idea over the 10th revision. To move fast without losing the plot.

So how do we fix this? That six-person social pod? It's now two generalists with taste, AI tools, and the permission to act. The old model was "I do my bit, then pass it along." The new model is "I can make the call, do the work, and hit publish before lunch." This isn't about finding unicorns. It's about removing the need for unicorns.

Our seven-layer approval chain isn't quality control. It's momentum loss. Sharp teams set boundaries, act fast, and review while it still matters. Not "Is this perfect?" but "Did it work?" As my dad, a doctor, used to say: "Don't wait for the autopsy to fix something you could've caught in triage."

Nobody cares how many banners we shipped. Start tracking how fast you move from insight to action. That's the real metric now. Every handoff slows you down. Every sync meeting is a tax. Every job title that starts with "liaison" should come with a warning label.

Smart teams won't make coordination more efficient. They'll make it unnecessary.

# 02

## The robots aren't coming. They've already moved in

I watched a junior marketer on the Luas, building a campaign in the notes app. He had one earbud in, the other dangling, and he was flicking between a prompt, a doc, and a half-finished landing page like it was nothing. At one point the train jolted, his thumb hit send, and the tool still spat back three usable headlines before we got to the next stop. It was weirdly calm. Not heroic. Just normal.

That's what messes with your head. The work isn't brilliant, but it is fast enough to change who needs to be in the room.

For years we built marketing like a relay race. Strategy hands to creative. Creative hands to design. Design hands to web. Web hands to analytics. Someone somewhere produces a deck to prove it's all connected. Each handoff adds cost, time, and a new chance to dilute the point.

Now a lot of the relay baton work is done by software. Drafts, variations, basic layouts, summaries, first-pass analysis, the endless "make it shorter, make it punchier, now do ten versions." That used to be paid labour. Now it's a button. So the threat isn't that "robots will take your job" in some dramatic overnight way. The threat is that the middle of the org chart, the part built to move work along and tidy it up, gets thinner because there's less to move and less to tidy.

You can feel it already in the kind of tasks that used to justify whole roles. First draft copy. Basic competitor tables. Social calendars. Slide clean-up. Even the polite meeting notes that everyone pretends to read. They can still be done by people, but the economic argument for needing so many people to do them has cracked. Most teams haven't faced the real implication yet.

They're still improving workflows that exist mainly because execution used to be slow, expensive, and annoying. They're still hiring for narrow slices of output, even as output becomes the cheapest thing in the room.

This doesn't mean the work goes away. It means the shape of the work changes. A smaller team can now ship a first version without waiting two weeks for the slot, the resource, the review, the rewrite, the reschedule. That speed is not automatically good, but it is real. And it forces a question most marketing teams don't like answering, which is: if the work can be produced quickly, what exactly are we being paid for. The answer isn't "more content." Nobody is crying out for more content. The answer is direction. Taste. The ability to make a call under uncertainty and stand over it. The ability to say no to the shiny wrong thing. The ability to spot the quiet risk before it becomes a public one.

Software will happily produce a hundred plausible routes. It will not carry the consequences of choosing the wrong one. It won't sit in the room when sales are down and the board asks why. It won't be the person whose name is on the decision.

That's why teams that win won't just "use AI." Everyone is doing that now. The teams that win will rebuild around fewer handoffs and clearer ownership. Fewer people passing work along, more people capable of taking a piece end to end, from question to draft to test to fix.

The robots have already moved in.

# 03

## The algorithm will see you now

A lad beside me in a café tapped “Add to basket” and stared at a spinning wheel. Four seconds, maybe five, then he stood up, left the coffee, and walked out like the website had slagged him off. The barista had written “DAVE?” on his cup in thick marker, all caps.

It’s not impatience, it’s training. We’ve all been trained by feeds and apps and one-click everything to treat friction like a fault, and we punish it instantly.

Offline, you used to fight for eye-level shelf space. Get into Tesco, pay for the end cap, make sure the packaging holds up under fluorescent lights, job done. Online, your customer types “best” and two other words, and a system decides what they see, in what order, and whether you even deserve to be there.

The brutal part is how quiet it is. There’s no manager telling you you’re underperforming. There’s a small drop in click-through, a few extra bounces, a handful of people who don’t wait for the page to load, and you’re out of the running before your brand story has even turned up.

My dad used to say a rushed diagnosis is just a polite misdiagnosis. Algorithms don’t do polite. They do fast. They make a call in microseconds off scraps of behaviour, and they never ring you afterwards to explain themselves.

And your logo, your long heritage, the year you were founded, the awards you won last year, none of it buys you much when the page takes a beat too long to appear. If the site feels like it’s running on dial-up, you’re not “losing attention.” You’re failing an entry test. A neighbour

of mine is convinced marketing is still mostly a sharp logo and the right golf handicap. Lovely man. He's arguing with a machine that doesn't play golf and doesn't care who you know.

So what do you do with that, other than have a small cry into your tea. Start with the boring stuff you've been ignoring because it didn't feel like "marketing." Speed. Clarity. The obvious answers in plain language. The bits that let a tired person get what they came for without working hard. If you want to see how your brand really performs, don't do the grand brand workshop.

Open your own site on bad Wi-Fi, on an old phone, with your thumb. Try to buy. Try to find the one thing customers always ask you in support tickets. Count how many times you swear.

Next, assume you are being scored all the time. Every delay, every confusing label, every dead-end page is a little vote against you. People don't leave you a review saying "your checkout was fiddly." They just vanish. This is where AI is useful in a very unsexy way. Let it help you spot the leaks. Summarise the complaints. Cluster the "I couldn't find" messages. Rewrite the help article so a human can skim it, and so a machine can lift the one sentence answer without guessing. Use it to clear the debris.

But don't fall into the trap of thinking the algorithm is your customer. It's not. It's the bouncer on the door. Passing the bouncer gets you into the room. Once you're in, you still have to be worth staying with. The machine can decide what shows up. It can't decide what gets remembered.

That's still the human job. Not "make content." Make meaning. Give people a reason to recognise you again when the options blur together, and the shortlist looks the same.

# 04

## Everything. Everywhere. All at once

My mate bought a vacuum between TikToks. Twelve minutes from "I need a hoover" to "done, bought." No trip to Carrickmines, no asking friends which one they'd recommend, no proper comparing of models. Just need, algorithm, click, confirm.

Funny thing is it feels normal. Search engines figured out what he needed. Reviews told him which one worked. Price alerts told him when to buy. The whole thing happened in the time it takes to scroll past three cat videos.

How many categories is this happening in? If your brand isn't in the algorithm's shortlist, you're not even in the conversation.

Most marketing departments are still planning like it's 2019. We map out customer journeys that take weeks. We optimize funnels that assume people research and consider and evaluate. We compete against brands in our category, on the same shelf, pitching to the same buyer persona.

Yet my nine-year-old twins won't even consider restaurants under three stars. Not because they've had bad experiences. They just expect the algorithm to do the choosing for them. When they need to know something, they open my ChatGPT app before they google. That's their normal.

I'm watching a generation grow up where human choice and algorithmic filtering have blurred into something I don't fully understand yet. And I think most of us are missing what this means.

Visibility used to be about reach. Buy enough media, get in front of enough eyeballs, build enough awareness. As if that wasn't hard enough, now the algorithm decides who gets seen. And it optimizes for signals we barely track: speed, reviews, clicks, freshness, engagement, behavioural data we can't see.

Your beautifully crafted positioning statement? The algorithm didn't read it. Your brand strategy deck? Irrelevant. The system cares about whether people click, how long they stay, what they do next.

The brands that win aren't the biggest or loudest. They're the ones with better ratings, faster loading times, clearer descriptions, more recent reviews. The algorithm doesn't care about your brand equity. It cares about click-through rates and user behaviour.

My mate didn't choose between Dyson and Shark and Henry. The algorithm chose three options. He just clicked the middle one. Not a bad decision. I've a Shark.

My neighbour still drives to Currys, asks seventeen questions about hoovers, then buys whatever the salesman recommends. Purchases still happen that way, slowly, with consideration, with human advice. But for how long?

The exceptions are spreading. And when they happen, everything we think we know about customer behaviour breaks down.

I won't lie. I'm not that pleased about all of this. But it's not hype. Of course there's AI hype, but it is also real.

Specs won't save you when this happens. Everyone has the same functional benefits. The algorithm will surface three options that solve the same problem, at similar prices, with similar ratings. The only thing left to compete on is meaning. Emotional relevance, clarity of stance, something that makes people remember you when the options blur together.

Three things seem clear. We need to understand how algorithms actually decide, not just reach and frequency, but the behavioural signals that drive visibility in the first place. We need to make meaning our competitive advantage, because when everything else is identical, emotional relevance is what breaks the tie. And we need to build for both worlds, because some customers still take weeks to decide while others choose between TikToks.

For now, the diagnosis matters more than the solution. We're not just competing with other brands anymore. We're competing with whatever the algorithm thinks is more relevant at that moment.

# 05

## Average is the new invisible

My family shops like extremists. Not because we're picky, or because we're cheap, but because algorithms trained us to forget that average things exist.

Our shopping basket tells the story: top-shelf pistachios next to rock-bottom cereal. Fancy runners beside discount flip-flops. Expensive shampoo (not mine) and €1.99 body wash. We're an ordinary, average household that ignores the average. By design, but not our design.

We're not alone. Google Maps has a button that lets you filter out anything below 4.5 stars. Just tap it. Every decent-but-not-exceptional business vanishes. That's not a search feature. That's a Delete Average button.

Seth Godin talked about being remarkable decades ago, with his wonderful Purple Cow book.

Clearly the robots have been swotting up on their marketing and decided to just go ahead and implement this reality. Seth's insight twenty-plus years ago: be remarkable or be invisible. I remember it. It's a good book. His general thesis was that the way to win was to create something worth talking about, a purple cow in a field of ordinary brown ones. Most of us don't. Because it's hard. Bit like fame campaigns. They are the most profitable, but not easy to pull off. So while we settle for ok, our robot friends are off building a world for us where only purple cows survive.

Amazon's first three search results capture 64% of clicks. Products with 4.5+ stars get three times more attention than ones slightly less impressive, the ones that get between 4.0 and 4.4 stars. Their Amazon Buy

Box dominates sales but only rewards extremes: lowest prices or highest performance.

This creates a squeeze. If you're not the cheapest option, you better be exceptional at delivery speed, customer service, and maintaining high ratings. There's very little room for "decent price with ok service." No middle ground. Just best or cheapest. No room for beige cows.

Long run, this might be a good thing. For one, let's stop testing into sameness. A/B testing has its value. I spent a decade in direct marketing. But this type of testing is about efficiency, not necessarily effectiveness. To stand out, we need remarkable, unforgettable, even weird.

Every time we optimize for "good enough," we're optimizing for invisibility. The middle wasn't murdered. It was simply made inefficient. And efficiency is all the robots care about.

This is the end of Part 1.

This was about the collapse. Our marketing structure collapsed. The layered, titled, ritualized system we built for a world that no longer exists.

Our execution advantage collapsed. AI made average production free, shifted value from polish to judgment. Our visibility collapsed. The algorithm became the gatekeeper. We lost control.

Now we compete with expectations, not just competitors. Our competition collapsed. We're not fighting in categories anymore. We're fighting for algorithmic shortlist status. The system decides who gets seen.

And the middle collapsed. The once-safe compromise between quality and cost where most brands lived.

\part\_02

# Better Together

# 06

## Drivers Wanted

AI doesn't replace great marketing. It makes average work free.

Nobody wants to make average work. Yet most is. We're worried about AI stealing our jobs when the real problem is simpler: AI just made mediocre marketing easy, fast and almost free. Not bad marketing. Mediocre marketing. The decent campaigns. The "good enough" strategies. Anyone can now produce passable content. Reasonable copy. Decent analysis.

What used to differentiate marketers from the competition is now available by subscription. This first hit me last September (2024) when I started seriously experimenting with Claude on marketing campaign ideas. What we got to, was pretty decent. Was it better than the best people I've worked with, doing their best work? No. But it was better than a lot of stuff I've seen and done throughout the years.

Somehow, we went from being called marketers, creatives, researchers, to being labeled "humans in the loop". Nobody used to say "Human decision required". We just called it work. If we're not intentional about this, we'll end up the backup singers in a song written by software. "Don't worry, there's still a human involved" is the professional version of "We swear someone checked this." This is absurd. This framing will rewrite our value before we've had a say.

The goal isn't to protect our role from machines. It's to define it. Own it. Drive it. Because when you strip away the hype, the simple truth is AI is brilliant at doing. But it's still bad at deciding. It doesn't know what matters. It just knows what's next.

So the question isn't "How do we keep humans in the loop?" It's "What should never be outsourced?" That's your real value. The things that only a human, a sharp, strategic, culturally attuned human, can bring. But

even "human value" is vague. What's yours? It could be taste. Timing. Ethics. Cultural nuance. Leadership. Emotional fluency. Be specific.

Two weeks ago, I watched this play out in real time. I put my phone on a client's desk, on speaker, and asked Claude to run a discovery call. We were physically in the client's office in Dublin. Two humans and one robot. Claude ran the meeting, and politely worked through a bunch of strategy-type questions, probing and pushing when not satisfied with the answers. Occasionally it didn't hear every reply, but overall, was pretty impressive.

"What's your core challenge?" "People don't know we exist." "How do current customers find you?" "Word-of-mouth mostly." "What's your customer acquisition cost for word-of-mouth versus paid channels?" "We don't really do paid channels." "Why not?" "Haven't needed to."

Claude kept drilling. Revenue per customer. Lifetime value. Market size. Competitive positioning. When answers got vague, it asked for specifics. Numbers. Examples. Did funnel conversion calculations on the spot. Even offered a decent diagnosis. Wrote up the entire meeting notes before I'd left the building.

Overall, both client and I were impressed with the quality of the meeting and the experience of this robot talking away with us. It's insane to me that it could do what it did. Claude was happy enough too apparently.

But it was also clear to me what it could not do. Claude needed to know which frameworks to use, when to push different directions, when to move on. AI doesn't replace the strategy work. It executes it. But someone still has to know what we should be doing, why, and how.

That's the role we must claim. Not a robot wrangler. Not a compliance officer. Driver.

# 07

## Robots don't have skin in the game

If you've never killed creative work that tested well, you've never really been accountable for the outcome.

About four years ago, when at The National Lottery, I was looking at System1 research results for our latest campaign. They suggested the campaign would be effective. Not a homerun, but better than the average. Definitely good enough.

But I killed it.

This was a difficult decision. My colleagues were waiting for this campaign. Sales were down. This was a bet we were counting on. I didn't kill it because it wasn't our normal 4 or 5 star ads. I'm a practical marketer. Not everything is a homerun.

What was niggling was some data points that suggested that a small minority would be offended by the advertising. I'm not precious and have been in the game long enough to not worry about keeping everybody happy.

Looking at the research, the ad was still a decent bet. But Irish people feel a certain sense of ownership over The Irish Lottery. It has political oversight. A small minority could create a lot of noise. So I pulled the campaign.

We like to believe that decision-making is rational. Sure, there are frameworks, first principle thinking, and techniques to help us. But we're kidding ourselves if we think we make decisions like robots. For good or bad, we have skin in the game.

Our experiences influence how we see things. They influence our decisions. Pulling that campaign wasn't fun. I went in front of the company, explained my decision, apologised and articulated how I'd fix

this. I had to make a decision between how likely that this will blow up, versus how painful is pulling this going to be for me. And decide what is the best outcome for the business.

I use Claude and ChatGPT every day. I share problems I'm working on. Brand challenges. Positioning thinking. I've trained them to use knowledge I use. Frameworks I like. I even have a doc called "How Dervan thinks about strategy" they use. If anybody wants it, let me know.

Claude and ChatGPT are both really useful in this area. They spot problems. They suggest alternatives. They help when I get stuck. They push me if I ask them to. But I can't delegate the decision to them.

They can't feel the weight of the decision. And you need that to make the decision.

# 08

## Kill bugs fast

Can AI devise advertising strategies? I ran an experiment.

A few months ago, I gave Claude one of my favourite case studies to tackle. Jon Steel's Porsche turnaround story from the 1990s. Porsche sales had tanked 59% by 1993. There were several reasons for the decline. But one was this: People still loved the cars, but they hated what they represented. Wall Street greed. Gordon Gekko types. 80s excess. They didn't want to be seen as the kind of person that drives a Porsche.

Among other things, Porsche came up with a very successful advertising campaign that helped reverse this. I wondered how Claude would tackle it.

Claude suggested five approaches:

1. Reconnect with engineering heritage. Focus on German craftsmanship and driving experience rather than status.
2. Profile the "true" Porsche enthusiast, create clear daylight between status-seekers and passion-driven owners.
3. Emphasize racing heritage, shift the narrative from Wall Street to racetrack.
4. Take a self-aware approach, acknowledge the stereotype but cleverly subvert it.
5. Position as timeless versus trendy, frame the 80s association as merely a passing trend.

All seemed like directions worth exploring. I was pretty impressed. And to be fair to our robot friend, this was only after a few moments. If I were doing this again, I'd give it more context and wrestle with it more, and try a few techniques I've done since with it, such as river-jumping.

Afterwards, I told Claude about the actual campaign that worked, from Jon Steel. The turnaround campaign was "Kill bugs fast" created by Goodby, Silverstein & Partners. I asked Claude if it would have gotten to this with more time and prompting.

Claude admitted it couldn't have made that jump: "Honestly, I might not have arrived at 'Kill bugs fast' even with more time. The approach is brilliant precisely because it's unexpected and doesn't follow conventional strategic reasoning. It represents a creative leap that bypasses the entire problem rather than solving it directly."

Again, to be fair, Claude's direction of "Profile the true Porsche enthusiast" is not far off this. Digging into a true enthusiast could have led to something about capturing the sheer joy and feeling you get from driving a Porsche if working with a creative team. Maybe.

That said, "Kill bugs fast" came from someone who knew that feeling. Somebody that understood that Porsche wasn't just about status or engineering heritage. Our robot friends can analyze what customers do. They can spot patterns in behaviour. They can even generate strategies based on perception gaps.

They just can't experience what customers experience. Every good marketer has these moments. When the breakthrough insight comes not from analysis but from experience. From using the product. From feeling the frustration when it doesn't work, from watching customers in the shops get confused. It's more than sentiment analysis or journey mapping. It's often about human understanding.

Maybe it's the difference between knowing about something and knowing something. My take on this is that AI might get us to something half-decent. But still needs a human behind the wheel.

For now anyway.

# 09

## Critical thinking has never been more critical

Ask AI for marketing help and what we get back feels polished, looks professional, and we nod along. We get recommendations on what to do, based on winning examples.

This has always been a challenge for marketers. We love case studies. I love case studies. But case studies are not proof of rules or marketing laws. They are written by winners, about winners. They leave out the late-night panics, the bad bets, the abandoned work. What we get is a tidy story.

AI makes this worse. Because everything looks plausible and appears credible. Until you dig a little.

But you got to dig. Critical thinking was a top behaviour I cared about when I wrote my first book. When I was leading The National Lottery marketing, we stuck up our behaviours on the walls in our office. Number 3 on this list was critical thinking: question assumptions, look for evidence, stay sceptical but not cynical. It mattered then because lazy thinking led to poor decisions.

Now that AI can produce endless confident nonsense, critical thinking has never been more critical. Critical thinking is about asking why something worked, not just accepting that it did. Spotting when you're looking at survivor bias dressed as strategy. And remembering that a pile of isolated marketing facts is trivia, not guidance.

But AI can also help in the pushback. Have it generate counter-arguments. Have it test your assumptions. But don't let it validate you. That's when we get lazy. The job isn't to collect answers. The job is to work out which answers don't stand up.

For me a few checks matter. Start by asking if it's even the right question. Try more than one model, let them clash. Look for the failures, not just the shiny winners. And always ask: what would change my mind? If the answer is nothing, then I'm defending, not thinking.

AI is brilliant at pattern recognition. It's terrible at pattern relevance. The edge isn't in the confidence of the answer. It's in having the judgment to call bullshit when everyone else is nodding along.

If you're going to use these tools, set them up to argue back. AI can absolutely amplify lazy thinking. Confident nonsense at scale. But it can also be the thing that sharpens us, if we set it up that way. Use it to test assumptions, generate counter-arguments, and find our gaps. It won't think for us. But it can force us to think.

# 10

## AI has turned us all into little Hemingways

I've been in awkward meetings where the authors of marketing documents sat down having not yet read their own work.

AI has turned us all into little Hemingways. Smooth lines, confident tone, not a typo in sight. Everything looks like it was written by someone who drinks long blacks and has opinions about linen. Except half the time the writing is saying absolutely nothing. It's like we gave a robot a thesaurus and it decided to run marketing.

You can see where this is heading. Marketers won't just use AI to help with a plan, they'll get it to actually do the plan. Their managers will get their own AI to review it. And the two systems will start emailing each other like a pair of polite civil servants in Merrion Square, discussing "alignment" while the humans are staring at the screen wondering why Outlook suddenly feels smarter than them. Whole projects will sail from kick-off to sign-off with nobody meaningfully involved.

Give it a bit longer and you'll have robots making work for other robots, with humans floating around the edge doing that nervous nodding we all do in meetings when we haven't a clue what's happening. AI will write the strategy. Another AI will summarise it. A third AI will send a meeting invite nobody attends, because overnight some automated workflow quietly approves the whole thing.

That's the risk. Not that AI will make the writing worse. It won't. But it'll make the nothingness sound good.

AI is actually brilliant at forcing this honesty. Ask it to explain the strategy to a non-marketer. In plain language. No fancy words. No poetry. Ask it to rewrite your paragraph without the buzzwords, and now you're

looking at something my 9-year old twins would understand. That's the test. Removes all the places to hide. Once the sentence is plain, you have to stand over it.

Our robot friends give us polish for free. It's kinda cool in one sense. But it also makes hiding easier than ever. The advantage now is in the marketer who can look at the polished thing and say, "Right, but what does any of this actually mean?" The one who can ask the daft questions without embarrassment. The one who uses AI not as a ghostwriter, but as a truth serum.

Because if everything sounds smart, the thing that matters is what's actually true.

# 11

## The friction is the point

AI can show you things you'd never think of on your own. That's handy. But there's potentially a longer-term downside.

You ask a question, and the robots suggest an angle you hadn't thought of, maybe a half-clever phrase you wish you'd written. Feels good. Feels like progress. And in some sense, it is progress. Except your brain doesn't know whether that insight came from work or luck. The buzz is the same, but when we work for it, it stays in our brains.

John Nosta calls this the "*Cognitive Corridor*". He says it's like driving at night with dipped headlights, then someone flicks the full beams for a second and you see the landscape ahead. That light helps you notice where you might go next. But it's not understanding. It's just illumination.

The friction is the point. Learning science backs this up. The struggle is where the wiring happens. The friction isn't a tax on learning. It is the learning.

Say you're learning to cook. You could ask AI for the perfect risotto recipe, follow it step by step, and end up with something smooth enough to post online. But you wouldn't know why it works, or what to do when it doesn't. The person who's scraped burned rice from the pan has something better than a recipe. They've got judgment. And a new pan.

Psychologists call this the generation effect. You remember more when you make the thing yourself. Every time you reach for an idea, you strengthen the path to it. Skip that bit and you've eaten without digesting.

We're doing this in marketing too. AI can give us campaign ideas, strategy frameworks, neat charts. But effectiveness lives in *why* things work, not *what* might work.

The strategist who's had to watch their ideas fail, who's tested and learned and changed tack, has built a kind of scar tissue. That's judgment. Someone running AI playbooks is borrowing someone else's.

The Cognitive Corridor isn't the destination. It's the hallway that shows you which doors might be worth opening. Let it spark ideas, ask odd questions, point you somewhere new. Then maybe close the chat...and think.

I discussed this with Claude. It agreed (naturally). It told me that the answer is almost a side effect. The learning is the goal. Wish I'd written that.

# 12

## A very tidy wrong answer

Richard Schwartzstein teaches at Harvard Medical School and he likes showing students how easy it is to reach for the wrong explanation. He took a 45-year-old man with obesity and shortness of breath and fed the case into ChatGPT. It came back confident: congestive heart failure.

Safe answer. Plausible. Wrong. He was obese and short of breath. Complications of obesity can create the same signs without heart failure.

That's the point, the label can be plausible and still be the wrong cause. ChatGPT saw a pattern and matched it to the nearest familiar answer. Obese, short of breath, it reached for heart failure because that's a common story in that neighbourhood. Humans do it too, by the way. We just do it with less confidence and worse formatting.

Mechanistic reasoning is different. It asks what is physically happening here, in this specific body, right now. What is the chain that links the symptom to the cause? In Schwartzstein's example, the chain is obesity changing how the body moves air and shifts fluid around, not a failing heart pump. Same outward signs, different reason.

I'm not having a go at AI. It did what it's built to do. It produced an answer that sounds like the kind of thing that belongs in a hospital note. The danger is that sounding right gets mistaken for being right, and when the work is serious, that mistake costs.

That's where the human value sits when you work with these tools. Not in getting them to write more, or faster, or smoother. In judgment. In being able to look at a confident answer and ask, "What would have to be true for this to be right?" and, "What else could explain the same signs?"

That little pause matters now because the tools are getting better at removing it for you. They give you a clean story and a neat next step, and your brain relaxes.

Marketing has the same trap. Feed AI a brief and you get something that looks like strategy. Structurally sound. Pattern-matched from thousands of campaigns. But it can't tell you whether this strategy will work for this brand, with this history, against these competitors, given what happened last quarter. That's mechanistic reasoning. That's judgment.

It's the difference between "our awareness is down so we need a bigger top-of-funnel push" and "our page got slower, the checkout got fiddlier, and we're bleeding people before the message even has a chance." Same symptom, different cause, different fix.

There's another problem sitting underneath all of this; you can't judge the output if you don't know what good looks like.

Alex Yuen at Harvard's design school waits until halfway through the semester before bringing AI into student workflow. A trick I'm going to nick for my classes. His reasoning is simple: if you haven't done the work yourself, you can't judge whether AI output is any good.

Skip the foundation and everything AI produces looks plausible. You lose the ability to tell the difference between a strong answer and a well-written one. This is why "humans in the loop" can turn into a comforting lie. The loop is only useful if the human can tell the difference between a good answer and a good sounding answer. Otherwise you're just nodding at the screen, grateful it saved you time, while it quietly picks the story for you.

Judgment is built from behaviours. Putting in the work. Knowing what good is. Poking past the first answer. Building real things so you see where reality pushes back. Staying hopeful enough to use the tool, and sceptical enough not to believe it.

Part 3 is about these behaviours.

[runwithfoxes.com](http://runwithfoxes.com)