twiraa

VOLUME 1

THE FOUNDER'S BRIEF

The Intelligence Paradox

WHY AI WON'T REPLACE GREAT LEADERS, BUT IT WILL EXPOSE AVERAGE ONES.

READ • REFLECT • DISCUSS

SAURABH KHANNA
FOUNDER, TWIRAA CREATIVES PVT. LTD.

MASTHEAD

About This Publication

Most business writing tries to hand you an answer. This one is built to hand you a better question — the kind that changes what happens in the room after you have closed it.

The Founder's Brief is a quarterly publication built around a single conviction: the leaders who win the next decade will not be the ones with the most AI tools. They will be the ones who got sharpest about what AI still cannot do for them. Each edition takes one shift already reshaping business and unpacks what it means for the people responsible for steering through it.

This is not a forecasting exercise and it does not pretend to universal answers. It is written to be read in fifteen minutes — on a flight, between meetings — and argued about for considerably longer than that, in a leadership offsite, a board pack, or a founder's own notes before a decision that actually matters.

Where the arguments here draw on published research, the source is credited. Everything else is first-hand observation, drawn from fifteen years of sitting across the table from founders, CXOs and boards wrestling with exactly these questions.

If one idea in these pages changes one conversation inside your organisation this quarter, this issue has done its job.

IN ONE SENTENCE *The Founder's Brief exists to help leadership teams ask sharper questions about the shifts already underway in their business.*

FROM THE FOUNDER

A Letter to the Reader

Dear Reader,

Over the past few months, one question has followed me into almost every conversation — with founders scaling their first real teams, CEOs running mature businesses, and marketing leaders trying to make sense of a tool landscape that reinvents itself every few weeks.

> *“How should we use AI?”*

It is a fair question, and an inevitable one. We are living through one of the fastest technology shifts in modern business history. New tools appear weekly. Tasks that took hours take minutes. Entire workflows are being quietly redesigned in the background of ordinary workdays.

The pressure to do something with AI has become universal. Every board wants the strategy slide. Every team wants a pilot to point to at the next town hall.

But the more of these conversations I sat in, the more convinced I became that we are starting from the wrong end. The more useful place to begin is not *how should we use AI*. It is:

> *“What should remain deeply human?”*

Because while technology keeps changing *how* organisations operate, it has not changed *why* people choose to trust them. Customers buy confidence, not just capability. Employees follow leaders they believe in, not the ones with the best prompt library. Investors back conviction over polish. Partners still value integrity over speed.

AI will keep reshaping the future of work. But leadership was never a race to complete more tasks. It has always been the discipline of making better judgments — particularly the ones no dataset can make for you.

That is where the real opportunity sits for the leaders reading this. Not in adopting AI faster than the competition, but in getting sharper about what only they can still do.

This first issue is deliberately uncomfortable in places. I hope it starts a useful argument inside your organisation.

Warm regards,

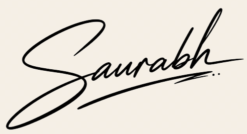

Saurabh Khanna

Founder, Twiraa Creatives Pvt. Ltd.

EXECUTIVE SNAPSHOT

The Argument, in Three Points

01

Execution variance has collapsed.

AI has not made anyone smarter. The repetitive, research-heavy, first-draft work that used to fill a leader's calendar is exactly what it absorbs. It cannot absorb the judgment underneath it.

02

Output stopped carrying information.

For twenty years, the quality of what someone produced was a usable proxy for the quality of their thinking. That proxy has been withdrawn — quietly, from every organisation at once.

03

The decision itself is now exposed.

Most organisations have never looked at it directly. The businesses that pull ahead will be the ones that got honest earliest about where their judgment actually sits.

ALSO EXAMINED IN THIS ISSUE

- The four kinds of leadership authority AI has just stripped the cover from
- Why sorting work by what AI can do expires within a year — and what to sort it by instead
- The apprenticeship problem: we automated the work that made people good, and kept the standard that assumed they still were
- The question no board has answered in writing: who is accountable when an AI-informed decision fails?

THE OPENING

Three Decks, All Excellent

A procurement committee I sat with earlier this year opened three vendor submissions for the same mandate.

All three were excellent. The formatting was flawless. The market sizing was credible. The competitive analysis was thorough, current and well argued. Six years ago, two of those three would have been visibly weaker than the third, and the committee would have known what to do within twenty minutes.

This time they sat in silence.

The screening criterion that had worked for a decade — *who did the better work?* — had stopped functioning. Not because the work got worse, but because it all got good. And when everything is polished, polish stops carrying information.

One of them eventually said the quiet part out loud: "I don't know how to tell these apart any more."

> *That sentence is the most consequential thing happening in business right now, and almost nobody is putting it on a board agenda.*

It does not stop at vendor selection. The same collapse is running through hiring, internal reporting, agency reviews, investment memos and performance conversations. Everywhere an organisation used the quality of output to infer the quality of thinking, the inference has quietly stopped working.

What is left, once polish stops carrying information, is the decision underneath it. **And most organisations have never had to look at that directly.**

IN ONE SENTENCE *When execution quality becomes universal, it stops being evidence — and the thinking it used to conceal becomes visible for the first time.*

THE MAIN ESSAY

The Scribes and the Scholars

A Familiar Panic, Selectively Remembered

The usual comfort offered at this point is a list. The printing press was expected to undermine scholarship. The calculator was believed to weaken mathematical thinking. The internet was predicted to destroy traditional business. Cloud computing was once considered too risky for the enterprise. Each fear, we are told, proved overblown.

It is a reassuring list, and it is selectively remembered.

It quietly omits the switchboard operator, the typesetter, the film processor and the travel agent — none of whom were being irrational. Their work did not get redistributed. It got deleted.

So both things are true at once, and this is the part most commentary skips. The technology did not destroy the profession. It destroyed **the part of the profession that had become mechanical** — and everyone whose value was concentrated there went with it. The scholars survived the printing press. The scribes did not.

Which reframes the question entirely. It is not whether AI will replace leaders. It is: *how much of your leadership is currently concentrated in the mechanical part?*

The Paradox Itself

AI is exceptionally good at producing answers. Leadership has never been about having more answers. It has always been about asking better questions, and knowing which answer to trust.

So as intelligence becomes cheap and abundant, the thing it was supposed to replace becomes scarce and expensive.

> *As intelligence becomes increasingly available, judgment becomes increasingly valuable. That is the paradox.*

There is a harder edge to this than the phrase suggests. Scarcity does not just raise the price of judgment — it makes the absence of judgment conspicuous. For most of the last two decades, an organisation could operate with mediocre decision-making and excellent execution, and the execution would carry it. That arrangement is over.

IN ONE SENTENCE *The more accessible answers become, the more an organisation's advantage depends on the judgment applied to them — and the more visible it is when that judgment is thin.*

THE MECHANISM

The Exposure Effect

Four kinds of leadership authority were protected by execution variance. All four are now exposed. None of them belong to bad leaders — which is precisely what makes this uncomfortable rather than satisfying.

01

The Volume Manager

Their authority came from throughput — the one who turned the deck around overnight, whose value proposition was capacity. Capacity is now a subscription available at forty dollars a seat.

THE TELL

Their contribution to a meeting is a status update rather than a position.

02

The Information Gatekeeper

Their power came from asymmetry — knowing the market, the benchmark, the competitor's pricing first. That asymmetry now has a half-life of about four seconds.

THE TELL

Visible discomfort when a junior arrives holding the same data they planned to present.

03

The Process Proxy

Some organisations mistook motion for progress. AI multiplies motion enormously — so the gap between activity and outcome, once deniable, has become embarrassing.

THE TELL

Output volume is sharply up this year and nobody can name a commercial number that moved with it.

04

The Consensus Chair

The most senior, most exposed of the four. Convening the people who had the analysis was the contribution — but when everyone has the analysis, having a view is the entire job.

THE TELL

They ask for more analysis when what the situation needs is a decision.

Most of the people described above are experienced, diligent, and were genuinely excellent at a game whose rules changed underneath them without notice. That is why the honest version of this conversation belongs inside your own leadership team — not in a keynote about somebody else's company.

IN ONE SENTENCE *AI did not lower anyone's capability. It removed the cover that made average judgment look like good judgment.*

THE FRAMEWORK

The Reversibility Grid

Most organisations sort work by asking *can AI do this?* Within eighteen months the answer will be yes almost everywhere, which means any framework built on that question has a shelf life of about seven months.

The durable question is different: **what does it cost us if this is wrong, and can we take it back?** Two axes. Cost of being wrong. Reversibility of the decision. Four quadrants, four different operating rules.

	LOW COST	HIGH COST
REVERSIBLE	**AUTOMATE** *AI owns it outright. No human review — reviewing costs more than the errors do.* Meeting notes, first drafts, formatting, scheduling, routine reconciliation.	**CONTEST** *AI proposes, a human is assigned to attack it. Structured disagreement, not review.* Pricing changes, campaign strategy, resource reallocation, forecast models.
IRREVERSIBLE	**VERIFY** *AI drafts. A named human signs. The signature is the control, not the reading.* Anything customer-facing, published, filed or sent externally.	**OWN** *A human decides. AI may inform, never author. Reasoning is written down first.* Hiring and exits, M&A, capital structure, market entry or exit.

Rows read by whether the decision can be undone. Columns read by what it costs if you're wrong.

Two things this grid does that a task list cannot. **It survives model improvement** — when the next release does something your "human column" claimed was impossible, nothing needs rewriting; a task may move quadrants, but the rules stay true. **It forces the accountability conversation** — every quadrant except AUTOMATE requires a named person, and most leadership teams discover they've been operating in CONTEST and OWN with nobody named at all.

IN ONE SENTENCE *Sort your work by what it costs you to be wrong — not by what the technology can currently do.*

THE FRAMEWORK

The Intelligence Spectrum

Organisations talk about “using AI” as though it were one decision made once, in one meeting. In practice it is four separate decisions, each demanding a different level of human involvement — and a fifth thing that is not a level at all.

Human leadership is not the top rung. It is the spine that runs through all four — the thing that decides which function belongs at which level, and stays accountable when the answer turns out to be wrong. Drawing it as a rung implies you graduate to it. You do not. You either exercise it throughout, or you exercise it nowhere.

01	02	03	04
Automation	**Augmentation**	**Collaboration**	**Decision Intelligence**
Rule-based work runs without people.	Individuals produce more, alone.	Teams and models work a problem together.	AI informs; accountability stays named and human.
FAILURE MODE	**FAILURE MODE**	**FAILURE MODE**	**FAILURE MODE**
Automating a broken process, faster.	*Output rises, quality variance rises with it.*	*The model becomes the loudest voice in the room.*	*“The model suggested it.”*

THE SPINE
Human Leadership

Deciding which function sits at which level, and owning the consequence — written down and revisited on a schedule, or made by default.

Most organisations deliberately chose levels one and two, then drifted into three and four without anyone signing off. **Drift is not a strategy. It is an unpriced risk.**

IN ONE SENTENCE *The question is not “are we using AI” but “which function sits at which level, who decided that, and when do we revisit it?”*

THE BLIND SPOT

The Apprenticeship Problem

There is an irony sitting inside every AI deployment plan, and it is almost never said out loud.

> *AI is best at precisely the work we used to hand to junior people in order to make them good.*

The research synthesis. The first draft. The model rebuilt four times. The competitive teardown nobody read. That work was never valuable in itself — it was the gymnasium. It was how a twenty-four-year-old accumulated the ten thousand small pattern recognitions that eventually became judgment.

We have automated the gymnasium and kept the exam.

Consider what a leadership team looks like in 2038. Its members entered the workforce in 2026. They have spent twelve years reviewing, editing and approving machine output, and almost no time producing anything from a blank page. They will have excellent taste and thin reps. They will be able to tell you that a strategy is wrong without being able to tell you what the right one is.

Every organisation currently optimising junior headcount is trading five years of margin for fifteen years of capability. Some will decide that trade is correct. Very few are making it deliberately, and almost none have priced it.

THREE THINGS TO DO BEFORE THE NEXT BUDGET CYCLE

1. **Design the reps back in, on purpose.** Pick two or three problems a quarter where juniors work without AI — the way a pilot flies manual approaches in a simulator. Name it as training, budget the extra time, and never apologise for it.
2. **Review reasoning, not output.** When a junior submits work you cannot fault, ask them to walk you through what they rejected and why. It is the only quality signal that still works.
3. **Make one senior person accountable for bench depth.** Not for hiring numbers — for whether the people three levels down are becoming capable of replacing the people two levels up.

IN ONE SENTENCE *We automated the work that made people good, and kept the standard that assumed they still were.*

GOVERNANCE

The Accountability Gap

A question worth putting to your leadership team before the pilot quietly becomes infrastructure:

> *When an AI-informed decision goes badly wrong in this company — who, by name, is accountable, and where is that written down?*

In most organisations the honest answer is: nobody, and nowhere. Which is workable right up until the day it is not.

THREE CONSEQUENCES THAT ARRIVE SOONER THAN MOST BOARDS EXPECT

— **Insurance and audit.** Underwriters and auditors are beginning to ask about AI governance in the places they asked about cyber three years ago. “We encourage responsible use” is not a control. A named signatory and a written reservation list is.

— **Regulatory asymmetry.** If you operate across jurisdictions, the same workflow can be compliant in one market and reportable in another. Most mid-market companies discover this after deployment rather than before.

— **Reputational half-life.** A machine-generated error that reaches a customer costs materially more than the identical error made by a person — because the story is not the error. The story is that nobody was watching.

The remedy is not complicated, which is exactly why its absence is hard to defend. **One page. The decisions you have reserved from AI. The person accountable for each. The date you will revisit it.** If your organisation cannot produce that page today, that is the finding.

IN ONE SENTENCE *Governance is not a brake on adoption. It is the only thing that makes fast adoption survivable.*

THE WORKING SESSION

Five Questions, and Five Weak Answers

Take these into your next leadership meeting exactly as written. The second column is the more useful one — it tells you what you are listening for.

01	**Where are we using AI only because everyone else is?**	*"We're running several pilots."* A pilot is a hedge. Ask which one has a P&L owner and a kill date.
02	**Which decisions have we reserved from AI — in writing?**	*"We use judgment on the important ones."* If it is not written down it is not reserved. It is untested.
03	**Are we becoming faster, or genuinely better?**	Any answer that cites output volume. Speed with no quality gain is noise on a shorter deadline.
04	**How will a customer feel the difference?**	*"It makes us more efficient."* Efficiency is an internal benefit. If the customer cannot feel it, it is a cost line, not an advantage.
05	**If every competitor has the same tools, where does our advantage come from?**	A tool name. Or a long pause. The pause is the more useful of the two — it is the finding.

Ask question five, then say nothing for a full minute. Who breaks the silence, and what they reach for, will tell you more than any AI roadmap on the wall.

FROM READING TO DOING

Five Actions You Can Take This Month

1

Audit the work before buying another tool

Most organisations do not have a tooling gap. They have a clarity gap about what is actually worth automating, and buying software is the faster way to avoid finding that out.

2

Write the reserved list

Five to eight decisions this company will not delegate to a model, whatever it becomes capable of. Fewer than five means you have not been honest. More than fifteen means you are protecting comfort rather than risk.

3

Protect thinking time deliberately

Strategic thought needs unstructured space, and that space is the first casualty of a calendar optimised purely for output.

4

Change what you review

Move one recurring review from output quality to reasoning quality. Ask what was rejected and why. Do it for a quarter and you will learn more about your bench than any appraisal cycle has told you.

5

Name the signatures

For everything in CONTEST and OWN, write a person's name. Not a team. Not a committee. Note every box you cannot fill — that list is your actual risk register.

FOUNDER FORECAST

Four Calls, and a Date to Grade Them

A prediction that cannot be proven wrong is not a prediction. These are dated and specific. They will be graded in public — including the ones that miss.

2028 BY JULY	“AI-enabled” disappears from vendor pitch decks as a stated differentiator — the way “cloud-based” vanished between 2012 and 2016. Claiming it will start to read as dated rather than advanced.	HIGH
2029 BY YEAR END	At least one large listed company discloses an AI-attributable decision failure as a material event. The governance question that follows — who signed off — drives more change across boards than the loss itself.	MEDIUM
2030 BY YEAR END	Entry-level hiring in professional services sits visibly below 2025 levels while revenue per firm rises. The firms that quietly kept hiring and training juniors will not look clever until around 2035, and then will look very clever indeed.	HIGHEST STAKES
— ONGOING	The premium for demonstrably human work — live, unedited, attributable, accountable — rises. First in categories where trust is the product, then everywhere else.	MEDIUM

These get scored in Volume 5, in July 2027. A forecast nobody grades is decoration.

TEAR-OUT

The Offsite Card

Ninety Minutes. Six People. One Whiteboard.

TEAR ALONG THIS LINE TO KEEP AT YOUR DESK

0–10 **The opener**

Everyone writes down, privately, one decision they made this quarter whose reasoning they could not fully explain. Nobody shares. This is calibration, not confession.

10–30 **Map the grid**

Take your twelve most consequential recurring work streams. Place each in AUTOMATE, VERIFY, CONTEST or OWN. Argue about the placements — the argument is the exercise.

30–50 **Name the signatures**

For everything in CONTEST and OWN, write a name. Not a team. A person. Note every box you cannot fill.

50–70 **The reserved list**

Agree the decisions this company will not delegate to a model, whatever it can do. Five to eight.

70–85 **The bench question**

Who, three levels down, is becoming capable of replacing someone two levels up? If the room cannot answer inside five minutes, that is the most important output of the session.

85–90 **One owner, one date**

All of the above becomes a single page with a review date ninety days out. If it does not, this was a pleasant conversation and nothing more.

FROM MY DESK

A Closing Thought

Leadership has never been a race to know more. It has always been a commitment to decide better.

That distinction is about to stop being a nice sentiment and start being a sorting mechanism. Organisations that treat AI as a shortcut to more answers will find themselves competing on a commodity everyone else has too. The ones that treat it as a way to free up better judgment will find themselves competing on something considerably harder to copy.

The uncomfortable part is that you will not get to choose which group you are in. Your customers, your best people and your competitors will work it out for you, and faster than they used to.

> *There is nowhere left to hide. Which, if your judgment is sound, is the best news you have had in years.*

ABOUT THE AUTHOR

Saurabh Khanna

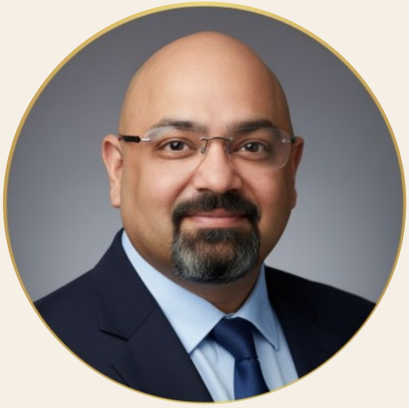

Founder, Twiraa Creatives Pvt. Ltd.

Saurabh Khanna is the Founder of Twiraa Creatives, a brand-led demand generation firm built on a simple premise: that the businesses winning the next decade will be the ones that fuse machine intelligence with human emotional depth, rather than choosing between them.

He has spent over fifteen years in brand strategy, go-to-market and performance marketing — working with large consumer and technology brands, high-growth enterprises and founder-led businesses across India and international markets. He now works closely with founders, CXOs and leadership teams on the part of growth that does not fit on a dashboard: positioning, judgment, and the decisions that determine whether the rest of the plan is worth executing.

The Founder's Brief is his independent publication, written for people who would rather leave with a better question than another tool.