INTENTIONAL EDGE

The Quiet Reckoning

What Organisations Owe Their People in the Age of AI

A white paper by Intentional Edge
March 2026

Foreword

Before the data, before the framework, before the argument — a word about where this paper comes from.

Two years ago, I closed a business. I made myself and my team redundant. I sat across from people I had recruited, developed, and genuinely cared about, and told them that the organisation we had built together no longer had a future in its current form. I led that process as honestly as I could — which meant holding their uncertainty and my own at the same time, maintaining the composure the situation required while privately experiencing the same loss.

I know what it costs to tell people the truth about what is coming. I also know that it was the right thing to do, and that it was possible to do it with dignity — for them and for myself.

What I fear is happening in boardrooms and on earnings calls right now is something different. An inclination — not yet named, not yet confronted — towards strategies already in motion, decisions already made. Leaders who know, privately, the scale of what is coming. And a workforce — including the children and younger siblings of those same leaders — that is largely not being told.

That gap is what this paper is about. Not the technology. Not the economics. The gap between what is known and what is said, and what that gap costs the people on the wrong side of it.

This is written from the vantage point of someone who has been on both sides of that conversation. My hope is not that it is useful. My hope is that it is thought-provoking — that it creates a moment of honest introspection, and that the introspection leads somewhere.

Nicole, Founder, Intentional Edge

Introduction

Most organisations deploying AI right now are not getting the returns they were promised. An MIT study found that 95 percent of enterprise AI pilots are failing to deliver measurable financial returns. IBM's research puts the share that deliver expected ROI at 25 percent. The gap is almost never the technology. It is almost always the organisational conditions surrounding it — the governance, the leadership capacity, the decision rights, the human judgment required to translate an AI recommendation into something the organisation actually does differently. Those conditions do not appear on a technology roadmap. They are built, deliberately, or they are absent. And their absence is expensive.

Something else significant is happening in boardrooms, on earnings calls, and at the world's most influential forums. The leaders of the world's most powerful companies are speaking openly — some with apparent candour, others with barely concealed urgency — about the scale of disruption that artificial intelligence will bring to the workforce. They know what is coming. Many are already acting on it.

And yet the people who will bear the cost of that disruption — the administrative workers, the junior analysts, the early-career professionals, the women disproportionately concentrated in the most exposed roles, the middle managers whose experience and judgment hold organisations together — are largely not part of those conversations. They are the subject of the strategy, not participants in it.

This paper holds both questions at once. It is a paper about performance — specifically, why the organisational conditions that determine whether AI investments deliver returns are the same conditions that determine whether the people inside the organisation can actually function during a transition of this scale. And it is a paper about responsibility — what organisations owe the people within them at a moment when both are under pressure that most leaders have acknowledged privately and very few have confronted publicly.

> ***The gap between what business leaders know and what they are doing about it is not a knowledge problem. It is a courage problem.***

That gap — between stated concern and actual behaviour — is the subject of this paper. It is also, we argue, the defining organisational challenge of this decade.

The System Behind the Disruption

Most conversations about AI and the workforce treat them as a bilateral relationship: a technology acting on an organisation, and an organisation deciding how to respond. That framing is too narrow — and the decisions made within it will be too narrow as a result.

What is actually under way is the collision of forces operating at four distinct levels simultaneously. At the geopolitical and macroeconomic level, the most aggressive monetary tightening cycle in four decades — driven by the US Federal Reserve from March 2022, and followed by central banks across Europe, Asia, and beyond — raised the cost of capital globally. Dollar strengthening increased borrowing costs for lower-income, dollar-exposed economies far beyond the US. Inflation compressed consumer spending. Investor patience shortened. The financial conditions in which organisations were operating changed fundamentally, and changed fast.

At the corporate level, that pressure met a technology — generative AI — that had been developing for decades but crossed an accessibility threshold in late 2022. For the first time, AI capable of performing cognitive, language-based work required no specialist resource to deploy. Organisations under simultaneous pressure from expensive capital, rising labour costs, and shareholder demands for efficiency found in AI not just a productivity tool but a cost recoupment mechanism. The pace of adoption was not driven purely by strategic intent. It was driven by financial pressure seeking the fastest available lever.

At the labour market level, those corporate decisions are landing on workforces already differentiated by structural vulnerability. Women overrepresented in the roles being automated. Young people entering a market whose entry pipeline is quietly contracting. Middle managers asked to lead a transition that may eliminate their own position. Workers in communities where administrative employment has been the primary source of stable, reasonably paid work for a generation. The costs of decisions made under macroeconomic pressure are not distributed evenly. They concentrate.

And at the individual level — the level that most strategic conversations never reach — there are people experiencing the uncertainty of not knowing whether their role will exist in two years, whether the skills they have built are still valued, whether the organisation they work for is telling them the truth about what is coming. Their engagement, their willingness to adopt new tools honestly rather than perform compliance, and their judgment — which the evidence consistently shows is the primary determinant of whether AI investments deliver returns — all depend on whether that uncertainty is met with honesty or managed with messaging.

The performance consequence of getting this wrong is direct. AI does not solve the organisational complexity of scaling — it inherits it. Deployed into an organisation where decision rights are unclear, where context does not travel between functions, and where leaders are making decisions without the full system in view, AI accelerates the confusion rather than the output. The same governance gaps that cause decisions to revert upward and momentum to stall as teams scale do not disappear when AI is introduced. They are amplified by it. The organisations that discover this tend to discover it at the worst possible moment: when a deployment has stalled, a team has disengaged, or an investment hasn't delivered and the board wants to know why.

> ***The forces driving AI-related workforce disruption are not organisational in origin. They are geopolitical and macroeconomic. But the decisions about pace, transparency, and who bears the cost of transition are entirely within organisational control.***

This matters because the question it poses to senior leaders is more demanding than the one most AI adoption frameworks ask. It is not: how do we manage our people through this change? It is: are we making decisions at the organisational level without accounting for the system we are operating inside — and what are the consequences of that, for our people, and for the organisation's own capacity to perform?

The organisations that treat AI adoption as a bilateral relationship between technology and strategy will optimise at the firm level in ways that degrade the system conditions their own future performance depends on. The entry pipeline they close is the one that produces the mid-level talent they will need in five years. The trust they erode through managed messaging is the same trust that honest AI feedback — the kind that tells you what is actually working — depends on. The experienced judgment they eliminate to recoup an

investment that has not yet delivered is precisely the capability the evidence shows is required to make it deliver.

This is not a moral argument, though the moral case is real. It is a compounding returns argument. The organisations that make decisions with the full system in view do not trade performance for responsibility. At sufficient time horizon, the two are the same thing.

The rest of this paper examines what is happening at each level, who bears the cost, and what it would actually mean to lead consciously within this system rather than react to it. The Five Conditions in Section 6 are not a change management checklist. They are the points at which a systems-aware leader intervenes differently — and where the gap between organisations that will navigate this well and those that will not is already opening.

The People in the Room

There is one thing that is rarely said directly in this conversation, and it should be.

The leaders making AI adoption decisions today are not making them about an abstract future workforce. They are making them about people they know. The class of 2026 — the graduates entering a labour market that is quietly contracting around them — are the children, the younger siblings, the nieces and nephews, the godchildren of the people in the room. The middle managers facing redundancy are former colleagues, mentors, friends, people who were at the wedding.

The distance that corporate decision-making creates between a strategy and its human consequences is, in this moment, an illusion. There is no distance. The people who will bear the cost of these decisions are sitting at the same dinner tables as the people making them. They are being handed the bill for choices made in rooms they were not invited into, about futures they were not consulted on.

That does not make the decisions easier. But it does make the question of honesty harder to avoid.

And yet most organisations are avoiding it. Not because their leaders are indifferent — most are not — but because the institutional architecture of corporate decision-making does not create space for this kind of reckoning. Strategy is set in one room. People are managed in another. The connection between the two is rarely made explicit, and almost never made personal.

This paper argues that making it personal is precisely what responsible leadership in this moment requires.

---

1. What Is Actually Happening

The public narrative and the private reality

The standard corporate communication on AI and jobs follows a familiar arc. AI will create more jobs than it destroys. Workers who use AI will thrive. The future belongs to those who

adapt. These messages are not entirely false — but they are carefully incomplete, and the people delivering them know it.

An Axios investigation found that every single CEO they spoke to, across companies of various sizes and industries, is actively working out when and how AI agents can displace human workers at scale. This is not a fringe view. It is the mainstream private consensus among senior leadership. The public message is managed. The private strategy is clear.

The executive statements that have begun to surface publicly are striking in their directness. Ford's CEO has warned that AI will replace literally half of all white-collar workers. Salesforce's Marc Benioff has claimed AI is already handling up to half of his company's workload. Walmart's CEO has stated it will change literally every job. JPMorgan has instructed managers to avoid hiring as AI is deployed across its businesses. Goldman Sachs has described taking a front-to-back view of how it organises people in response to AI — prioritising productivity and efficiency throughout.

Microsoft's AI chief has predicted office jobs will crumble within 18 months. Anthropic's CEO Dario Amodei has estimated that AI could eliminate roughly half of all entry-level white-collar positions within five years — and has stated that those producing this technology have a duty and an obligation to be honest about what is coming.

What makes this particularly striking is not the scale of the predictions but the gap between those predictions and organisational behaviour. Nearly 60 percent of hiring managers surveyed in 2025 said they emphasise AI's role in reducing hiring because it is viewed more favourably than admitting financial constraints. AI is simultaneously a genuine driver of workforce change and, in some cases, *a convenient cover story* for decisions that would otherwise attract greater scrutiny.

The global data underneath the narrative

The scale of what is under way is not a US story. It is a global one — and the data, when read without the reassuring headline aggregate figures, is more concerning than it is usually presented.

The World Economic Forum, drawing on surveys of over 1,000 employers representing 14 million workers across 55 economies, projects 92 million jobs displaced globally by 2030, with 170 million new roles created — a net gain of 78 million. That net figure is repeatedly cited as reassurance. It deserves scrutiny. Net gains tell us nothing about which people will occupy the new roles, or whether they will be the same people who lost the old ones, or whether the transition between the two will be manageable within a human career. Aggregate optimism is cold comfort to the individual experiencing the disruption.

The entry-level picture is already visible and already global. Since January 2024, entry-level job postings have fallen 29 percent based on analysis of 126 million postings worldwide. Youth unemployment stands at 17 percent in India, 16.5 percent in China, 15.3 percent in the UK, and 57 percent in South Africa. South Africa2019s figure warrants a specific note: its youth unemployment crisis is structural and predates AI by decades, rooted in the long economic legacy of apartheid, persistent skills mismatches, and a labour market that has never absorbed its young population at scale. AI did not create that crisis. What AI does 2014 in South Africa as elsewhere 2014 is close the entry-level roles that have historically served as the most accessible first rung into formal employment, for precisely the workers least able to absorb that loss. AI adoption does not cause the same harm everywhere. It lands on existing vulnerabilities and amplifies them. In the UK, 1.2 million graduates competed for just under 17,000 entry-level positions in 2024. These are not future projections. They are present realities in economies at very different stages of AI adoption

2014 which is why the cause cannot be read as simply cyclical, and why the response cannot be simply technological.

McKinsey estimates that today's existing technology — not future iterations — could theoretically automate approximately 57 percent of current work hours globally. Goldman Sachs estimates that 25 percent of global work hours could be automated with current AI capabilities. The ceiling is high. The pace of approach to that ceiling is accelerating.

> ***'If AI does to white-collar work what globalisation did to blue-collar work, we need to confront that directly — not with abstractions about the jobs of tomorrow, but with a credible plan for broad participation in the gains.' — Larry Fink, CEO BlackRock, Davos 2026***

The perception gap: what leaders think is happening versus what people experience

Running beneath the public narrative gap is a second, equally important one — between what leaders believe about how their people experience change, and what those people actually feel.

Boston Consulting Group and Columbia Business School research from November 2025 found that 76 percent of executives believe their employees feel enthusiastic about AI adoption in their organisations. The reality? Only 31 percent of individual contributors express such enthusiasm. Leaders are more than twice as far off the mark as they believe.

This is not a new problem specific to AI. Gartner research consistently shows that around 74 percent of executives believe they include people meaningfully in change processes, while only 42 percent of employees feel that way. McKinsey has found that while 80 percent of senior executives believe their change management initiatives are successful, only 30 percent of frontline employees agree. The pattern is remarkably consistent: leaders systematically overestimate how well they are communicating and how supported their people feel.

William Bridges' foundational work on transition management offers a useful lens here. Bridges drew a sharp distinction between change — the external event, the new system, the restructured team — and transition, the internal psychological process people must go through to genuinely adopt new behaviours. Most organisations manage the change. Almost none manage the transition. The result is what Bridges described as the neutral zone — a state of suspension in which people have let go of the old way but have not yet committed to the new one — where resistance hardens, anxiety rises, and the visible metrics show green while the actual behaviour on the ground remains unchanged.

Bain's research on organisational change found that most change programmes direct roughly 90 percent of their energy into external mandates — policies, training programmes, governance decks, process maps — while only 10 percent goes into the internal reinforcement that actually changes behaviour: real feedback loops, consequence management, visible leadership modelling, environment design. The organisations that invert that ratio are four times more effective. The implications for AI adoption are direct: most organisations are spending almost all of their effort on the thing that matters least.

Applied to AI adoption, the question this raises is simple: what would it actually mean to bring your people into the reality of this moment honestly? Not a town hall with prepared answers. Not a FAQ document. A genuine conversation — at the organisational level —

about what AI adoption means for roles, for career paths, for livelihoods, and what the organisation is prepared to commit to in response.

---

2. Who Bears the Cost

The uneven distribution of disruption

One of the most important — and least discussed — features of AI-driven workforce disruption is that its costs are not distributed evenly. The aggregate headline figures mask a pattern of concentration that, left unaddressed, will structurally deepen existing inequalities across gender, age, class, and geography. Understanding that concentration is not an optional add-on to the AI strategy conversation. It is the conversation.

Gender: a double exclusion

Globally, 4.7 percent of female jobs fall into the highest-risk automation category, compared with 2.4 percent for males. In high-income countries — where financial and professional services are most concentrated — that disparity sharpens: 9.6 percent of women's jobs carry the highest automation risk, versus 3.2 percent for men.

The roles most immediately exposed — administrative support, customer service, data processing, clerical work, paralegal and junior financial analysis — are disproportionately held by women. This reflects decades of occupational segregation: the systematic concentration of women into precisely the white-collar roles that AI is now automating most aggressively.

The consequence is a double exclusion. Women are overrepresented in the roles being displaced and underrepresented in the roles being created. Women currently make up just 28 percent of the global STEM workforce and only 22 percent of AI professionals. As the economy restructures around AI capability, women are being pushed out of the jobs being automated while remaining largely locked out of the jobs that are growing. Any organisation that describes itself as committed to gender equity while pursuing aggressive AI adoption without a parallel commitment to equitable transition is operating a contradiction it has not yet been asked to explain publicly. That question is coming.

Age and career stage: the missing ladder

Entry-level roles are not just first jobs. They are where expertise is built — where people develop the judgment, contextual knowledge, and professional relationships that make them valuable at senior levels a decade later. When that pipeline contracts at the current speed, the consequences ripple forward: organisations are quietly dismantling the infrastructure of their own future leadership while focusing on near-term efficiency gains.

Goldman Sachs data shows that unemployment among 20 to 30-year-olds in AI-exposed occupations has risen by almost three percentage points since early 2025 — notably higher than for older workers in the same fields. Harvard Business School research found that since the launch of generative AI, job postings for occupations involving structured and repetitive tasks fell 13 percent, while demand for analytical and creative work grew 20 percent. The market is bifurcating in real time.

Middle managers: the overlooked generation

Middle managers are those who sit between frontline employees and senior leadership — the people who translate organisational strategy into operational reality. They manage teams, typically of between three and fifteen people. They hold budget and resource accountability. They are the primary point of human connection between an organisation's intentions and its people's daily experience. In practice, they are the team leads, department heads, programme managers, operations managers, and people managers who carry organisational culture from the boardroom to the desk.

They are also the people AI is coming for next. By the end of 2026, research suggests that 20 percent of organisations will use AI to flatten their hierarchy — eliminating more than half of current middle management positions. These are not entry-level workers in their twenties with decades to adapt. These are experienced professionals in their thirties, forties, and fifties, often with mortgages, families, and deep professional identities built around the work they do.

Consider the specific position this creates.

The people most commonly tasked with implementing AI adoption in their teams — communicating the strategy, managing the questions, absorbing the anxiety — are middle managers. The same people whose roles are most directly in the sightline of the flattening that AI enables. They are being asked to lead a transition they know may end their own position, while maintaining the confidence of the people below them and meeting the performance expectations of the people above them.

That is not an organisational design oversight. It is a request for a particular kind of courage that most organisations are not acknowledging, let alone supporting.

When organisations remove this layer, they do not just reduce a cost centre. They remove the human connective tissue that makes organisations function: the person who notices when a team member is struggling, who translates strategy into something actionable on a Tuesday morning, who holds relationships across functions and knows where the real knowledge lives. That loss does not appear on a productivity dashboard. It appears, quietly, in the years that follow — in organisations that became more efficient and less capable at the same time.

Class, geography, and the limits of aggregate optimism

The workers with the least adaptive capacity are those in routine roles, with limited savings, in labour markets with fewer alternative employment options. The displacement of administrative and clerical roles does not primarily affect workers in major financial centres with strong safety nets and diverse local economies. It affects workers in mid-sized towns, regional offices, and communities where those roles have been among the most reliable sources of stable, reasonably paid employment for a generation.

The productivity gains from AI flow to shareholders and to the cities and sectors where AI capability concentrates. The displacement flows outward — to the communities that have already experienced the erosion of manufacturing, retail, and now administrative employment. The inequality this creates is not just economic. It is geographic and political. And it is building faster than most institutional responses can accommodate.

Does tighter employment legislation make a difference?

The evidence from Europe offers a genuinely instructive counterpoint to the displacement patterns visible in lighter-touch labour markets. A study of more than 12,000 European firms found no evidence that AI reduces employment in the short run — with AI adoption increasing labour productivity by around four percent on average while leaving headcount unchanged. The mechanism identified by researchers is capital deepening: AI augmenting worker output rather than replacing workers outright.

This is not simply a function of employment protection law making it harder to reduce headcount. The distinction that matters most is cultural and structural. Studies across Europe show that countries with stronger forms of workplace consultation — the Nordic countries and Germany in particular — see workers who are significantly more open to technological adoption. Social dialogue between employers and workers, through works councils and collective bargaining, encourages redeployment and retraining over job loss, and is identified as the primary mechanism through which technology is integrated in ways that reduce displacement.

The Nordic flexicurity model — combining meaningful workforce protection with well-funded transition support — offers a different answer to the question of how AI adoption can be managed responsibly. It is not that displacement does not occur. It is that the pace is deliberate, the support is real, and the conversation between employers and employees about what is changing and why happens before the decisions are made rather than after.

This matters for the paper's central argument: the workforce displacement currently visible in markets with lighter employment protection is not an inevitable consequence of AI adoption. It is a consequence of a particular approach to AI adoption — one that treats labour as a cost variable rather than an organisational asset. Europe's experience demonstrates that a different approach is both possible and consistent with genuine productivity gains.

---

3. The Systemic Contradiction

Businesses are eroding the conditions for their own success

Here is the logic of AI-driven cost reduction as it is currently being pursued. Automate tasks. Reduce headcount. Lower costs. Increase margins. Gain competitive advantage. Return value to shareholders. At the firm level, in the short term, the logic works.

Follow it far enough, however, and it collapses on itself.

Larry Fink has observed that the vast majority of wealth has moved to people who own assets, not to those who earn their income through work. If AI systematically removes income from the people who currently constitute the consumer base — the middle-income workers, the administrative professionals, the junior analysts, the middle managers — then the companies benefiting from AI productivity gains are simultaneously hollowing out the market they depend on. Companies are optimising at the firm level in ways that are destructive at the system level. Because every firm is doing it simultaneously, no individual actor feels responsible for the collective outcome.

> ***Henry Ford understood that his workers needed to be able to afford his cars. The AI economy has yet to produce a comparable insight — and the silence on this point is becoming harder to ignore.***

The Ford analogy is instructive, but only partially. Ford's manufacturing economy distributed wages broadly, across large workforces, in geographically dispersed communities. The AI economy concentrates gains in equity — in the hands of technology owners and shareholders — while distributing the labour costs downward. The mechanism for broad recirculation that made Ford's insight work does not exist in its current form. That is the problem that no amount of productivity data resolves.

The reskilling illusion — and the robotics problem

The dominant institutional response to this contradiction has been reskilling. Retrain displaced workers for the jobs of the future — the electricians, welders, engineers, and technicians needed to build the physical infrastructure that AI requires. The argument is appealing in its tidiness. It does not survive examination.

The first problem is assumption: that workers displaced from administrative, analytical, or professional roles want to retrain for trades. Some will. Many will not — whether for reasons of personal identity, physical capacity, age, geography, or family circumstance. Reskilling programmes that do not account for the actual human beings they are designed for are not solutions. They are comfort stories.

The second problem is arithmetic: the volume of infrastructure jobs created by the AI buildout is not equivalent to the volume of white-collar roles that will be automated. The AI infrastructure buildout is a finite construction project. You build the data centres, staff the power grid, lay the cables — and then that wave of demand passes.

The third problem is the most fundamental. The trades being proposed as the destination for displaced white-collar workers are themselves in the sightline of the next wave of automation. Humanoid robotics — physical AI systems capable of performing manual labour — are not science fiction. Tesla's Optimus programme is targeting 50,000 units in production in 2026. Elon Musk has predicted that work itself will become optional within 10 to 20 years as physical labour is automated at scale. Credible experts caution that the timeline for general-purpose humanoid robots replacing skilled trades is genuinely uncertain — the physical and coordination challenges are substantial, and ambitious production targets in this field have repeatedly been revised. But the direction of travel is not in doubt.

The same capital that is funding the automation of white-collar work is simultaneously funding the automation of the manual work being proposed as its replacement. The people being retrained are not part of that conversation. Reskilling as currently framed places the entire burden of adaptation on the individual worker while leaving the structural conditions that created the displacement entirely intact.

---

4. The Return on Investment Question

What the evidence actually shows

Here is the situation many scaling organisations are in right now. Budget is constrained. The capability and talent you need to grow is expensive and hard to find. AI is being positioned — by vendors, by investors, by competitive pressure — as the bridge: do more with less, move faster, reduce your dependency on headcount that is hard to recruit and harder to retain. The logic is appealing. The evidence for it is weaker than the conversation suggests.

If the disruption being caused by AI adoption were clearly and demonstrably delivering the returns organisations are projecting, the conversation would at least have a harder economic logic to it. The evidence suggests that for most organisations, it is not — and that fact deserves to sit alongside the human cost data in any honest account of what is happening.

The investment numbers are significant. Gartner projects spending on AI application software will approach $270 billion globally in 2026, with organisations spending between $590 and $1,400 per employee annually on AI tools alone. US companies spent $37 billion on generative AI in 2025. This is a capital commitment of extraordinary scale, made under intense shareholder pressure to show returns quickly.

The returns, for most organisations, are not arriving at the speed or scale projected. An MIT study from summer 2025 found that 95 percent of enterprise AI pilots are failing to deliver measurable financial returns. IBM's research found that only around 25 percent of AI initiatives deliver expected ROI, that just 16 percent have scaled enterprise-wide, and that only 29 percent of executives say they can measure ROI with any confidence. Deloitte's survey of 1,854 senior executives across Europe and the Middle East found that the average payback period for a typical AI use case is two to four years — significantly longer than the seven to twelve months typically expected for technology investments.

This creates a dynamic that is rarely named directly. Organisations have made large capital commitments. The financial returns are taking longer than projected to materialise. The board and investor pressure to show returns within a two-year window is acute. Headcount reduction is the fastest, most legible lever available — it shows up immediately on the balance sheet in a way that improved decision quality or better customer experience simply does not. Workforce displacement, in many cases, is not the primary strategic intent of AI adoption. It is the financial mechanism by which organisations recoup an investment that is not yet delivering on its other promises.

> ***Organisations are paying the human cost of disruption against a promise that has not yet been kept.***

Scrutinising the optimistic figures

Some statistics circulating in AI investment conversations deserve scrutiny before being accepted as evidence. A widely cited figure claims that companies investing significantly in AI see an 82 percent increase in revenue compared to those that do not. The figure is real — but its provenance matters. It comes from research commissioned by Basware, a financial software company with a direct commercial interest in demonstrating AI's value, surveying 400 finance executives about their own perceptions of performance. It measures correlation between AI investment levels and reported revenue growth — not causation. Larger, more sophisticated organisations tend both to invest more in AI and to grow faster for reasons that may have nothing to do with each other.

The independent research tells a more measured story. A working paper from Duke University's Fuqua School of Business, conducted with the Federal Reserve Banks of Richmond and Atlanta, surveyed nearly 750 executives and found what the researchers describe as a productivity paradox: companies reported AI-driven productivity gains averaging 1.8 percent in 2025, but when the researchers calculated implied gains using actual revenue and employment data rather than executive self-reporting, those gains were considerably smaller across all major industries. As the lead researcher put it plainly: the returns are not really hitting the top line yet in full force.

The gap between what executives believe AI is delivering and what the financial data actually shows is not a minor discrepancy. It is the difference between a story organisations

are telling themselves — under significant pressure to justify large capital commitments — and the reality that independent researchers can observe when they look at the numbers directly.

Where genuine innovation returns are emerging — and what they share

This is not to say AI is delivering no value beyond cost reduction. In pharmaceuticals and biotech, AI is compressing drug discovery timelines from years to months and enabling research that was previously impossible. In financial services fraud detection, pattern recognition at scale reduces losses that no human team could identify at speed. In retail personalisation and supply chain optimisation, predictive capability creates commercial value that did not previously exist. In each case, the value comes not from doing existing work more cheaply, but from doing something previously impossible or prohibitively costly.

What distinguishes these successes from the majority of deployments is not the technology. It is the organisational conditions surrounding it. McKinsey's analysis of more than 200 at-scale AI transformations identified workflow redesign as having one of the strongest contributions to meaningful business impact of all factors tested. BCG's research across 1,400 executives found that 10 percent of AI success comes from the algorithm, 20 percent from technology and data, and 70 percent from people, processes, and organisational change. The soft stuff, as BCG puts it, turns out to be the hard stuff.

Hill, Tedards and Wild, writing in Harvard Business Review in March 2026, identify a specific and frequently overlooked failure point: the gap between generating an insight and getting that insight into operational workflow. Innovations fail to scale, they argue, not because the ideas are flawed but because no one owns the crossing between insight and execution. What is needed is a named, funded role — what they call the bridger — whose explicit accountability is to translate across the boundaries where handovers typically fail. In AI terms: the gap between what the model recommends and what the organisation actually does is not a process problem. It is an identity and accountability problem. Treating it as the former is why so many AI deployments produce dashboards that nobody uses.

The conclusion across every major research source is consistent and directly relevant to this paper's central argument: the organisations cutting their people to recoup their AI investment are eliminating the very capability that the evidence shows is required to make that investment pay off. The experienced middle manager who understands the workflow, the junior analyst who knows where the data quality problems are, the administrative professional who holds the client relationship — these are not inefficiencies to be automated away. They are, according to the most rigorous research available, the primary determinant of whether AI creates value or simply creates disruption.

> ***AI is 20 percent algorithms and 80 percent organisational rewiring. The organisations cutting their people are eliminating the 80 percent.***

---

5. The 'Why' Question

Efficiency for whom, and to what end?

The question that is almost never asked in the AI adoption conversation is the simplest one: why? Not how will we implement this, or what will the ROI be — but why are we doing this, and who does it serve?

The honest answer, in most cases, is that AI adoption is being pursued primarily to reduce costs and increase margins, and that the benefits will flow predominantly to shareholders and senior leadership, while the costs will be borne by the workforce and, eventually, by the societies in which those workers live. That is not a conspiracy. It is the ordinary operation of corporate incentive structures in the absence of countervailing obligations.

But it raises a question that conventional CSR frameworks have been almost entirely inadequate to address: what do businesses owe the societies that make them possible? Businesses depend on educated workforces trained in schools funded by public money. They depend on infrastructure built by public investment. They depend on stable communities, functioning legal systems, and solvent consumers. A strategy that systematically degrades all of those things in pursuit of short-term margin is not a growth strategy. It is extraction. And extraction, historically, ends badly — including for the extractors.

> ***The question is not whether AI will transform your workforce. It will. The question is whether that transformation will be something you designed with your people — or something that happened to them.***

There is a harder version of this question that deserves space in the boardroom: if we automate significant portions of our workforce, and those workers are unable to find comparable employment, and their reduced purchasing power ripples through the communities where they live — are we comfortable with that outcome? Is that the organisation we want to be? Is that the legacy this leadership team is prepared to own?

For most leaders, the honest answer is no. But the incentive structures, the reporting frameworks, the investor expectations, and the competitive pressures all point in the opposite direction. The gap between intention and action is not primarily a moral failure. It is a structural one. And structural failures require structural responses.

Pace is a choice

One clarification matters here, because it has direct implications for what organisations can actually do. The pace of AI development itself is largely outside any individual organisation's control. The pace of AI deployment — the speed at which organisations choose to automate roles, eliminate positions, and restructure workforces — is not.

Governments around the world are grappling with exactly this distinction, arriving at different answers that reflect different values and different assessments of risk. Some are moving toward prescriptive frameworks requiring transparency, human oversight, and advance worker notification for AI-driven employment decisions. Others are prioritising innovation speed and competitive positioning, taking a lighter regulatory touch. Organisations operating across multiple jurisdictions are navigating this divergence in real time — and will increasingly need to make their own position explicit, regardless of what any single jurisdiction requires.

The point is not which approach is correct. It is that pace is a variable organisations control. An organisation that deploys AI to eliminate a layer of roles in a single financial quarter is making a different choice from one that deploys over three years with parallel investment in transition support. Both are adopting AI. Only one is exercising responsibility alongside it.

That distinction — between speed and care — is one that employees, communities, and eventually customers will notice and remember.

---

6. The Five Conditions

What makes responsible AI adoption actually stick

Acknowledging the problem is not the same as addressing it. Most change efforts — including most AI adoption programmes — focus on the visible surface: the tools deployed, the training delivered, the communications sent. The research is unambiguous that this is where 90 percent of organisational energy goes, and where the least value is created.

Drawing on established models of change management and psychological transition — including Bridges' work on transition, Bain's research on push versus pull dynamics, and the emerging literature on AI adoption — and tested through Intentional Edge's work with scaling organisations, we have found that five conditions consistently determine whether change of this magnitude actually sticks. They are not independent disciplines. They function as an ecosystem. One missing condition creates a gap the others cannot close, regardless of how much resource is directed at the remaining four.

Most AI adoption programmes are designed as though addressing one or two of these conditions is sufficient. It is not. The question for any leadership team is not whether they have thought about each condition, but whether all five are genuinely present — and what the honest answer reveals about why progress is stalling.

A useful diagnostic question, widely used among change practitioners and grounded in Bridges' insight that real transition produces changed behaviour rather than completed activities, is this: if we stopped the AI initiative tomorrow and disbanded the team, would the new behaviours continue? If the honest answer is no, the conditions have not been built. The change has been launched. The transition has not been managed.

1. Strategy — clear, prioritised direction with honest boundaries

Most organisations have an AI strategy. Fewer have one that honestly addresses what AI adoption will mean for the people within them. Strategy in this context means more than a technology roadmap or an efficiency target. It means a clear, shared account of why the organisation is adopting AI, what it expects that adoption to deliver, who will benefit, and what the organisation is committing to for those whose roles will change.

Without this, every other condition is built on unstable ground. People cannot navigate a transition whose destination has not been honestly described. And leaders cannot hold themselves accountable to outcomes they have not articulated.

2. Governance — single-threaded ownership for the human transition

Hill, Tedards and Wild's research identifies a consistent failure point in scaling innovation: nobody owns the crossing between insight and workflow. The same is true of AI adoption and its workforce consequences. Decisions about technology deployment sit in one function. Decisions about people sit in another. The accountability for what happens in between — for the human experience of the transition — belongs to no one.

Responsible AI adoption requires that someone — a named individual with real authority and real accountability — owns the human transition alongside the technical deployment. Not as

a welfare function responding to decisions already made, but as a strategic input shaping how and at what pace those decisions are taken. Without this, the most important work consistently falls between the gaps.

3. Safety and accountability — making truth-telling possible

After repeated waves of change programmes, employees develop what might be called learned futility — the understanding that this too shall pass, that the safest response to a new initiative is quiet compliance while waiting for it to fade. The consequence is what Bridges identified as the dashboard illusion: metrics that show green while actual behaviour on the ground remains unchanged.

This dynamic is acutely relevant to AI adoption. If it is not safe for people to say 'this AI tool does not work for my role' or 'I am worried about what this means for my future here', those concerns do not disappear. They go underground. They surface as resistance, disengagement, and eventual attrition of exactly the people whose knowledge and judgment the organisation most needs to make AI work. Psychological safety is not a soft benefit. In the context of AI adoption, it is the condition that determines whether the organisation gets honest feedback or managed compliance.

4. Leadership behaviour — modelling the transition, not just the destination

There is no neutral position for leadership during an AI transition. People follow what leaders do, not what they publish. A leadership team that communicates honestly about uncertainty, acknowledges what is changing and why, and visibly engages with the transition rather than managing it from a distance creates the conditions for genuine adoption. A leadership team that operates behind managed messaging, delegates the human conversation to HR, and measures only the efficiency outcomes creates the conditions for the dashboard illusion at scale.

This is particularly acute for middle managers, who are being asked to model a transition they may be experiencing themselves. The organisations that acknowledge this — that name the specific courage required of this layer, and invest in supporting it rather than simply expecting it — are the ones that will retain the connective tissue that AI adoption actually depends on.

5. Bridging capability — owning the handover across boundaries

The final condition is the one most consistently absent. Every significant AI adoption initiative involves a handover: between the technology and the workflow, between the strategy and the people, between the insight and the action. These handovers are where AI recommendations feel like threats to competence, where dashboards go unused, where pilots succeed and scaled deployment fails.

Responsible adoption requires that these handovers are explicitly owned — not assumed to happen through goodwill or process compliance, but actively managed by people whose role is to translate across the boundaries where momentum is most often lost. This is not a project management function. It is a negotiation of identity and a management of consequence. Organisations that treat it as the former consistently discover, too late, that it was always the latter.

---

7. The Leadership Invitation

There is a version of the next decade in which the AI transition is remembered as a moment when organisations chose efficiency over people — when the gains of technological transformation were captured by those who already owned assets, while the costs were borne by those who could least afford them, and when the warnings that were clearly visible were largely ignored because acting on them was commercially inconvenient.

There is another version. In this one, some organisations — not all, but enough to matter — chose to take their responsibilities seriously. They told their people the truth, early enough to matter. They invested in transition before displacement made it urgent. They asked who was bearing the cost of their decisions and allowed that question to change the decisions. They built the five conditions — the strategy, the governance, the safety, the leadership behaviour, and the bridging capability — that allowed their people to adapt rather than simply survive.

These organisations did not sacrifice performance for principle. They discovered, as the evidence consistently shows, that how you treat people during a moment of profound disruption is precisely the condition that determines whether you come through it with the capacity to innovate, attract talent, earn trust, and lead. The same people whose judgment, relationships, and institutional knowledge determine whether AI investments actually deliver value.

The leaders making AI adoption decisions right now are not making them in a vacuum. They are making them in the full knowledge of what is coming — knowledge that their peers have shared openly, that the data makes undeniable, and that is visible, for many of them, in the lives of the young people closest to them. That knowledge carries an obligation.

> ***The organisations that will define responsible leadership in the age of AI are not those that moved fastest. They are those that moved most honestly.***

Closing the gap between what is known and what is done is not comfortable. It requires honesty about what is coming, courage to act on what is known, and the organisational will to build the conditions that make responsible decision-making the default rather than the exception.

That is the leadership invitation this moment presents. The organisations that answer it well will not just survive the transition. They will define what it means to lead when the stakes are genuinely high — and when the people affected are not strangers.

About Intentional Edge

Intentional Edge is a consultancy focused on the people side of organisational performance. We work with scaling companies to build the leadership capability, cultural conditions, and organisational systems that allow businesses to grow without losing what makes them effective — or what makes them good.

Our work is grounded in a simple conviction: most organisations don’t have a change problem. They have a system problem. The blockers that prevent progress — unclear ownership, decisions that don’t land, effort that doesn’t translate — sit between teams and between functions, in the spaces that no single leader owns. Our role is to diagnose that system, resolve what is stuck, and build the capability and conditions that prevent it from recurring.

If this paper has raised questions that feel relevant to your organisation, we would welcome a conversation.

intentionaledge.com

A Note on How This Paper Was Made

This white paper was developed in collaboration with Claude, an AI assistant made by Anthropic. The research was gathered, the arguments were built and stress-tested, and the drafts were shaped through an extended, iterative conversation between human judgment and AI capability. The ideas, the instincts, the lived experience, and the editorial decisions are mine. The AI brought research capacity, drafting speed, and a willingness to be pushed back on.

We have chosen to say this openly — not despite this paper's subject matter, but precisely because of it.

If the argument of this paper is that organisations should be honest with their people about how AI is changing the nature of work, then the least we can do is model that honesty. This is what responsible AI adoption looks like in practice: transparent about its use, clear about what the human brings and what the AI contributes, and unwilling to present the result as something other than what it is.

It is also, if you look at it squarely, a demonstration of the central tension the paper describes. AI made this work faster, broader in its research, and more rigorously evidenced than it might otherwise have been. It did not replace the judgment, the experience, or the point of view. Whether that balance holds — as AI capability grows and the line between augmentation and replacement becomes harder to locate — is exactly the question this paper is asking organisations to take seriously.

We are asking it of ourselves too.

Nicole, Founder, Intentional Edge

Sources and attributions: William Bridges, Managing Transitions (1991, updated editions 2003 and 2016) — transition model, Endings/Neutral Zone/New Beginning framework; Bain & Company research on push vs pull change dynamics; Gartner research on executive-employee perception gaps in change management; Hill, Linda A., Emily Tedards, and Jason Wild, 'Why Great Innovations Fail to Scale,' Harvard Business Review 104, no. 2 (March–April 2026): 74–85, adapted from Genius at Scale (HBR Press, 2026) — bridger concept and innovation handover accountability; BCG AI Radar and Build for the Future 2025 — 10-20-70 principle and AI value gap research; McKinsey State of AI 2025 (1,993 respondents, 105 nations) — workflow redesign and high performer analysis; Deloitte AI ROI Survey 2025 (1,854 executives, Europe and Middle East); IBM CEO Study and Q4 Think Circle 2025; MIT GenAI Divide: State of AI in Business 2025; Duke University/Federal Reserve Banks of Richmond and Atlanta working paper on AI productivity 2026; Basware/FT Longitude AI to ROI Report 2025; BCG/Columbia Business School AI Sentiment Research November 2025; BlackRock Chairman's Letter to Investors 2026; Larry Fink, Davos opening remarks January 2026; Axios CEO interviews May 2025; WEF Future of Jobs Report 2025; WEF Labour Market Stories 2025; ILO/Poland National Research Institute GenAI and Employment Study May 2025; Goldman Sachs AI Workforce Analysis August 2025; Harvard Business School working paper (Srinivasan, 2026); CEPR research on AI adoption and employment across European firms 2026; Carnegie Endowment for International Peace, How Europe Can Survive the AI Labor Transition, February 2026; Perceptyx Transformation and Change Benchmark 2025; National Bureau of Economic Research 2025; Challenger Gray and Christmas 2025; Harvard Law School Forum on Corporate Governance December 2025; Built In/Tesla Optimus reporting 2025–2026; Fortune, CNBC, Financial Times reporting 2025–2026.