



## PAPER

Everyone

# What the 5% Do Differently

The organisations seeing real AI returns aren't using better technology - they've built four disciplines that make the same AI productive. Paper 2 introduces the ARGS framework: Agency, Architecture, Governance, and Scaling.

## The Same Tools, Different Results

---

In Paper 1, we diagnosed the problem: \$40 billion spent, 95% seeing zero return. The technology works. The implementation does not. But that diagnosis raises an obvious question: if 95% are failing, what are the 5% doing that everyone else is not?

The answer is uncomfortable for anyone hoping to solve this with a purchase order. The 5% are not using better AI. They are not on some secret waitlist for GPT-6. They use the same models, the same cloud infrastructure, the same subscription tiers as everyone else.

The difference is what happens between the AI and the decision.

BCG studied this divide in 2025 and found that the organisations generating 5x the revenue growth of their peers had done something structurally different: they had redesigned how work happens, not just what tools people use. Bain confirmed it from a different angle - tool-only deployment yields 10-15% efficiency gains. Workflow redesign delivers 25-30%.

The 5% invested in their people and their processes. The 95% invested in software licenses and waited.

---

## The Maturity Gap You Cannot See

---

Before we look at what the 5% built, we need to understand what most organisations look like from the inside - because the view is deceptive.

Imagine a company where 80% of staff use AI daily. Impressive adoption numbers. The dashboards look great. But look closer at how they use it. Most are doing one of two things: asking the AI to draft something (an email, a report, a summary) or asking it to answer a question. In both cases, they accept the first output, make minor edits, and move on.

This is Level 1-2 on a maturity scale that goes to 6. And here is the dangerous part: Level 2 users - the ones who have become fluent with prompts and feel productive - are the highest-risk group in the organisation. They generate the most content. They verify the least. They are the primary source of the \$9 million annual workslop cost we discussed in Paper 1.

The 5% organisations look different. They have the same Level 1-2 users doing daily tasks - but those users are operating inside structured systems that were designed by someone at Level 4 or 5. The AI is not a blank chat window. It is a guided process with built-in checks that catch errors before they reach a client or a board.

The difference is not the people. It is the architecture around them.

---

## Three Bands of Capability

---

The maturity gap maps to three distinct bands. Understanding them is the key to understanding why some organisations capture value and others do not.

### Band 1: Explorers (Levels 0-2)

The organisation has adopted AI. Usage is high. Verification is low. Leaders celebrate adoption metrics - how many seats are active, how many prompts are being sent - while the profit-and-loss statement tells a different story. Returns are linear at best. AI is a faster typewriter, not a force multiplier.

Most organisations are here. Most do not know it, because the output looks professional and the dashboards show green.

**Before AI maturity:** A consulting team writes client reports manually. Each report takes 8 hours.

Quality depends on the individual consultant's expertise. **After AI adoption (Level 2):** The same team uses AI to draft reports in 2 hours. They produce 4x the volume. But senior partners are now spending 3 hours reviewing each report because the AI introduces subtle errors - wrong regulatory references, logical leaps that sound right but are not, recommendations that do not fit the client's specific situation. Net time saved: almost zero. Net risk added: significant.

Addy Osmani, who leads developer experience for Google Chrome, named this the "70% Problem" in late 2024: AI gets you 70% of the way to a finished product rapidly, but the remaining 30% - edge cases, integration, quality assurance - takes just as long as it ever did. Worse, the AI-generated 70% often introduces problems that would not have existed if a human had built it from scratch. He called it the "AI speed tax" - the hidden cost of reviewing, debugging, and refactoring output that looked complete but was not. The 70% Problem is not limited to software. It applies to any knowledge work where "looks done" and "is done" are two different things.

## Band 2: Architects (Levels 3-4)

This is where the value curve bends. The organisation has shifted from buying tools to building systems. People are trained not just to use AI but to question it. The critical shift: leaders begin to own their AI-informed decisions rather than pointing at the machine.

The transition from Band 1 to Band 2 is painful. It requires admitting that speed is not the same as value. It means slowing down verification even as generation speeds up. It feels like a step backward - and for a few months, productivity often dips. This is normal. It is the investment period before the returns compound.

**After maturity investment (Level 3-4):** The same consulting team now works with structured templates. When a consultant opens the report tool, it already knows the client's industry, regulatory environment, and the firm's quality standards. The AI draft is not a blank-page generation - it is guided by the firm's institutional knowledge. A secondary check automatically flags logical inconsistencies before a human sees the output. Senior partners review half the reports they used to - and the ones they review are already 80% right, not 50%.

## Band 3: Orchestrators (Levels 5-6)

This is the Power Law upturn. Output is decoupled from headcount. One domain expert - someone who deeply understands the business and has learned to design AI systems - manages the verified output of what previously required a department. Not because the machine replaced the team, but because their deep knowledge now scales through the architecture they built.

**After orchestration (Level 5-6):** A supply chain director with 20 years of experience does not write reports or review drafts. She designs the system. She has built a verification engine that automatically handles routine logistics updates without human review. Complex supplier negotiations get flagged for her team. Only genuinely novel situations - a new regulatory regime, a geopolitical disruption - require her direct attention. Her output has tripled. Her stress has halved. Her expertise is encoded into the system so it works even when she is on holiday.

---

## Where Value Leaks Out

---

Even organisations that understand the maturity gap often fail to capture value because of three structural leaks. Recognising them is the first diagnostic step.

**The Architecture Leak** happens when AI is used for judgment-dependent work without structural verification. A consulting firm ships AI-drafted client reports without partner review. The first three clients say nothing. The fourth takes the engagement to a competitor - because the report contained a regulatory reference that did not apply to their jurisdiction, and the AI's confident tone meant nobody checked. **The Infrastructure Leak** happens when AI works from bad data. A sales team feeds their CRM data into an AI forecasting tool to generate quarterly projections. The projections look precise and professionally formatted - but 30% of the CRM records are duplicates with conflicting close dates. The AI does not know this. It produces a confident, beautifully structured forecast built on noise. **The Talent Leak** happens when organisations buy capability they cannot use. A company licenses an advanced AI coding assistant for 500 developers but invests nothing in teaching them when to override its suggestions. The tool generates code faster. The developers accept it faster. Six months later, the codebase has a maintenance debt that takes a year to unwind. A lead developer wrote publicly in late 2024 about deliberately removing all AI integrations from his code editors after noticing he felt "worse at his own craft" and "less competent at doing what was quite basic software development than a year before." The decay started, he said, with "stopping reading documentation - why bother when an LLM can explain it instantly?" Then he stopped thinking through problems before coding. Then he stopped recognising patterns he used to catch instinctively. This is the Talent Leak in slow motion - the tool was working, but the human was degrading. Anthropic's own research confirmed the pattern: in a controlled study, developers using AI coding assistants scored 17% lower on skill assessments than those who learned without AI. The largest gap was in debugging - the exact skill you need most when AI-generated code breaks.

Most organisations are leaking from all three simultaneously.

---

## The Four Disciplines That Close the Gap

---

The 5% have built their advantage on four disciplines that work together. Remove any one of them and the system fails in a predictable way. We call them **Agency, Architecture, Governance, and Scaling** - or ARGs.

### Agency: The Decision to Think

Agency is the prerequisite. Without it, nothing else works.

Agency is the shift from accepting what the AI provides to questioning what it provides. It means recognising that no matter how polished the output, it is a first draft - never a final answer. It means

providing the context, the intent, and the domain knowledge that the AI does not have and cannot generate.

For a leader, Agency means refusing to celebrate speed when you have not verified logic. It means asking "why did the AI produce this answer?" before asking "how fast can we ship it?"

**Without Agency:** The organisation builds efficient systems that produce unverified output at scale. Workslp accumulates. Errors are caught by clients, not by staff.

## Architecture: The Structure That Makes AI Reliable

Most organisations interact with AI through unstructured conversation - typing into a chat window and hoping for the best. This produces output that feels useful in the moment but creates invisible debt: functional systems that nobody fully understands, built on logic nobody has verified.

Architecture replaces chatting with structure. It means building documented, repeatable chains of reasoning where each step is traceable and each output can be explained. It means establishing clean data foundations so the AI works from accurate inputs, not organisational noise.

**Without Architecture:** A few brilliant individuals do great work with AI, but their methods are personal and unrepeatabe. When they leave, their capability leaves with them.

## Governance: The Accelerator, Not the Brake

Governance is the most misunderstood discipline. It is not compliance. It is not a committee that meets quarterly to review AI policy. It is not a document you file and forget.

Think of it like an airline pilot's relationship with safety regulations. Pilots do not experience safety checks as restrictions on their freedom. They experience them as the infrastructure that lets them fly faster and more confidently. The pilot who operates within well-designed guardrails flies faster than the one who must constantly worry about the unknown.

The same applies to AI governance. When your team knows where the boundaries are, they play more aggressively within them. Governance as compliance kills speed. Governance as living infrastructure accelerates it.

**Without Governance:** The organisation moves fast but is one hallucination or data breach away from a reputational crisis.

## Scaling: Decoupling Output from Headcount

Scaling is where the investment pays off. The economics of AI have inverted the cost of production - generating a draft, an analysis, or a code module costs almost nothing. But verifying that output remains entirely human.

The real bottleneck is not production. It is judgment. Scaling means designing systems where one expert can manage the verified output of what previously required a department. The machine handles the language. The human owns the meaning.

**Without Scaling:** The work is excellent and safe, but the gains are small and localised. The organisation fails to capture market share.

Scaling also carries a warning. Organisations that cut headcount before building the supporting architecture are not innovating - they are creating fragility. And organisations that stop hiring juniors today are building a talent gap: no entry-level pipeline means no senior capability in five years. The 5% use AI to accelerate junior development, not replace it.

---

## The Evidence

---

The case for ARGUS is not theoretical. It is visible in the financial data of the early movers.

- **PwC** invested \$1 billion in AI capability - not just tools, but training and workflow redesign. The result: 95% voluntary engagement (people chose to use the system because it worked) and 20-30% efficiency gains that were immediately redirected into new service lines.
- **JPMorgan** doubled its AI-augmented operations from 3% to 6% of the workforce in a single year. Their operations specialists - people who combine domain expertise with architectural thinking - saw 40-50% efficiency gains on complex tasks.
- **EY** reinvests 47% of its AI-driven productivity savings back into deeper capabilities and human development. This is the compound cycle in action: savings fund capability, capability generates more savings.
- Organisations that combine workflow redesign with human development see 25-30% productivity gains. Those that deploy tools alone see 10-15%. Same technology. Different returns.

The most rigorous study of this divide came from Harvard and BCG in 2023. They gave 758 consultants access to AI and measured what happened. The results split cleanly. Consultants who

used AI saw 12% more tasks completed, 25% faster, at 40% higher quality - but only when they maintained what the researchers called an "active posture." They identified two successful patterns: "Centaur" who strategically divided work between human judgment and AI generation, and "Cyborg" who integrated AI into every step but stayed critically engaged throughout. Both patterns share a common trait: the human never stopped thinking.

The study also revealed the danger. On tasks that fell outside AI's capability boundary, passive users were 19 percentage points more likely to produce incorrect answers than people not using AI at all. The tool made the passive users worse. The active users knew when to override it.

The relationship between these four disciplines is clear: organisations that have NIST (the "building code") but not ARGs have compliance without capability. ARGs provides the builder, the architect, and the development programme for the people who will live in the building.

---

## What This Means

---

The gap between the 95% and the 5% is not about budget, technology, or talent. It is about whether the organisation has done the structural work of building four disciplines that make AI productive rather than just fast.

Most organisations have skipped this work. They bought the tools, ran the pilots, and waited for the productivity magic to happen. It did not, and it will not - because the bottleneck was never the technology. It was always the system around the technology.

The good news: ARGs is buildable. It does not require replacing your workforce or buying a different AI platform. It requires investing in the capability to use what you already have.

In Paper 3, we will get specific. What does Monday morning look like when you start building this? Where do you begin? What do you build first? And how do you know it is working?

---

*This is Paper 2 of a three-part series from the Centre for AI Leadership (C4AIL). Paper 1: "Why Your AI Investment Isn't Working" diagnoses the problem. Paper 3: "Monday Morning: Where to Start" provides the practical playbook. For the full research framework, see "Orchestrating Intelligence: A Maturity Framework for Realising Human-AI Potential in the Age of Automation" - available from C4AIL on request. **Contact:** [hello@c4ail.org](mailto:hello@c4ail.org) | [centreforaileadership.org](https://centreforaileadership.org)*