

AI Anarchy vs. AI Autocracy

Two Adoption Paradigms, One Unresolved Risk – and
Which Approach Keeps You in Control

ANDING PERSPECTIVE | MAY 2026



Target Audience

Private Equity Funds and their **portfolio company leadership teams, COOs and CEOs** of mid-sized **European industrials, CHROs** and **Chief Transformation Officers** at **larger enterprises**

Summary

Most companies are not executing an AI strategy. They are reacting to one - competitor announcements, board pressure, vendor pitches - and calling it transformation. The result is activity without architecture.

We observe **two dominating adoption models** in practice. **AI Anarchy** hands out tools and waits. **AI Autocracy** designs the future state and mandates it. Both are presented here with their real trade-offs - not the sanitized version consultants use to justify a preferred recommendation.

The **more consequential question**, however, is not which model to choose. It is what each model does to the **organization's ability to function independently of AI** - and to negotiate with the small number of providers who will shortly control access to it. On that dimension, the two models produce opposite outcomes. One preserves human intelligence as a strategic asset. The other systematically destroys it in pursuit of efficiency.

This paper closes the loop: The adoption model question and the dependency question are the same question. Organizations that treat them separately will get the answer to both wrong.

Why Companies Are Getting AI Wrong

The AI adoption problem is not technical. The mostly technology works, and where it is not yet good enough, it will likely be soon. The problem is that most organizations are deploying a general-purpose capability without a framework for where it creates value, how it changes work, and what it costs when the terms of access change.

The Comprehension Gap at the Top

C-suite leaders making AI investment decisions typically have **limited direct experience** with the technology. They have **seen demonstrations**. They have read the productivity estimates. They have **not personally built a prompt chain**, debugged an agent workflow, or watched a large language model hallucinate a plausible-sounding contract clause. This matters because **AI capability is highly specific** - a model that drafts legal summaries fluently will fabricate case citations with equal fluency. Leaders who **cannot distinguish between these outcomes cannot govern the risk**.

The comprehension gap means AI strategy defaults to the people who procure the tools, not the people who own the processes. IT sets the governance framework. Innovation teams run the pilots. The **business leaders whose workflows are changing are consulted late**, if at all. The result is AI adoption that **optimizes for deployment speed and tooling cost, not for business value or operational integrity**.

The FOMO Trap

AI investment decisions are disproportionately driven by what peers appear to be doing. This is rational under uncertainty - but it produces convergent, undifferentiated adoption. When every company in a sector deploys the same tools for the same use cases at roughly the same time, the **productivity gains are competed away at the industry level**. The companies that capture lasting advantage from AI are not the ones that moved first with standard tools. They are the ones that **applied sharper thinking to better-defined problems** - and built the organizational capability to keep doing so as the technology evolves.

As the CompAller Imperative argues: AI is the compiler, your people are the programmers. The technology is becoming a commodity available to all at comparable cost. **The differentiating variable is the quality of human intelligence directing it**. That variable is what both adoption models handle differently - and what the dependency risk threatens to eliminate.

Two Adoption Models

AI Anarchy: Distribute and Hope

Provide employees with professional-tier AI subscriptions. Set minimal policy. Wait for value to emerge. The model is built on one defensible premise - that **frontline workers understand inefficiencies better than leadership does** - and one convenient assumption: That capability will translate into value without design.

Example	<i>A PE fund issues Claude Pro licences to all employees across a portfolio company. No use-case mandate. Within four months, the finance team has automated monthly reporting packs, procurement is running AI-assisted contract review, and one analyst has built a news monitoring agent for portfolio signals. None of this was directed.</i>
----------------	---

Example	<i>A 600-person professional services firm rolls out Copilot firm-wide with a light-touch policy. Six months in, adoption has clustered spontaneously in two areas — proposals and research synthesis. Engagement elsewhere is near zero. The firm has learned where AI creates natural pull. It has no mechanism to convert that insight into standardised capability or to quantify what the non-adopting 70% of the workforce is costing in forgone productivity.</i>
----------------	--

ADVANTAGES	DISADVANTAGES
<ul style="list-style-type: none"> ▪ Surfaces unexpected use cases. ▪ Immediate learning. The organization accumulates real usage data from day one, rather than assumptions from a design workshop. ▪ Genuine fluency develops. Employees who own their AI workflows understand what the tool actually does - and where it fails. ▪ No upfront design cost. Deployment begins before the organization has to make hard decisions about process change. 	<ul style="list-style-type: none"> ▪ No cost control. Token consumption scales with curiosity, not value. Enterprise spend can reach multiples of a structured deployment. ▪ No institutional knowledge. Individual discoveries stay individual. The organization learns in parallel, not cumulatively. ▪ Invisible risk exposure. Data governance, confidentiality, and compliance violations happen in the spaces between acceptable-use policies and actual usage. ▪ Management cost deferred, not eliminated. Ungoverned adoption requires expensive intervention later - cultural change programs, retroactive governance, process rationalization.

AI Autocracy: Design and Mandate

Map the operating model and **identify the highest-leverage AI use cases** top-down. Design the future-state process and respective AI tool usage. The model is built on one correct observation - that **undirected experimentation is inefficient** - and one **dangerous assumption**: That the **processes being redesigned will remain stable**, and that the AI embedding them will remain affordable.

<p>Example</p>	<p><i>A machine manufacturer runs a structured AI program. Four processes are prioritized - quality control, spare parts forecasting, technical documentation, and complaint routing. Future-state workflows are designed with AI roles explicitly specified. Tooling is integrated into existing systems. Adoption is mandatory. Twelve months in, all four processes are performing against KPIs. The organization has also removed the human inspection capability that the AI replaced.</i></p>
-----------------------	--

<p>Example</p>	<p><i>A regional bank treats AI deployment as a risk-governed program. The CRO, COO, and CISO co-own a use-case classification framework. Only tier-one cases - full audit trail, no autonomous decisions, complete explainability - are approved. A central AI team owns design and deployment; business units are users. The framework is slow. It produces zero compliance incidents. It also produces an organization that cannot articulate what it would do if the approved tooling became unavailable or doubled in price.</i></p>
-----------------------	---

ADVANTAGES	DISADVANTAGES
<ul style="list-style-type: none"> ▪ Controlled cost. Scope is defined upfront; token spend is bounded and forecastable. ▪ Measurable ROI. Value creation is specified before deployment, not inferred after. ▪ Risk is manageable. Defined guardrails, approved tooling, and clear ownership make governance tractable. ▪ Institutional process capability accumulates. Blueprints, governance templates, and performance benchmarks compound over time. 	<ul style="list-style-type: none"> ▪ Filters out bottom-up innovation. The central design team's imagination is the ceiling for what AI can contribute. ▪ Removes human capability by design. Processes redesigned around AI are stripped of the people who previously ran them. That capability does not come back quickly. ▪ Creates structural vendor dependency. Deep process embedding means switching providers or reverting to human operation is not a quarterly decision - it is a multi-year transformation. ▪ Leadership bottleneck. The model requires genuine AI literacy at the top. Organizations that lack it produce autocracy in form but anarchy in substance.

The Comparison

The two models **differ on every operationally relevant dimension**. One dimension - reversibility - has been systematically underweighted in how organizations evaluate their options:

Dimension	AI Anarchy	AI Autocracy
Speed of Deployment	Immediate - no design phase	Weeks to months - design first
Cost Profile	Uncontrolled - token sprawl is the norm	Bounded - defined scope, measurable ROI
Innovation Capture	High - crowd intelligence at work	Capped - central filter misses edge cases
Risk Exposure	High - data, compliance, quality	Manageable - guardrails built in
Human Capability Retained	Yes - individuals stay close to the work	No - automated functions hollow out fast
Vendor Dependency	Diffuse - many subscriptions, easy to cut	Structural - embedded in core operations
Reversibility	High - remove licenses, resume prior state	Low to none - process redesign is a sunk cost
Leadership	Low upfront, high to rescue later	High upfront, predictable, thereafter

The reversibility row is the one that matters most over a five-year horizon. Anarchy is recoverable. Autocracy, executed thoroughly, is not.

The Dependency Risk: The Argument Neither Camp Makes

Both adoption models **generate AI dependency**. The **anarchy** model creates **diffuse individual reliance** - workers who have restructured their daily output around AI assistance will take a productivity hit if access is removed. The **autocracy** model creates something qualitatively different: **Structural organizational dependency**, embedded in core processes, with the **human fallback capacity removed**. The distinction matters because of where AI market is heading.

The Oligopoly in Plain Sight

Three to five organizations, incl. Anthropic, OpenAI, Google DeepMind, Microsoft, control the frontier AI market. The **infrastructure cost of training and serving state-of-the-art models is a structural barrier to entry** that open-source alternatives are narrowing but have not removed. This is not a competitive market. It is an **oligopoly in formation**, and the pricing behavior that characterizes early-stage oligopolies - aggressive on adoption, extractive on retention - is already visible in enterprise contract structures.

Current **token pricing reflects a strategic decision to drive adoption**, not equilibrium economics. The analogy to enterprise software is exact: **Oracle and SAP spent twenty years building irreplaceable process dependencies** into corporate infrastructure, then extracted the pricing power those dependencies created. **AI providers are running the same playbook on a faster timeline**. The difference is that **AI dependencies are being built into cognitive processes** - how organizations think and decide - not just administrative ones. The switching cost is correspondingly higher.

The Reversibility Problem

When an organization redesigns its quality control line around AI visual inspection, it does not add a tool. It **removes the trained workforce that previously did the job**, reorganizes the physical process around the AI's throughput and error tolerance, and integrates the AI's output into downstream quality systems. Reversing that - **reinstating human capability, reconfiguring the process, rebuilding the institutional knowledge** - is not a procurement decision. It takes years and costs more than the original deployment.

The **anarchy model creates a softer version of the same trap**. Knowledge workers who have offloaded drafting, research synthesis, and analysis to AI over two or three years have not kept those muscles in shape. They can produce output faster with AI than without it. Without AI, output quality drops and volume drops further. The individual dependency is real - it just does not appear on any balance sheet until access is disrupted.

The question organizations are not asking: if your primary AI provider doubled prices next year, what would you do? Most have no answer. That is not a technology risk. It is a strategic failure.

What to Do About It

Dependency cannot be fully avoided - the productivity gains are too material to forgo in competitive markets. It **can be managed**. Three requirements:

- **Maintain human fallback capability in critical processes.** Any process embedded with AI must retain a documented, periodically exercised human operating mode. Not as a business continuity formality - as a negotiating asset. A provider that knows it can be replaced, even at cost and delay, has meaningfully less leverage than one facing an organization with no alternative.
- **Architect for portability.** Where AI is integrated into core systems, use **abstraction layers - provider-agnostic agent frameworks**, standardized prompt interfaces - that decouple the intelligence from the infrastructure. Switching models should be a configuration change, not a re-implementation.
- **Stress-test the cost model.** Model the operating cost impact of a two-to-five-times increase in AI access pricing. If the answer is 'the business case breaks', the dependency is already too deep. Fix that before the provider knows it too.

Closing The Circle: Which Model Keeps You in Control?

The dependency question and the adoption question are the same question. Which model is better positioned to resist provider pricing power, retain the organizational intelligence to challenge AI outputs, and reverse course if access deteriorates?

The **answer is uncomfortable for advocates of structured transformation**: Anarchy has a structural advantage on every dimension of strategic resilience. Not because it is a better adoption model. It is not. But because it **preserves the one asset the autocracy model systematically destroys** - human capability that remains close to the work.

Why Anarchy Preserves Strategic Resilience

In the anarchy model, **individuals discover and own their AI workflows**. They understand what the tool is doing at a task level. They know where it fails - because they have encountered those failures directly and absorbed them as professional knowledge, not as an incident report. This means **they retain the cognitive baseline to challenge AI outputs**, to recognize when a model is producing confident-sounding nonsense, and to work around failures when they occur.

More importantly, the processes themselves are not redesigned around AI. A knowledge worker using AI to accelerate their existing work can stop using it and continue at a lower speed. A manufacturer whose quality control line was rebuilt around AI cannot. **The anarchy model's apparent weakness - it does not eliminate human effort - turns out to be its strategic strength.** Human capability is not removed, but augmented. That distinction is the difference between dependency and leverage.

Why Autocracy Creates a worse Dependency Profile

The **autocracy model is designed to be thorough**. That thoroughness is precisely what makes the **dependency hard to escape**. Processes are redesigned and headcount is reduced. The model will likely deliver better business outcomes and P/L impact. However, the institutional memory of how the process worked without AI atrophies within eighteen months. When the AI fails - through a model deprecation, a price increase that breaks the business case, or a provider outage - there is nothing to fall back on.

This is not a theoretical risk. **Enterprise AI contracts are already showing price escalation clauses**, usage-based pricing that penalizes volume growth, and model deprecation timelines that force expensive re-integration. Organizations that signed three-year transformation programs around GPT-4 capabilities are now discovering that the model they designed for no longer exists in the same form. The automation they built is not broken - but the prompting logic, fine-tuning, and workflow integrations require rework on a timeline set by the provider, not the organization.

The **autocracy model is the dominant choice if one believes that the AI market will be competitive and AI models will remain interchangeable**. Hedging against this potential outcome will require a smart approach that merges anarchy and autocracy

The Right Answer: Structured Anarchy

Neither pure model is the correct answer. The **right approach borrows the anarchy model's preservation of human capability and the autocracy model's discipline around value creation** –and adds an explicit dependency management layer that neither model currently includes.

In practice this means four things:

- **Deploy AI as augmentation, not replacement.** The unit of value is not 'headcount eliminated by AI'. It is 'output per person increased by AI while the person retains the ability to produce output without it'. The moment AI replaces rather than augments, the dependency clock starts.
- **Invest in structured AI literacy, not just tool access.** The CompAller argument applies directly: an organization full of people who can direct AI precisely - decompose problems, construct prompts with rigor, recognize failure modes - is resilient to provider changes in a way that an organization full of tool users is not. Tool users need the tool. CompAllers need a compiler - and compilers are interchangeable.
- **Build a use-case portfolio with explicit reversibility ratings.** Not every AI integration needs to be reversible. Some processes are genuinely better redesigned around AI capability. But that should be a deliberate choice, made with eyes open to the dependency it creates, not an accidental outcome of a transformation program optimized for efficiency.

- **Treat the AI provider relationship as a procurement risk**, not an IT vendor relationship. Negotiate multi-provider optionality. Avoid exclusive integrations. Maintain parallel capability - even at cost - in the processes that matter most. The organizations that will negotiate effectively with AI oligopolists in 2028 are the ones building negotiating leverage now, before they need it.

What Leadership Must Do Differently

Three decisions, made explicitly, separate organizations that will manage this well from those that will not:

- **Decide the adoption model consciously.** Not by default, not by procurement cycle, not by which vendor runs the best executive briefing. Anarchy is not free. Autocracy is not safe. Both have a cost structure that extends well beyond the subscription line in the P&L.
- **Add reversibility to every AI business case.** Before any process is redesigned around AI, quantify the cost of reversing that redesign. If it is prohibitive, either reconsider the depth of embedding or explicitly provision for the dependency - as a risk item on the balance sheet, not an asterisk in a footnote.
- **Build CompAller capability, not tool fluency.** The differentiating variable across both adoption models is the **quality of human intelligence directing the AI**. That is a **capability that must be built deliberately** - In how people are trained, how work is structured, and how performance is measured. Tool fluency - the ability to use the interface - is not the same thing. It is the floor, not the ceiling.

The window to build AI-resilient organizations - ones that extract value from the technology without surrendering the capacity to function, adapt, and negotiate independently - is narrowing. Every month of deep process integration without a reversibility strategy is a month of negotiating leverage transferred to the provider.

Conclusion

Anarchy OR autocracy is not the right question. The right question is: What kind of AI-adopting organization do you want to be in five years - one that uses AI as a competitive multiplier it controls, or one that depends on it in ways it cannot?

Anarchy, done without investment in human AI capability, produces individual dependency that is invisible until it matters. Autocracy, done without a reversibility strategy, produces structural dependency that is very visible - to the providers who will price accordingly.

The organizations that navigate this well will **not be the fastest adopters or the most disciplined designers**. They will be the ones that **treated AI adoption as a strategic question**

about organizational intelligence - and made sure that intelligence remained in the organization, not just in the model.

Further reading

If you are interested in further case studies and guidance on transforming organizations, we recommend these thought papers and case studies:

- [Steering organizations towards a modern operating model - The DIPOC Framework](#)
- [Replace and Accelerate: How to Run a Smart HR Software Transition that Creates Value](#)
- [Accelerating organizational learning with an AI point-solution](#)
- [Taking the Art out of Artificial Intelligence](#)
- [The CompAller Imperative](#)

All papers are available on our website: www.anding.company

Author



DR. MARKUS ANDING
Founding Partner & Managing Director

Anding & Company, Munich
markus.anding@anding.company

Imprint

Editor: Anding & Company GmbH, Bavariaring 24, 80336 Munich

Responsible Partner: Dr. Markus Anding

Photo: Midjourney

Copyright © Anding & Company GmbH. All rights reserved.