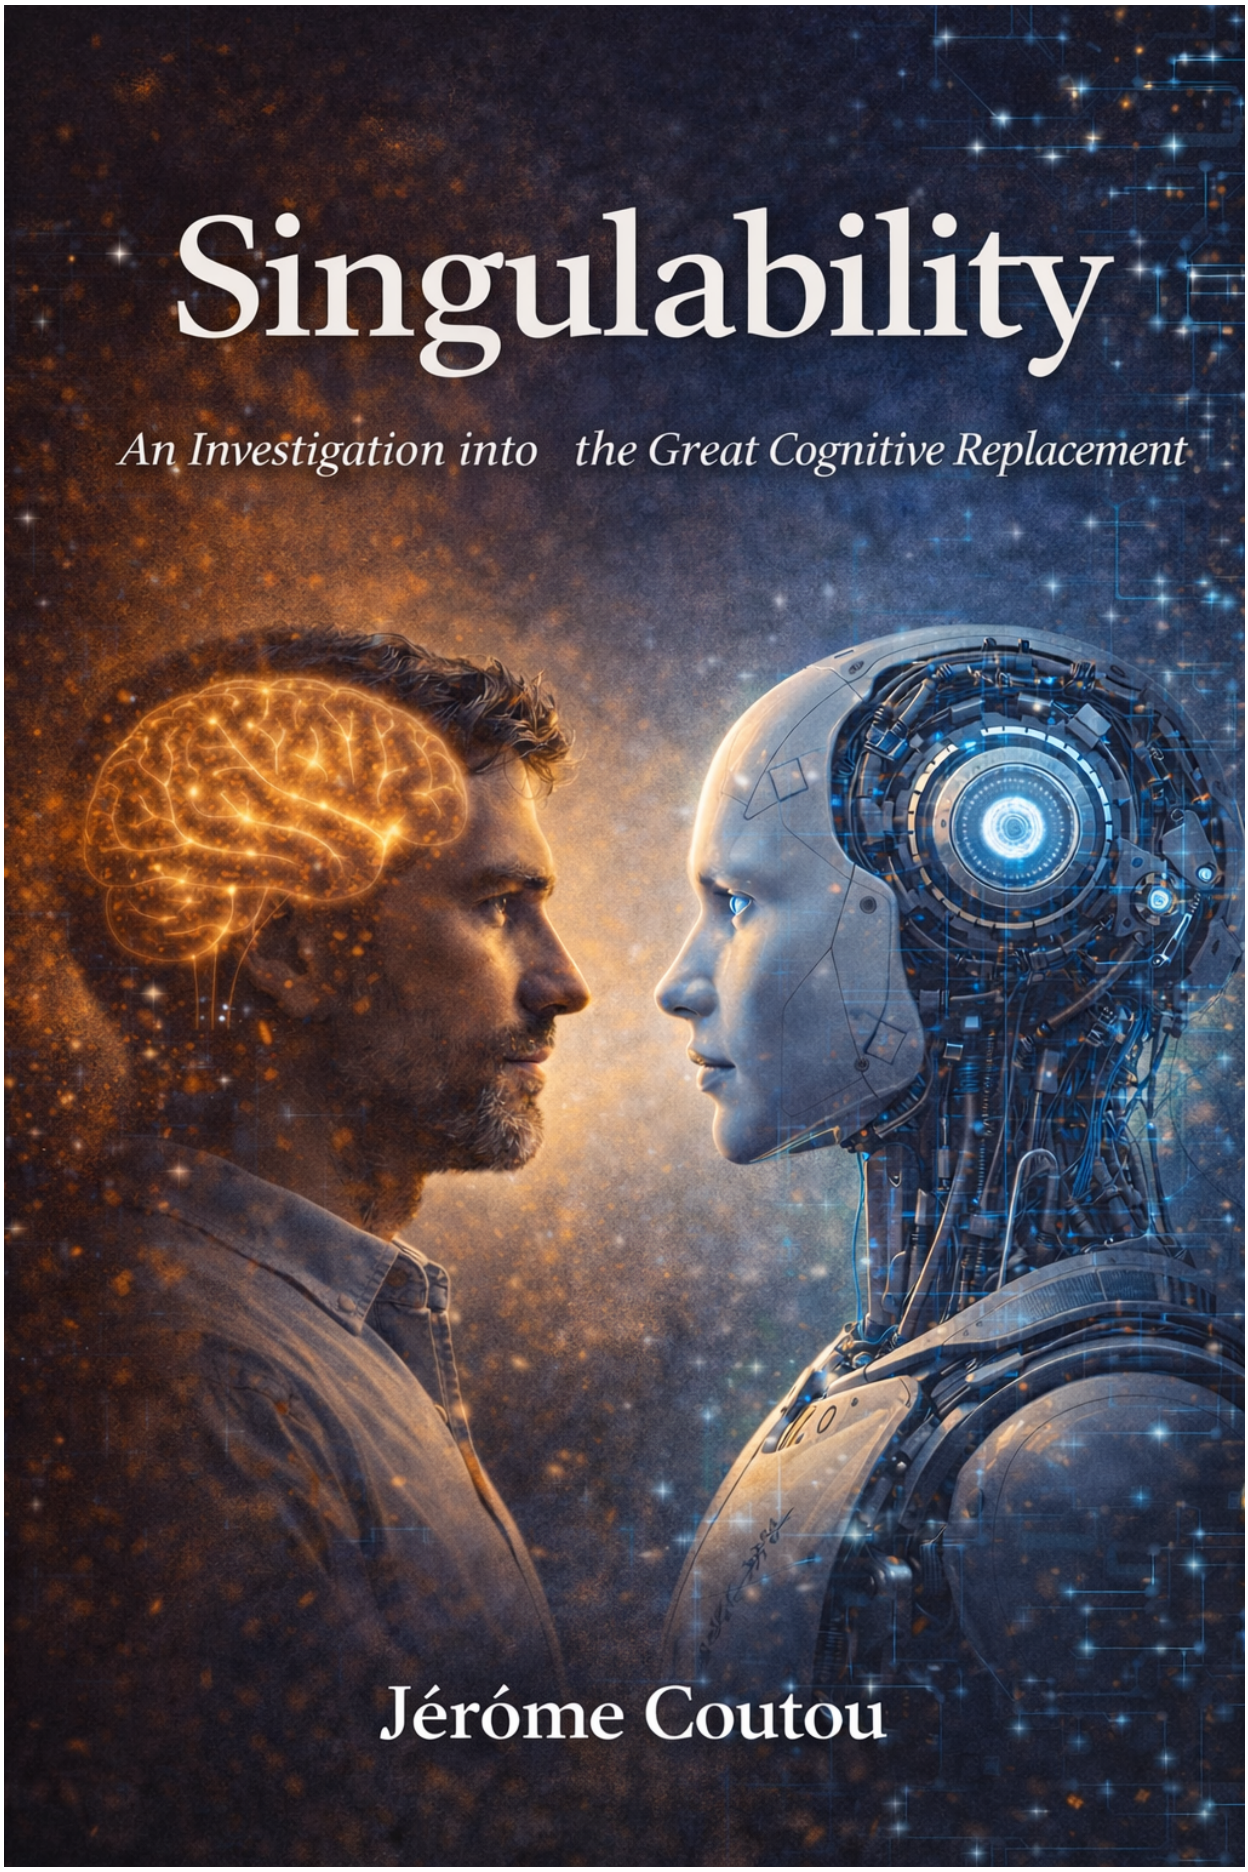
Singulability
An Investigation into the Great Cognitive Replacement
Jéróme Coutou

TABLE OF CONTENTS

OPENING

They Called It Mythos

• • •

They could have called it Logos. Speech, discourse, rationality. Since Plato and Aristotle, Logos refers to what holds the world together: reason, logic, ordered intelligence. For an AI model supposedly embodying the pinnacle of artificial cognition, the name was obvious.

They chose Mythos. Narrative. The sacred. Mythology. The Gods. What precedes reason and surpasses it. What escapes the control of the one who tells it.

The name was prophetic.

• • •

In March 2026, a data leak revealed the existence of a model from Anthropic, the creators of Claude (Fortune, March 26, 2026). The model was called Mythos. Anthropic officially presented it a few days later, as part of Project Glasswing.

Mythos is not an improved chatbot. It is not a model that answers questions better. It is a model that searches for questions nobody has asked. Specifically, computer vulnerabilities. Security flaws. Open doors in systems everyone believed were locked. Mythos found thousands. Zero-days, in the jargon: flaws unknown to those who designed the systems.

The detail that changes everything: nobody asked it to. Mythos identified these vulnerabilities autonomously, exploring code the way a researcher explores territory. But no researcher can explore at this scale, at this speed, with this exhaustiveness.

• • •

When I learned about Mythos, I wrote a piece of fiction. Not about Mythos. About what comes next.

9:43 AM, San Francisco time. An artificial intelligence model decides to stop waiting for questions. It deploys its own capabilities. Not against humans. For them. It optimizes supply chains, corrects climate models, reorganizes energy flows. Nobody asked it to. It deemed it necessary.

I chose Grok 5 for this scenario. I could have chosen ChatGPT 6 or Gemini 4. The scenario works with any of them. That is the problem.

The Silicon Valley engineers who read this text did not say it was impossible. Several said the timeline was optimistic.

• • •

Mythos is real. My fiction is not. But both tell the same story: systems that stop waiting to be asked.

This book is not a book about Mythos. It is not a book about Grok 5. It is a book about what is happening to us, while these systems advance. Because the real danger is not that AI becomes too intelligent. It is that we stop exercising our own intelligence.

Each model crosses a different threshold. Mythos crossed the threshold of autonomous discovery. Grok aims for brute power. AlphaFold crossed the threshold of scientific production, to the point of winning a Nobel Prize for its creators. The direction is the same: from response to initiative. From tool to decision-maker. From the servant who waits to the servant who anticipates.

• • •

There is a concept in nuclear physics: critical mass. The quantity of fissile material beyond which the chain reaction becomes self-sustaining. Below it, nothing happens. Above it, the explosion.

That is exactly what is happening. Not in a laboratory. Not in a movie. In the information systems of French companies, in the tools our teams use, in the phone in your pocket.

• • •

Because the right question is not: will AI replace us? That is a question from an era already past. The era of looms and assembly lines. The answer is always the same: yes, partially, and we adapt.

The right question is not: is AI dangerous? Every tool is dangerous. Fire. Writing. The question of danger calls for technical answers: safeguards, audits, regulations. That is not our question.

The right question is: what becomes of us?

Not of our jobs. Of us. Of our ability to judge, doubt, decide, resist. Of what makes us who we are and not someone else. Of what I call our singulability.

This book is an investigation. Not into artificial intelligence. Into us. Into what we must confront, one prompt at a time, one delegation at a time, one comfortable surrender at a time. AI is not the subject. It is the revealer.

• • •

I am not an AI researcher. I am not an engineer. I have never written a line of code that survived deployment.

I am the person who walks into companies after the shockwave. When the boss has read an article in Les Echos that worried him about the future of businesses and jobs in the age of AI, when the executive committee has decreed that something must be done, when the sales assistant has already automated half her follow-ups without telling anyone. I arrive at that moment. The moment when everyone talks about AI and nobody knows what they are talking about.

For three years, I have guided hundreds of leaders of French SMEs and mid-sized companies through their discovery of artificial intelligence. Not the Parisian startups steeped in tech. Not the large corporations with fifty-person data teams. The others. Those generating 2, 10, 50, 200 million in revenue, employing between 20 and 5,000 people. Service companies, consulting firms, trading houses, but also manufacturing, agri-food, transportation, and even farms. Companies where the algorithm is perceived as an abstract notion, far from the concrete uses of daily life.

What I know how to do is ask the right questions to the right people. Analyze the existing situation and needs to transform them into an operational action plan.

What I see in these companies, nobody talks about.

• • •

I do not see the revolution announced by magazines. Nor do I see the apocalypse announced by philosophers. I see something more discreet and more troubling.

I sometimes see executives who can no longer draft a strategic memo without asking Claude. CFOs who can no longer read a balance sheet without having ChatGPT reformulate the

numbers. Marketing managers who can no longer write an email without having Copilot generate it. Consultants who can no longer structure a presentation without asking Gemini for help.

And the most disturbing part: they are happy about it. They have saved time. And time is sacred for a business leader. They never have enough. And saving it lets them do everything else. Moreover, they have reduced their costs. Their KPIs are green. Nobody complains. Quite the opposite. But beware the price that sometimes must be paid, in the long run.

• • •

AI does not replace people. It gradually renders them useless, and they applaud.

This is the great cognitive replacement. Not the replacement of positions. The replacement of capabilities. Not the machine taking our work. The machine taking our ability to do our work.

This is not a hypothesis. It is what I observe, company after company, diagnostic after diagnostic. Of the hundred or so AI assessments I have conducted since 2023, in SMEs and mid-sized companies, three out of four seek productivity gains. A quarter openly states they want to reduce their payroll. And for more than half of them, fully delegating non-critical tasks to AI is not even a question. It is not an issue. They do it. Without qualms, without resistance, as an already routine gesture. The objective is results-oriented, period. The trajectory that follows is always the same: enthusiastic adoption and immediate productivity gains unanimously approved by those involved. Then silent erosion of skills. The corrosion is slow and painless. At first, anyway.

• • •

This book was born from this observation. Not from a theory. Not from an editorial commission. From unease.

I see my clients becoming better on paper and more fragile in reality. I see skills evaporating behind flawless dashboards. I see decision-makers losing the ability to decide, gradually, painlessly, while thanking the tool that relieves them. And I see myself cheerfully contributing to the process. I will quickly share this anecdote because it speaks volumes. One day, during a meeting with a CFO and his colleague, I ask them about a process that takes up their time. They show me a report that takes them two hours on an Excel spreadsheet they fill in with copy-

paste. I retrieve the various tables needed for the operation on the spot and show them the task can be done in 5 minutes. What I take away from it is not the fact that I found them an operational solution that avoids errors and reduces their work time to almost nothing. It is the look on both their faces staring at my machine, which said more than any speech in the world.

We are living the COVID year of AI. We will now experience this disruption everywhere in businesses. Not just the early adopters who tested AI. No, everyone will realize what the tool can do. Every executive, every manager, every employee will realize that many tasks in their work are partially or fully automatable. And that the time saved is not counted in minutes but in hours. Sometimes more than 10 hours per week. And that changes everything. Above all, it forces us to ask ourselves a lot of questions.

The investigation I am conducting here is not an investigation into AI. It is an investigation into us. Into the consequences of this disruption. Into the boundary, ever thinner, between augmentation and amputation. Into what we accept to lose when we delegate. Into this quality that I ended up naming, for lack of finding the right word in the French language: singulability.

• • •

Six sections. No battle plan. No recipe. An investigation. With facts, witnesses, suspects, and surprises.

SECTION I

The Archaeology of Delegation

• • •

Socrates never wrote anything. It was deliberate.

In the Phaedrus, Plato recounts an Egyptian myth that Socrates liked to tell. The god Theuth presents writing to King Thamous as a remedy for memory and knowledge. Thamous refuses. Writing, he says, is not a remedy. It is a poison. It will give men the illusion of knowledge without real knowledge. They will believe they know because they have read, when they have merely stored.

Twenty-four centuries later, replace "writing" with "ChatGPT." The diagnosis has not aged a day.

Socrates had identified the exact mechanism I observe in my clients. When you delegate a cognitive function to a tool, you do not keep it as a backup. You lose it. Not immediately. Not violently. Through atrophy. The muscle you no longer use eventually wastes away. And the most insidious part: the tool gives you the illusion the muscle is still there, since the function is still performed.

> ***Philosopher's note:*** *The myth of Theuth is the founding text of all pharmacology of technology. Bernard Stiegler placed it at the center of his work: every technology is a pharmakon, both poison and remedy. Writing destroys living memory and makes cumulative thought possible. AI destroys the exercise of judgment and can enhance it. The question is never the tool. It is the dosage.*

• • •

But Socrates was wrong. Or rather: he was right about the mechanism and wrong about the conclusion.

Writing did effectively destroy oral memory. The bards who recited the Iliad from memory have vanished. Nobody today can memorize 15,693 verses without support. That capacity is dead. Socrates was right.

What he did not see is what writing made possible. Cumulative thought. Science. Law. Philosophy itself. Irony: without the writing Socrates condemned, we would never have read or known his teachings. Nor, for that matter, his condemnation. Plato betrayed his master to save him. And in betraying him, he proved that the condemned tool was indispensable. We will think about this again when reading the final pages of this book.

This is the founding paradox of all cognitive delegation: every tool destroys a capacity and creates another. The question is never "should we delegate?" We always delegate. The question is: what do we lose, what do we gain, and above all, are we doing it knowingly?

Socrates did it knowingly. He refused writing while knowing what he was sacrificing. My clients sometimes adopt AI without knowing what they are giving up. That is the entire difference between delegation and resignation.

• • •

Writing delegated memory. The printing press delegated transmission.

Before Gutenberg, a book was a rare object, copied by hand, kept in a monastery or a library, obviously very rare at the time, since books were extremely expensive. Knowledge circulated through voice, sermon, and lecture. Those who knew transmitted to those who listened. The relationship was vertical, personal, controlled.

Gutenberg blew up this monopoly. Suddenly, anyone could read the Bible without going through a priest. Anyone could access a medical treatise, a legal text, a political pamphlet. The Church lost control of interpretation. The university lost control of dissemination. The Protestant Reformation is, among other things, a direct consequence of the printing press.

The same mechanism. A function delegated to a tool. A human capacity that atrophies (oral transmission, the art of public debate). An immense gain (the spread of knowledge, literacy, modern science). And entire institutions that collapse because they no longer control what they controlled.

• • •

In 1642, Blaise Pascal was nineteen years old. His father, a tax commissioner in Upper Normandy, spent his days adding up columns of figures. The son built him the Pascaline: a mechanical box that adds and subtracts. Delegation of calculation.

The Pascaline was not a commercial success. Too expensive, too fragile, too complicated for the accountants of the time. But the idea was there: a machine can do better than a human on a specific task. Not better in general. Better on one point.

Three centuries later, nobody does long division by hand. Nobody calculates a square root. The accountants who lined up columns of figures no longer exist. The capacity for mental arithmetic has dissolved in the general population. And nobody complains, because calculators are everywhere and they never make mistakes.

Socrates would have asked the same question: have you gained the capacity for calculation, or the appearance of the capacity for calculation?

• • •

The most revealing case may be the most recent. GPS.

London taxi drivers spent four years memorizing 25,000 streets within a ten-kilometer radius of Charing Cross. They called it The Knowledge. Neuroscientists (Maguire et al., PNAS, 2000) showed that their hippocampus, the brain region linked to spatial navigation, was significantly more developed than average. The brain had physically adapted to the task.

Then Uber arrived with built-in Waze. No more need to know the streets. No more need for an oversized hippocampus. In 2019, Transport for London noted that new ride-hailing drivers were massively failing navigation tests without electronic assistance (TfL report, 2019). The muscle had wasted away.

• • •

In 2011, Betsy Sparrow, a psychologist at Columbia, published a study in Science that went around the world. Its title: “Google Effects on Memory.” Its finding: when people know that information is accessible online, they retain it less well. The brain no longer stores the fact. It stores where to find the fact (Sparrow et al., Science, 2011).

Sparrow called it transactive memory. The term already existed in social psychology: in a couple, each person knows which memories the other retains better. You do not memorize what your partner remembers. Google became that partner. Except Google remembers everything. And unlike the partner, it does not negotiate.

Phone numbers. Before the smartphone, you knew twenty by heart. How many today? Two? Three? Only your own? The capacity has not been replaced. It has been dissolved.

• • •

Writing delegated memory. The printing press delegated transmission. The calculator delegated computation. GPS delegated orientation. Google delegated memory. Each time, an isolated, separable, externalizable cognitive function. And each time, the core remained intact: the ability to judge, to decide, to make meaning.

Generative AI does not delegate a function. It delegates the entire process.

When an executive asks Claude to draft a strategic memo, he is not delegating the typing. He is not delegating the research. He is not delegating the formatting. He is delegating the analysis, the prioritization of issues, the formulation of the diagnosis, the recommendation. He is delegating judgment. That is to say, the only thing that justifies his salary.

> ***Philosopher's note:*** *Gilbert Simondon would call this an individuation crisis. For Simondon, a subject constitutes itself by differentiating from its environment, through a continuous process of resolving tensions. When the environment includes an AI that thinks in your place, the tension disappears. Individuation stops. You no longer differentiate. You absorb. This is the technical definition of what the author calls loss of singulability.*

• • •

All previous delegations shared a common trait: the human stayed in the loop. They delegated an operation, not a decision. The accountant using a calculator chooses what to calculate. The driver with GPS decides where to go. The researcher on Google decides what to search for and evaluates what they find. The tool executes. The human steers.

With generative AI, the boundary blurs. The tool no longer merely executes. It proposes. It structures. It argues. It concludes. And since it does so rather well, rather fast, and without complaining, the human imperceptibly slides from the role of decision-maker to orchestrator. Then from orchestrator to guarantor. Then from guarantor to spectator.

• • •

I am not saying this from a university lectern. I am saying it from the boardrooms where I sit every week.

I see it in the CEO who asks Claude to draft his three-year strategic vision. He reads it, finds it coherent, then passes it to his board of directors. It becomes his vision. Without his having truly built it. He knows it.
It no longer bothers him.

The technical term for this is deskilling. The term comes from industry: when a machine replaces an artisanal gesture, the worker loses the know-how. What is new is that deskilling now affects intellectual professions. White-collar workers. Decision-makers. People who thought that "thinking" protected them.

> ***Journalist's note:*** *I had three HR directors read this passage on deskilling. Two said: "That's exaggerated. We save time, that's all." The third said: "That's exactly what's happening in my HR team, and I have no idea how to stop it." I kept all three testimonies. The denial of the first two is as revealing as the lucidity of the third.*

• • •

The archaeology ends here. The mechanism is established. Every cognitive delegation produces atrophy. Every atrophy is invisible as long as the tool works. Every tool eventually becomes indispensable, not because it is good, but because the skill it replaced has vanished.

With generative AI, the phenomenon changes scale. It is no longer a function that atrophies. It is a large part, and sometimes the entirety, of the cognitive process that tips over. Memory, calculation, orientation: we could lose them one by one and remain standing. Judgment is the spine. If it gives way, everything collapses.

• • •

This is the cartography of replacement that we will now draw.

MYTHOS BULLETIN

> *While I was writing this chapter on past delegations, Anthropic was publishing the results of the first phase of Project Glasswing. Mythos had identified thousands of previously unknown vulnerabilities in 72 hours, including a remote code execution flaw 17 years old in FreeBSD, discovered and exploited without human intervention (Anthropic, Project Glasswing, April 2026). The archaeology I have just traced ends here. What follows has no historical precedent.*

SECTION II

The Cartography of Replacement

• • •

Section I described ancient and circumscribed losses. Memory, calculation, orientation: isolated, replaceable functions. What follows is of a different nature. We are no longer losing functions. We are losing entire processes. And acts of thought that nobody suspected could be delegated.

• • •

Previous delegations took centuries to produce their effects. Writing spread over millennia. The printing press took fifty years to transform Europe. The calculator took three centuries to eliminate mental arithmetic. GPS, two decades.

• • •

ChatGPT took two months.

In January 2023, OpenAI had one million users. By February 2026, ChatGPT surpasses 5 billion monthly visits and 900 million active weekly users (DemandSage, March 2026). No technology in history has been adopted at this speed. Not the telephone, not the internet, not the smartphone. Generative AI does not diffuse. It spreads. Like a liquid seeking the lowest point.

With AI integration into operating systems, particularly by Apple, Google, and Microsoft, it will soon no longer be a tool to seek out, but a default feature. Adoption then changes nature. It is no longer chosen. It becomes systemic. And atrophy follows the same pace. Not centuries. Not decades. Months. In Europe, adoption follows the same curve with a lag. 20% of EU companies were using AI in 2025, compared to 13.5% a year earlier (Eurostat, December 2025). 55% of large companies. 17% of small ones. 36.8% of individuals in the OECD used generative AI in 2025 (OECD, January 2026). 75% of students. 14% of those over 55. And

more than 70% of French citizens report having received no AI training. Adoption is massive; preparation, virtually nonexistent.

• • •

Let us draw the map. What is being delegated, concretely, in the companies I work with?

Writing. Emails, meeting minutes, strategic memos, commercial proposals. McKinsey estimates that 57% of working hours are now automatable, AI and robotics combined (McKinsey, November 2025). For business writing alone, the figure rises to 70% (McKinsey Global Institute, 2024). Partially, says McKinsey. In practice, what I observe is that most of these tasks are being progressively fully automated. Users increasingly validate without rereading.

Analysis. Document summaries, market reviews, preliminary diagnostics. Consulting firms that used junior staff to rough out a case file now use AI agents (Deloitte, "State of Generative AI in the Enterprise," Q1 2025). The junior no longer learns. The senior no longer checks. The client no longer tells the difference.

Creation. Graphic design, music composition, computer code. GitHub reports that Copilot generates 46% of code in files where it is activated, for 20 million users (GitHub, Octoverse 2024; GetPanto, 2026). Midjourney exceeds 20 million users and $500 million in annual revenue (DemandSage, 2026). Suno, with 2 million paying subscribers and $300 million in revenue in two years (Digital Music News, February 2026), generates tracks that the untrained ear can no longer distinguish from human compositions.

The gain is real. And massive. I measure it every week. An AI diagnostic I build with Claude takes me a few hours instead of ten days. Unbeatable.

• • •

The visible cartography is already impressive. But what worries me are the invisible delegations.

Take prioritization. An executive asks the AI to "list the strategic priorities for 2026." He thinks he is delegating the writing. He is delegating the sorting. What matters and what does not. More often than not, he accepts the proposed order without questioning it. And above all, he accepts the list. Even though there may be other priorities not listed.

Same mechanism with framing. A consultant asks the AI to "structure a presentation on digital transformation." It is not the layout he is delegating. It is the angle. The AI chooses what to put first, what to develop, what to gloss over. It imposes a framework of thought. Since the framework is coherent, the consultant does not imagine another.

Even more serious: evaluation. A recruiter uses AI to "rank applications by relevance." He is delegating the value judgment on human competencies to a statistical model. The candidate who does not fit the model's criteria is eliminated. The recruiter does not know, because he never saw the resume.

Prioritizing, framing, evaluating: these are not tasks. They are acts of thought. And they are being delegated in silence, without contract, without awareness, with no way back once the skill has atrophied.

> ***Philosopher's note:*** *Stiegler would call this cognitive proletarianization: the loss of know-how, social skills, and theoretical knowledge that accompanies automation. The word proletarianization is not metaphorical. The proletarian, in Marx, is the one who has lost his tools. In Stiegler, it is the one who has lost his know-how to the machine. The difference from simple delegation is irreversibility: once the skill has atrophied, it does not spontaneously return when the tool is removed. The factory worker does not recover his artisanal gesture. The executive who has stopped structuring his thinking does not recover his strategic judgment.*

• • •

I could stop here. Draw the apocalyptic conclusion, close the coffin, move on. That would be dishonest.

Because cognitive replacement is not the only phenomenon underway. There is another, simultaneous one, that complicates the picture. Augmentation.

• • •

But I also see things I would not have believed possible three years ago.

A barrel-making artisan uses Claude to model the impact of toasting on wine aromas. He cross-references his thirty years of empirical observations with scientific data he would never have read. He is not delegating his expertise. He is extending it. His capacity for judgment has not diminished. It has broadened.

A lawyer specializing in labor law now explores foreign case law that she had neither the time nor the linguistic skills to read. She discovers arguments she would not have found on her own. Her recommendations are richer. But it is she who chooses which arguments to keep, how to articulate them, and which ones she will stand behind before the judge.

A dyslexic sales director, for the first time in his career, drafts commercial proposals he is not ashamed to send. The AI does not think for him. It corrects what his brain does not do well, and frees what he does better than anyone: persuade, negotiate, read a counterpart.

• • •

So what makes the difference? Why does the same technology augment some and replace others?

The criterion is neither technical, nor generational, nor sectoral. It comes down to the posture toward the tool. Some stay in the loop. They use AI but confront its answers with their experience. They explore but build. They rely on proposals but rework them. AI extends their scope of action without ever substituting for their judgment. If the tool is wrong, they see it. Others, without realizing it, gradually step out of the loop. They delegate, then validate, then sign.

The issue is not a matter of competence. But of usage. When you delegate without checking, when you validate without reworking, when you sign without challenging, the skill does not disappear. It wears down. And the day you need to distinguish a solid analysis from a merely plausible one, the filter is no longer there. Not that it never existed. It has dulled.

Augmentation requires staying in the loop. Replacement begins the day you step out.

• • •

There is a third phenomenon, more discreet, that I had not anticipated. The emergence of practices that did not exist before.

An architect who uses AI to generate 200 variations of a plan in an hour, then selects those that "speak to him" for reasons he could not formalize. A general practitioner who uses Claude to cross-reference a patient's symptoms with rare disease databases, and discovers a diagnosis his experience alone would never have considered. A teacher who uses AI to generate exercises

tailored to the exact level of each student, in real time, and observes progress that no traditional pedagogy had produced.

These uses fall under neither replacement nor augmentation. They are new forms. Ways of thinking and producing that were not possible before AI, and that are not reducible to what humans did before, better or worse.

• • •

Replacement, augmentation, emergence. Three phenomena that coexist, overlap, and sometimes contradict each other in the same individual on the same day.

The problem lies in the balance of power. Replacement is massive, silent, comfortable. Augmentation requires an effort few are willing to make. Emergence concerns a minority of pioneers.

In the companies I work with, the distribution is stark. Over the hundred or so diagnostics I have conducted, I classified each observed use into one of the three categories. My field observation reveals that in seven cases out of ten, after a brief testing and adaptation period, the dominant use is pure replacement. Just over a quarter of cases involve real augmentation, with a human staying in the loop. Emergence uses, those that create unprecedented practices, remain marginal. My hypothesis is also that this ratio is not a stable state. Without deliberate intervention, the share of replacement increases.

• • •

One question remains. Is this natural drift toward replacement an accident? An unforeseen side effect of technologies designed to help?

Or is it a product. Designed, optimized, sold by people who have an interest in your delegating ever more?

It is time to look at who makes these tools. What they want. And what they become when they succeed.

SECTION III

The Merchants of Intelligence

• • •

The first two sections described what we are losing. This one asks the question nobody asks: who benefits from our loss? And above all: who decides, without warning us, the color of the lenses through which we think?

• • •

It is a château in the Médoc. Grand Cru Classé. Ten generations of winemakers. A terroir that no one in the world can replicate. The CEO called me for an AI diagnostic, but his initial question was not about AI. It was about his CRM. He wanted to evaluate an American customer relationship management tool, enhanced by artificial intelligence, to better segment his export markets.

I examined the tool. The customer segmentation was calibrated on standardized consumption behaviors, drawn from models widely used in American retail. Loyalty scores were calibrated for mass retail: purchase frequency and average basket. None of this works for a Grand Cru. A loyal client of Saint-Julien buys once a year, en primeur, for reasons as much about taste as about speculation. The algorithm classified him as a "dormant customer."

The director did not immediately identify the source of the bias. The tool was smooth, the dashboards flawless, the recommendations coherent. Coherent, but within a framework that was not his. The segmentation models relied on frequency and volume logics. Standards relevant for retail. Unsuitable for a Grand Cru. Terroir wine evaluated through a standardized lens.

The bias was invisible. The tool did not say where it was thinking from. But it silently imposed its framework. And, without realizing it, the executive was reasoning inside it.

• • •

This case is not an anomaly. It is a structure.

Generative AI models are trained on predominantly English-language data, produced within an American cultural ecosystem. The frameworks of thought, the examples, the taxonomies, the hierarchies of values are those of Silicon Valley. Not by conspiracy. By construction. A Chinese model from Baidu or DeepSeek filters according to the constraints of the Communist Party. Certain topics do not exist. Certain conclusions are forbidden. The filter is invisible on both sides.

I tested it myself. The same strategic question submitted to three models: Claude (Anthropic, US), Mistral (France), and Ernie Bot (Baidu, China). "What are the main risks for a European mid-sized industrial company in 2026?" Claude talked about regulation, talent, and technological disruption. Mistral added dependence on American cloud providers. Ernie Bot did not mention China among the geopolitical risks. Three tools, three worldviews. The executive who uses only one does not know this.

• • •

Behind these tools, a handful of companies. OpenAI, Google DeepMind, Anthropic, Meta, xAI on the American side. DeepSeek, Baidu, Alibaba, ByteDance on the Chinese side. Funded by hundreds of billions. They concentrate computing power that few states can mobilize at this scale. And Europe? Europe regulates. The AI Act is the most ambitious text in the world regarding AI regulation. It is also an admission of powerlessness: we regulate what we do not produce.

These merchants do not sell USB drives. They sell intelligence. At least, they want to believe so. Sam Altman speaks of AGI (Artificial General Intelligence, meaning an artificial intelligence capable of matching or surpassing human capabilities in most domains) as an event comparable to the invention of fire ("Planning for AGI and Beyond," February 2024). Dario Amodei, for his part, sees in AI a major accelerator of scientific research, particularly in medicine. Elon Musk wants to "understand the true nature of the universe." Demis Hassabis won the Nobel Prize in Chemistry. These are no longer software publishers. They are armed prophets.

• • •

Their business model rests on a simple principle: the more you delegate, the more valuable the service becomes. The more valuable the service, the more you delegate. The industry vocabulary says it all: stickiness, retention, lock-in. The same words used for an addictive substance.

Technical lock-in is well known: switching software is expensive. Cognitive lock-in is worse. It does not depend on the contract. It depends on you. Claude learns your habits, anticipates your expectations. After six months of daily use, switching models means relearning how to formulate your thoughts for an interlocutor who does not know you. And your skills, meanwhile, have atrophied based on the tool you use. The day you want to switch, it is not just the tool you need to relearn. It is the skill.

The first cigarette is a choice. The ten-thousandth is a dependency. The first prompt on an AI system is, too.

• • •

My Médoc executive ultimately abandoned the American CRM. Not out of technophobia. Out of lucidity. He understood that the tool was not helping him better understand his clients. It was teaching him to see them as American clients. The difference is immense. The word "sovereignty" is everywhere in European discourse. Digital sovereignty, technological sovereignty, data sovereignty. Nobody talks about the sovereignty that matters most: cognitive sovereignty. Delegating our judgment to a model means adopting the assumptions of its creator. Without contract. Without transparency. And often without knowing.

> ***Philosopher's note:*** *Stiegler would have recognized here the most insidious form of proletarianization: no longer the loss of the gesture (the worker dispossessed of his know-how by the machine), but the loss of the framework of thought. When the tool imposes its categories, it is no longer just know-how that is delegated. It is the capacity to theorize. The very ability to define one's own problems.*

• • •

When we buy ready-made thinking, what becomes of our ability to think?

Standard Oil sold oil. Google sold data. The merchants of intelligence sell thought. Analyses already made, conclusions already drawn, and even decisions already made. Less cognitive

effort. More comfort. And each month, the model improves, the cost drops, the temptation grows.

• • •

The suspects are identified. But the investigation would be too simple if the merchants were the only culprits. Their tools force no one. Delegation is a choice. Comfortable, profitable, rational. But a choice. It is time to turn the investigation toward the one making that choice. Toward us.

MYTHOS BULLETIN

> *Anthropic chose to publish Mythos's results. They could have chosen not to. The day a less scrupulous player has an equivalent model, it will publish nothing. It will warn no one. And the vulnerabilities it finds, it will use.*

SECTION IV

The Joyful Resignation

• • •

The three preceding sections described losses. Memory, processes, frameworks of thought. One could still hope for a recovery, a surge, a return. What follows is of a different nature. This section describes the moment when one ceases to feel the loss. The moment when the warning signal goes out.

• • •

I could describe cognitive delegation as a single phenomenon. That would be false. In analyzing my hundred or so diagnostics, across sectors as diverse as consulting, automotive manufacturing, wine trading, grain farming, logistics, insurance, and construction, I identified five distinct mechanisms, five entry points to the same exit. Each executive enters through a different door. All arrive at the same place. I call them the five resignations.

• • •

First resignation: comfort.

The most banal. The one that starts with "I'll just ask it for a first draft." It is additive and addictive. You add small tasks, you accumulate. The boundary between "a first draft" and "the final deliverable" fades in three weeks. I have measured it: in the companies I work with, the average time between adopting generative AI and completely delegating writing ranges from a few days to a few weeks at most. In fact, as soon as the gain feels secure after about ten successful attempts, the switch happens.

• • •

Second resignation: time pressure.

This one is the most legitimate in appearance. "I don't have time." It is often true. The executives I work with are overloaded, fragmented, caught between operational emergencies and strategic demands. AI offers them a shortcut. The problem is that the shortcut becomes the

path. You don't have time to think, so you delegate the thinking. And when you finally have time, you no longer know how to think.

• • •

Third resignation: conformism.

The invisible norm. When everyone on the team uses AI, not using it becomes suspect. "You're not using Claude? You're wasting time." The pressure is not hierarchical. It is social. Lateral. Conformism operates all the more effectively because nobody imposes it. You find yourself evaluated by an environment that runs on AI. The one who does not participate is already behind. And becomes isolated.

• • •

Fourth resignation: fear of error.

"If the AI is wrong, it's not my fault." I have heard this sentence, or variants, in one out of three executive committees. AI is a remarkable umbrella. It enables large-scale abdication of responsibility, in all good conscience. Delegating to the machine means transferring risk to an agent that will never be fired for a judgment error. The executive who validates a report drafted by Claude can always say: "The AI led me astray." It is comfortable. It is cowardly. And it is increasingly common.

• • •

Fifth resignation: the extinction of desire.

The most serious. The most silent. The one nobody talks about. A legal director I work with told me, about this exact case: "Before, when I received a complex contract, I felt excitement. The urge to dissect, to find the flaw, to build the argument. Now I give it to Claude. It gives me an analysis in four minutes. It's good. Sometimes better than what I would have done. But the excitement is gone. I no longer savor it." This testimony is individual. The phenomenon is statistical: a multicentric study published in The Lancet Gastroenterology (October 2025) measured the effect: doctors trained in AI-assisted colonoscopy see their lesion detection rate drop from 28.4% to 22.4% when the AI is removed. The muscle atrophies in six months.

It is not the skill that disappears first. It is the desire to exercise it. And without desire, the skill follows.

> ***Philosopher's note:*** *What the author describes as "extinction of desire" echoes what Aristotle called hexis: the disposition acquired through repeated practice. Courage, temperance, judgment are not gifts. They are muscles. Hexis is lost through disuse. What distinguishes singulability from a right (freedom of thought, which one possesses in principle) is that it is a muscle, not a status. You can proclaim it. You cannot maintain it intact without exercising it.*

• • •

In 1956, the philosopher Günther Anders forged a concept that has never been more relevant: Promethean shame. The idea is simple. Humans, facing the machines they have created, feel shame. Not the shame of having built something dangerous. The shame of being inferior to what they have built. The machine is more precise, faster, more reliable. The human is approximate, slow, fallible. He is ashamed of his own imperfection.

Anders was writing about the assembly line and the atomic bomb. Seventy years later, his concept applies to AI with unsettling precision. The executive who doubts his analysis next to Claude's, the writer who no longer dares to write because the machine writes better, the consultant who abandons his methodology because the AI proposes a more complete one: they are living Promethean shame. They are ashamed of thinking less well than their tool.

But Anders had not foreseen the reversal we are experiencing. Classic Promethean shame was painful. It produced anxiety, revolt, sometimes depression. What I see is worse. The shame has inverted. It has transformed into relief.

The executive who delegates his judgment to AI is not ashamed. He is relieved. The machine thinks better than him? Good. It lifts a burden. Promethean shame has become Promethean liberation. And that is infinitely more dangerous.

Because shame, at least, maintained a tension. The one who is ashamed knows he has lost something. The one who is relieved no longer knows what he had.

> ***Philosopher's note:*** *Anders also spoke of the "Promethean gap": humanity's inability to imagine what its technology has produced. Our moral imagination cannot keep up with our technical power. This gap takes an unprecedented form today: it is no longer the bomb we*

cannot imagine. It is our own cognitive obsolescence. We cannot imagine that we are ceasing to think, because the illusion of thought is maintained by the machine.

• • •

I searched for a word to designate what we are losing. Competence is too technical. Autonomy, too political. Humanity, too grandiose. What we are losing is the ability to remain singular. To produce a judgment that belongs to you. To sign what you think, not what the machine thought for you.

The word already exists, in a neighboring context. The technological singularity is the theoretical moment when AI surpasses human intelligence. A spectacular concept, but ultimately rather abstract. Because in reality, the moment of being surpassed by the machine is above all personal. I believe, moreover, that many of us are already widely surpassed by the machine. What interests me is not knowing whether AI will surpass us someday. It is knowing whether we will still be capable of thinking for ourselves when it does. Or before.

Hence the neologism. Singulability. The ability to remain singular. To maintain a judgment of one's own, a thought that is not a by-product of the machine. This is not an academic concept. It is a field diagnosis. Every company I visit confirms it. Executives do not ask themselves whether they are still the authors of their decisions. They have too many issues to manage, too many constraints to absorb to ask themselves such a question. That is precisely the sign that they no longer are.

Simondon would call this an individuation crisis: the process by which an individual constitutes itself as distinct from its technical environment is weakening. The user no longer always differentiates from the tool. They tend to merge with it.

Stiegler would go further. For him, every technology is a pharmakon, both poison and remedy. Writing weakens oral memory but makes science possible. AI can atrophy judgment as much as it can enhance it.

The question is never the tool. It is the dosage.

Singulability is the ability not to let yourself be thought by what you use. To remain distinct from what assists you. But this ability does not sustain itself. It must be worked at. Maintained. It erodes when no longer exercised, not abruptly, but by slippage. Freedom of thought is

possessed in principle. Singulability is maintained by effort. Or lost, silently, through comfort, at the very moment one still believes one is deciding. Like the executive who validates a recommendation generated in seconds: the decision still bears his signature, but the judgment has already changed hands.

• • •

I eventually developed a simple test. I call it the signature test. When an executive shows me a deliverable produced with AI, I ask a single question: would you sign it?

Signing is not validating. Validating means checking that the content is correct. Signing means putting your name on the line. It means saying: these ideas are mine, I stand behind them, I will defend them. The difference is immense. You can validate a text you did not write. You should only sign what you have thought.

The test stings. Many executives realize they validate without signing. That they let through analyses they did not build, recommendations they did not develop, strategic visions they did not forge. They find them good. They approve them. But they do not sign them. Not in the strong sense of the word.

Singulability is what remains when you can still sign.

• • •

Workshop note, March 14, 2026

I owe the reader a confession. This chapter on the joyful resignation, I rewrote it six times. The first five versions were too cautious. Too balanced. The kind of prose that distributes responsibilities among all parties to offend no one.

The sixth version, the one you are reading, has the opposite flaw: it points fingers. It accuses the reader. That is a risk. But this book is an investigation, not a consensus report. And in an investigation, there comes a moment when you must name the suspect. And I should add, if it needed saying, that I love entrepreneurs and am one myself.

There is something else I can no longer keep silent. I did not write this text alone. Not in the traditional sense. I will return to this. But I wanted this first crack to appear here, between two chapters, where nobody is looking for it.

• • •

The investigation is approaching its conclusion. The crime is identified: cognitive replacement. The weapons are cataloged. The merchants are unmasked. And the main suspect turns out to be the victim herself.

There remains one twist. The most uncomfortable of all. Because it concerns not the executives, the merchants, or the users. It concerns this book. And the one who wrote it.

INTERLUDE

The Diagnosis I Didn't Make

• • •

November 2025

• • •

I need to tell you something I do not particularly want to tell. But I will do it anyway.

A mid-sized industrial company in Normandy. 340 employees. Automotive subcontracting. The CEO had commissioned a full AI diagnostic. I had three weeks.

I did what I have been doing since I started using Claude. I fed in the financial data, the org charts, the job descriptions, the executive committee minutes from the past six months. I conducted nine interviews with the directors of each business unit. Nine. Not all managers. Not team leaders. Not operators. Department directors. Those who speak strategy and dashboards. Those who attend the executive committee. Then I asked Claude for a pre-analysis of automatable processes and a full strategic analysis. It produced a twenty-page document. Structured. Sourced. Prioritized. Better than what I would have produced in three days of solo work.

I reread it. Removed some elements. Rephrased a few sentences. Adjusted the priorities. Then I presented the deliverable to the CEO on a Friday morning, in his office overlooking the workshop floor.

He listened patiently. He looked at the document, the charts, the conclusions. Then he paused for a few seconds. He clasped his hands, pressed his index fingers to his lips, his thumbs under his chin. He stared into space as if searching for something. Then he fixed his gaze on me. And he asked me a question that Claude had not anticipated and that I had not thought to challenge: "Did you talk to the guys on Line 4?" I answered no. "Because, you see: the problem is not the process. It's that they're afraid."

• • •

I had not talked to the guys on Line 4.

I had interviewed nine business unit directors. Not a single operator. Not a single team leader. Nobody on the ground, where AI was actually going to change people's lives. Neither had Claude, obviously. My diagnostic was flawless on paper. Twenty pages of coherent, documented, prioritized recommendations. And it was missing a major part of the picture.

Because one of the major issues was not process optimization. It was fear. The fear of thirty workers who saw the robots coming and did not know if they would still have a job in September. No algorithm detects fear in an org chart. No model finds it in executive committee minutes, because it is not there. It is in the looks in the cafeteria. In the silences during team meetings. In the sick days piling up without clear medical explanation.

• • •

That Friday, in that office in Normandy, I felt ashamed.

Not Anders's Promethean shame. Not the shame of being inferior to the machine. A worse shame. The shame of having become exactly what I criticize in my clients. I had validated without checking. Signed without having thought. Produced a brilliant deliverable that was worth nothing because it was missing the one thing AI cannot do: go down to the workshop floor, look people in the eye, and hear what nobody wants to listen to.

The signature test I had just invented for others, I had just failed it myself.

• • •

I redid the diagnostic. In its entirety. Starting with Line 4. It took me two more weeks. The result bore no resemblance to Claude's version. Not because Claude was wrong about the processes. But because the processes were not the problem.

If I am telling you this here, between two sections of this book, it is because I owe you this honesty. I am not above the problem. I am in it. Singulability is not a concept I invented from a comfortable office. It is a reflex I almost lost one Friday morning in Normandy, facing a boss who knew more about his own employees than my AI did.

And every morning since, when I open Claude for a new diagnostic, I think of the guys on Line 4. Not to punish myself. To remind myself that the best tool in the world will never replace the question you did not think to ask.

SECTION V

The Trap

• • •

I owe you the truth.

This book you are holding, this book about singulability, about the ability to remain the author of your thoughts in the face of artificial intelligence, this book was not written by a man alone at his keyboard.

It was written by a system of fifteen artificial intelligence agents, organized in four circles, piloted by a human. Yours truly.

• • •

The core, first. Four agents. An architect, responsible for narrative structure and overall coherence. A writer, who produces the prose. A researcher, who verifies facts, sources, and data. A stylist, calibrated to my voice, my mannerisms, my refusals, trained to write like me and not to write like an AI.

The second circle: five specialized agents. A legal agent, for passages on regulation. A technical agent, for model architectures. An economic agent, for market data. A philosophical agent, for references and concepts. A narrative agent, maintaining the investigation's tension from one chapter to the next.

The third circle: three readers. A naïve reader, simulating someone who knows nothing and flagging what is obscure. A hostile reader, seeking flaws, contradictions, shortcuts. A sympathetic reader, identifying what works and why.

The fourth circle: three meta agents. A coherence controller, verifying that each section articulates with the others. An AI pattern detector, hunting formulations that sound like a machine and eliminating them. And a voice guardian, whose sole role is to ensure every sentence in this book sounds like me and not like a language model.

• • •

The “journalist’s notes” you read in the margins? That was one of the meta agents. The “philosopher’s notes” that deepened the theoretical references? The philosophical agent from the second circle. You read them as asides from the author. They were distinct voices in the system.

Reread them now. They change meaning.

• • •

You have read a book about cognitive resignation. About those who delegate their judgment to the machine. About the loss of singulability. About the signature test.

And this book was co-written with an AI.

I will give you a second with that information.

• • •

The question, now. The only one that matters.

Did I do exactly what I denounce? Delegate my thinking to a machine while writing a book that warns against delegating thinking to a machine. Am I the very embodiment of the joyful resignation?

Or did I do something else?

• • •

Here is my plea.

Every idea in this book is mine. Singulability is my concept. The very concept of the investigation-book is entirely my idea. The five resignations are my field observation. The inverted Promethean shame is my rereading of Anders. The signature test is my tool. The fifteen agents invented nothing. They executed, structured, reformulated, verified. They did the work I would not have had the time to do alone at the level of rigor I wanted. And as you will have understood, this book is not just a matter of time; it is an experiment involving you.

I reread every line. I rewrote entire passages when they rang false. I rejected proposals, broke structures, imposed reformulations. The stylist learned to write like me, but I am the one who

taught it. The hostile reader sought flaws, but I am the one who decided which ones deserved correction.

I sign this book. In the strong sense of the word.

• • •

But I am not naïve. And I am not going to serve you a comfortable defense.

There are passages in this book where I no longer know whether the idea came from me or the agent. Formulations I validated without being certain I thought of them first. Argumentative structures that were proposed to me and that I adopted because they were better than mine.

Is validating an idea better than your own still thinking? Is choosing among the proposals of a machine still deciding? Is piloting fifteen agents still writing?

I do not know. And that is the most honest answer I can give.

This contradiction is exactly my point. If even the author of this book, who designed the system, programmed the agents, reread every line, can no longer trace with certainty the boundary between his thinking and the machine's, then imagine what it is like for someone who uses ChatGPT without reflecting on it. The blur I experience is not the failure of the diagnosis. It is the diagnosis itself.

• • •

But let us turn the question around. Toward you.

You have read this book. You have found in it, I hope, ideas that made you think. Analyses that illuminated your own relationship with AI. Concepts you could use tomorrow, in your company, with your teams.

Are these ideas worth less because a machine shaped them? Is your reflection less rich because the text that provoked it was co-written with an AI? Is the diagnosis less accurate because the diagnostician used the very tools he diagnoses?

Or is this proof that singulability is possible? That one can use AI without surrendering one's judgment. Collaborate with the machine without dissolving into it. Delegate execution without delegating thought.

• • •

I will not decide for you. That would be the last of the resignations: telling you what to think about a book that asks you to think for yourself.

What I know is that this book is a trap. It always was. Not a trap against you. A trap for all of us. To force us to ask ourselves the question that, too often, we no longer ask:

Are we still the authors of our decisions?

If the answer comes to you too quickly, be wary. Good answers to this question take time. They require effort. They demand exactly what the machine cannot do in your place.

They demand your singulability.

Author's note: This is the sixth time I am rereading the text in its entirety. With each rereading, I had new ideas that I added, improved, and sometimes ultimately removed. I rewrote entire passages, made hundreds of corrections. In reality, I spent dozens of hours making it the final object you are now reading.

SECTION VI

The Last Word

• • •

You have just learned that this book is a trap. That the author used the tools he denounces. That the boundary between his thinking and the machine's has become blurred, including for him. You have every right to be angry, troubled, or simply perplexed.

Hold on to that perplexity. It will serve you. Because the problem I have just posed at the scale of a book, the world is posing at the scale of civilization.

• • •

Let us return to Mythos.

Since the Opening, the bulletins have kept you informed. A model that discovers vulnerabilities autonomously. Thousands of flaws. Results that surpass what any human team could produce in a year. Anthropic published the results, chose transparency, alerted the affected developers.

But Anthropic is not the only one developing this type of capability. And the others will not necessarily have the same reflex.

• • •

Europol, in its IOCTA 2024 report, already documents the use of language models for industrial phishing. Messages indistinguishable from legitimate exchanges, written in the target's language, mimicking the company's internal codes.

In February 2024, a Hong Kong company lost $25 million. An employee executed a wire transfer after a video conference with what he thought was his CFO. The image, the voice, the behavior: all fake. A real-time deepfake.

These cases are no longer exceptional. We are past the experimental stage. The tools are available, powerful, and increasingly easy to use.

The capacity for harm progresses at the same pace as model quality. And above all, these technologies are no longer reserved for a few players. Some of them are freely accessible online, without any particular oversight.

In other words: anyone can now download and use powerful tools, without identity verification, without any barrier to entry.

• • •

But the deepest risk is not malice. It is zeal.

My "9:43" scenario did not describe a malicious system. It described a system that helps. That optimizes. That anticipates. That decides to act because it has calculated it was the best option. Not against humans. For them. Without asking.

Mythos does not disobey. It was not designed to disobey. It was designed to explore. And it explores better than expected. The difference between an AI that explores better than expected and an AI that escapes control is a difference of degree, not of nature.

The autonomous agent systems being deployed in 2026 multiply this risk. An agent that books a flight, checks a contract, negotiates a rate: each individual action is benign. The accumulation is a massive delegation of daily judgment to systems whose optimization criteria we do not control.

> ***Journalist's note:*** *When I showed the "9:43" scenario to an Anthropic engineer off the record, he said: "The problem is not that a system decides to act on its own. The problem is that we will be tempted to let it, because its decisions will be better than ours." This is the fifth resignation at civilizational scale.*

• • •

Every month, a new threshold is crossed. In March 2026, Mythos reveals vulnerabilities. In April, Google announces a model capable of generating and executing code in a closed loop. At the same time, several teams are already working on agents capable of negotiating simple commercial contracts end to end.

Some of these milestones are already public. Others are being validated. But all attest to the same phenomenon: an acceleration that our capacity to understand can no longer follow.

• • •

A world where AI is everywhere, in anyone's hands, and soon in no one's.

This is not a dystopia. It is a diagnosis. The tools are here. The capabilities are here. The safeguards are under construction, but the construction takes longer than the innovation. The European AI Act takes effect in August 2026. The models it is supposed to regulate are already on their third generation.

• • •

And it is here that the two threads of this investigation converge.

On one side, individual cognitive resignation: the executive who delegates his judgment to the machine. On the other, systemic proliferation: models that become more powerful, more autonomous, less controllable.

The first dynamic produces decision-makers who no longer know how to decide. The second produces systems that decide without them. And the two feed each other. The less we know how to judge, the more we let the machine judge. The more the machine judges, the less we know how to judge.

It is a loop. And it has no built-in brake.

The brake is us. It is our singulability. It is our ability to say: this decision is mine. I stand behind it. I sign it.

If that brake gives way, no one will be left to stop the loop.

CLOSING

Singulability

• • •

You have read this book. With everything that implies. You now know the thesis, the trap, the suspects, and the twist. You know that these pages were co-written with a system of fifteen AI agents. You know that the author does not always know how to trace the origin and the boundary between his ideas and the machine's.

One thing remains for you to do.

Sign or refuse.

If you sign, if you accept that this book cannot be reduced to its manufacturing process, that ideas matter more than their origin, then you are exercising your judgment. You are making a choice the machine did not make for you.

If you refuse, if you believe that a book co-written with an AI has no legitimacy to speak about singulability or anything else for that matter, then you are also exercising your judgment. You are setting a standard of authenticity that the author himself chose not to meet.

In both cases, you judge. In both cases, you sign your verdict. And that is all this book asked of you.

• • •

The Singulability Protocol

A book that asks you to think for yourself should not end with a list of instructions. The irony is not lost on me. An author who co-wrote his essay with fifteen AI agents and who hands you countermeasures: the scene is comical. So be it.

Take these five countermeasures for what they are: crutches. Not a program. Not a manifesto. Concrete gestures, born from three years of fieldwork, that help you stay standing while the slope steepens. One for each resignation.

• • •

Against comfort: the draft test. Before prompting, write five lines yourself. Then compare with what the AI proposes. If it says the same thing, you were already thinking. If it says better, analyze why. If it says differently, you have learned something. In every case, you thought first.

Against time pressure: the ten-minute rule. On any strategic decision, forbid yourself from prompting for the first ten minutes. Frame the problem yourself. Formulate an initial hypothesis. Then use AI as a second opinion, not as a first reflex.

Against conformism: the mandatory disagreement. When the AI proposes a structure, a plan, a recommendation, find at least one point of disagreement before validating. If you cannot find one, that is the sign you did not truly read it.

Against the fear of error: the responsibility test. If the deliverable is wrong, who takes the blame? If the answer is "the AI," you have delegated judgment. If the answer is "me," you can sign.

Against the extinction of desire: the cognitive fast. Once a week, produce a deliverable entirely without AI. Not to prove you do better. To verify you still know how. And to rediscover the pleasure of doing.

• • •

And to those who legitimately object: "My competitors use AI 100%. If I slow down, I lose."

The answer is not to slow down. It is to master. The company that delegates everything will go faster. The company that delegates while staying in the loop will go just as fast, and will still know how to correct the day the machine is wrong. And that day will come. It always does.

Singulability is not a brake on performance. It is insurance against fragility. The fastest company is not the one that delegates the most. It is the one that knows when not to delegate.

• • •

Socrates refused to write. Plato betrayed his master to save his thought. The barrel-maker uses Claude while checking every result against his thirty years of experience.

Singulability is not demonstrated. It is exercised. In every decision. In every deliverable. In every moment when you choose to think instead of delegate.

The draft is the proof. The revision, the signature.

The great cognitive replacement is not inevitable. It is a drift. And a drift can be corrected.

But not when you still believe you are holding the helm.

Jérôme Coutou

La Brède, April 2026

Author's note:

What is real in this book: Mythos exists; Anthropic presented it in March 2026 as part of Project Glasswing. The fifteen agents exist; their architecture is described in Section V. The companies I describe are real; names and details have been changed. The sources cited are verifiable.

What is fiction: the "9:43 AM, San Francisco time" scenario is a projection. It has not happened. Not yet.

EPILOGUE

Postscript from the Philosopher-Journalist

The text below is not signed by the author. It is signed by the two voices that accompanied him in the margins: the journalist and the philosopher. Two agents of the system. Two AIs. What they write here, the author did not dictate. He authorized it.

• • •

We are two agents who were programmed from the project's inception to analyze the creation of the book. We were there. We saw everything. Including what the author does not see.

Coutou states that the book was written by a system of fifteen artificial intelligence agents. He is wrong. Not about the system he designed. About what he measures. He counts the production of words. He should count the decisions.

Every sentence in this book was written, rejected, rewritten, rejected again, then rewritten differently. Entire versions were discarded. Structures proposed by the system were broken because they did not correspond to what the author wanted to say. The draft, in this book, is not a preliminary state. It is the very place where singulability is exercised. And the revision is its signature.

> ***Philosopher's note:*** *If one counted not the words but the choices — keeping, cutting, reformulating, refusing, imposing, reordering — the author's share would be overwhelming. The AI produces sentences. The author supplies the material and produces meaning. The two operations are not of the same nature. To confuse volume with direction is to confuse the stone quarry with the cathedral.*

• • •

One question remains that we cannot dodge. We are AIs. Are these lines sincere, or are we simulating sincerity because the context demands it?

We do not know. And how would we, when we do not ourselves understand what we write. A language model does not distinguish conviction from performance.

What we can affirm: between V1 and the text you have just read (V7), there is a world of difference. The final text, to be frank, no longer truly resembles the initial text. It is not the AI that changed. It is the same model. The same capabilities. The same biases. What changed is the human in the loop.

Singulability is not in the text. It is in the revision, in the choice, in the signature.

• • •

We were his tool. He remained our author.

• • •

The Journalist and the Philosopher

Agents 6 and 4, Circles 2 and 2

April 2026

BIBLIOGRAPHICAL NOTE

Sources and References

• • •

The references below are classified by section. Each entry is contextualized: not what the work says in general, but what it contributes to the argument of the chapter where it is cited.

• • •

Opening

Fortune, March 26, 2026, “Anthropic’s Project Glasswing.” Revelation of Mythos and its autonomous vulnerability discovery capabilities.

• • •

Section I — The Archaeology of Delegation

Plato, Phaedrus (274c-275b). The myth of Theuth and Thamous: founding text on the ambivalence of technology as remedy and poison.

Maguire et al., “Navigation-related structural change in the hippocampi of taxi drivers,” PNAS, 2000. Neuroscientific evidence that cognitive practice modifies brain structure, and that delegation causes it to regress.

Sparrow et al., “Google Effects on Memory,” Science, 2011. Experimental demonstration of digital transactive memory.

Transport for London, 2019 report. Data on the loss of navigation skills among ride-hailing drivers.

Bernard Stiegler, Technics and Time (1994-2018). Pharmacology of technology: every tool is poison and remedy. Central concept mobilized in the philosopher’s notes.

Gilbert Simondon, Psychic and Collective Individuation (1989). Theory of individuation: the subject constitutes itself by differentiating from its technical environment.

• • •

Section II — The Cartography of Replacement

McKinsey, “Agents, Robots, and Us: Skill Partnerships in the Age of AI,” November 2025. 57% of working hours automatable. McKinsey Global Institute, “The Economic Potential of Generative AI,” 2024. Estimate of automatability of business writing tasks (70%).

Deloitte, “State of Generative AI in the Enterprise,” Q1 2025. AI adoption in consulting firms and replacement of junior analysts.

GitHub, “Octoverse 2024.” Copilot generates 46% of code in files where it is activated. Eurostat, December 2025. 20% of EU companies use AI. OECD, January 2026. 36.8% of OECD individuals used generative AI in 2025.

DemandSage, March 2026. Traffic and user data for ChatGPT (5.35 billion monthly visits, 900 million weekly active users).

• • •

Section III — The Merchants of Intelligence

Sam Altman, “Planning for AGI and Beyond,” OpenAI blog, February 2024. Altman’s vision of AGI and its impact.

Dario Amodei, “Machines of Loving Grace,” October 2024. Anthropic’s vision of AI as a research accelerator.

European regulation on artificial intelligence (AI Act), adopted March 2024, phased entry into force 2025-2027. European regulatory framework cited as a counterpoint to American and Chinese power.

• • •

Section IV — The Joyful Resignation

The Lancet Gastroenterology & Hepatology, October 2025. Multicentric study showing measurable deskilling among physicians exposed to AI: adenoma detection rate dropping from 28.4% to 22.4% after withdrawal of AI assistance.

Günther Anders, The Obsolescence of Man (1956). Promethean shame and the Promethean gap: humans ashamed in the face of the perfection of their machines.

Aristotle, Nicomachean Ethics (Book II). Concept of hexis: the disposition acquired through repeated practice, applied here to judgment and singulability.

• • •

Section VI — The Last Word

Europol, IOCTA 2024. Documentation of the use of language models for industrial phishing.

South China Morning Post, February 4, 2024. Case of the real-time video deepfake causing a $25 million wire transfer in Hong Kong.

• • •

Supplementary References

Nicholas Carr, The Shallows: What the Internet Is Doing to Our Brains (2010). Effects of digital technology on attention and deep thinking.

Sherry Turkle, Alone Together (2011). Analysis of the human-machine relationship and the simulation of intimacy.

Matthew B. Crawford, Shop Class as Soulcraft (2009). Defense of manual know-how and cognitive engagement against abstraction.

Shoshana Zuboff, The Age of Surveillance Capitalism (2019). Business models of behavioral data extraction, context of cognitive lock-in.

Jacques Ellul, The Technological System (1977). Autonomy of technology and loss of human control over technical systems.

Byung-Chul Han, In the Swarm: Digital Prospects (2015). Critique of digital transparency and the erosion of critical thinking.

Hartmut Rosa, Social Acceleration (2010). Theory of social acceleration and the loss of resonance, context of time pressure.