# Daily use, shallow fluency.

*Why Artificial Intelligence still shows up more in conversation than in results inside Brazilian companies.*

BY **BERNARDO PRECHT** AND **FELIPE LOURENÇO**

**SCOPE**

50+ hours of interviews with HR directors, C-levels and Stanford MBAs, cross-referenced with 19 secondary sources and three anonymized Brazilian company portraits. Details in chapter 09.

**ANALYTICAL FOCUS**

A dual lens: the workforce now entering the market, and the leaders who recruit, manage and make decisions about it.

**AUDIENCE**

C-levels, board members, HR leaders, HR Tech investors.

CHAPTER 01

# Executive summary

*Four findings that frame the rest of the report*

---

The right question for 2026 is no longer whether AI is changing work in Brazil. It is *where* it has changed, and where it has not yet arrived. The answers, after eight months of research, are uncomfortable for anyone expecting linear transformation.

The report triangulates in-depth interviews (Brazilian HR managers and directors, C-levels and Stanford MBAs) with secondary sources from the Stanford Digital Economy Lab, McKinsey, Anthropic, PwC, BCG, Economist Impact, INSEAD, OECD and Drake Star. The framing is intentional. The MBAs show the workforce now entering the market; the leaders interviewed show the people recruiting, training and making decisions about that same workforce. What emerges is a country with adoption and anxiety in equal measure, and a management layer split between enthusiasm for pilots and the inertia of existing processes. Twelve numbers describe the landscape. Four findings frame the report.

# The state in 12 numbers

**4×**

Demand for AI skills quadrupled in Brazil over three years. From 19k postings (2021) to 73k (2024).

PWC AI JOBS BAROMETER BRAZIL

**−14%**

Decline in the *job-finding rate* for young workers aged 22 to 25 in AI-exposed occupations, post-ChatGPT.

ANTHROPIC NOWCASTING (2026)

**33%**

Only one third of theoretically automatable Computer & Math tasks are covered by real-world use today

ANTHROPIC ECONOMIC INDEX

**+49pp**

Performance gain for non-technical consultants using GenAI on a *data cleaning* task

BCG / BU / OPENAI

**56%**

Global wage premium for professionals with AI skills. The premium more than doubled in a year.

PWC AI JOBS BAROMETER GLOBAL

**94%**

Routine HR tasks automated at a *big tech* cited as a global benchmark

ECONOMIST IMPACT, "AI GLASS FLOOR"

**7%**

Mid-sized and large manufacturing and ICT companies in São Paulo that use any AI application

OECD / BCG / INSEAD · SÃO PAULO

**67%**

Brazilian public servants say their AI knowledge is entirely or mostly self-taught

PUBLIC SECTOR AI ADOPTION INDEX 2026

**17%**

São Paulo companies using AI in the survey that created a senior role or team dedicated to the technology

OECD / BCG / INSEAD · SÃO PAULO

**53h30**

Weekly time consuming online media in Brazil, one of the highest levels in the global sample

DATAREPORTAL / GWI · DIGITAL 2026

**98.4%**

Brazilian internet users who access the internet by mobile phone

DATAREPORTAL / GWI · DIGITAL 2026

**93.9%**

Brazilian internet users who used WhatsApp in the past month

DATAREPORTAL / GWI · DIGITAL 2026

## The four findings

**1**

### AI is part of the daily conversation, but not yet a driver of results.

Brazilian leaders and Stanford MBAs use ChatGPT multiple times a day, but almost no one builds agents or treats AI as infrastructure. The Stanford Digital Economy Lab's *Enterprise AI Playbook* (2026) confirms it: **77% of enterprise AI challenges are invisible** (change management, data, process), and NANDA/MIT estimates that **95% of GenAI pilots fail to deliver measurable financial impact**.

## 2 Brazilian companies are both further ahead than they look and further behind than they admit.

A financial-sector company is training 40 leaders to build agents. A pharmaceutical company cut its engagement survey cycle from three months to five days. A healthcare company allows only Copilot because of data-leakage concerns. The average hides the extremes. The OECD/BCG/INSEAD survey in São Paulo confirms the heterogeneity: only **7% of mid-sized and large companies** in manufacturing and ICT use AI, and among those that do, **58% have only one or two applications**. The bottleneck is not cultural familiarity with technology. It is turning a mobile-first, WhatsApp-first population into processes, data and governance inside companies.

## 3 The *ladder without a first rung* is structural in Brazil.

AI is absorbing the *grunt work* that used to train junior analysts, what Economist Impact called the *glass floor*. Without that rung, the white-collar entry door closes. With a young demographic pyramid and high informality, Brazil feels the effect before developed countries do.

## 4 Leadership cannot define the path for AI because it lacks fluency itself — and governance has no clear owner.

McKinsey (*Superagency*, 2025): **47% of the C-suite think their company is "too slow"**, and they are **2.4× more likely to blame employees than to take ownership of the bottleneck themselves**. But employees use AI 3× more than their leaders imagine. In Brazil, AI circulates between HR, IT and Compliance with no formal owner. In the public sector, the pattern appears even more starkly: **67% of Brazilian public servants say they learned AI on their own**, **68% say leaders do not provide clear direction**, and **49% would not know who to turn to if something went wrong**.

METHODOLOGICAL NOTE

All primary sources are anonymized by the authors' choice (not because of contractual confidentiality). The goal is to keep the focus on the pattern, not on the name. Where you see *"company in sector X"*, there is a real organization behind it, usually a segment leader. All external quantification is attributed to the original source.

CHAPTER 02

# Expectation × Reality

*What AI promises versus what the data show*

---

Every discussion about AI at work contains two parallel stories. The first is the expectation: a tsunami sweeping through offices, mass automation, immediate transformation. The second is the data, and it is more uncomfortable. Across every lens that matters (digital penetration, demand for *skills*, model capability, occupational exposure), there is a consistent distance between what is promised and what is measured. This chapter examines that distance from four angles.

## The Brazilian paradox: fluent on mobile, slow in the enterprise

Brazil's starting point is not technological lag. Quite the opposite: Brazil has one of the populations most accustomed to digital interfaces in the world. According to DataReportal/GWI, **98.4%** of Brazilian internet users access the internet by mobile phone, **93.9%** used WhatsApp in the past month, and the country spends **53h30 per week** consuming online media. INSEAD points in the same direction: more than **85% of Brazilians** use the internet regularly, the country leads Latin

America in GenAI adoption on LinkedIn with **110 thousand professionals** listing GenAI skills, and São Paulo concentrates more than **540 thousand** digital and AI professionals.

The paradox appears when this everyday fluency enters the enterprise. Demand for AI skills quadrupled in Brazil over three years, but organizational adoption is still narrow: in the OECD/BCG/INSEAD survey in São Paulo, only **7% of mid-sized and large companies** in manufacturing and ICT were using AI; among users, **58% had just one or two applications**. In other words: Brazilians already live inside digital interfaces, but Brazilian companies have not yet turned that familiarity into process, governance and productivity.

## Demand quadrupled. Total share is still 1.1%.

PwC's AI Jobs Barometer for Brazil shows a country where the curve already exists, even if the base is small: job postings demanding AI skills went from **19 thousand** in 2021 to **27 thousand** in 2022, **56 thousand** in 2023 and **73 thousand** in 2024. As a share of the total job market, that means going from around 0.28% to **1.1%** in three years.

### Job postings with AI skills in Brazil (2021 to 2024)

TOTAL JOB POSTINGS CITING AI SKILLS, PER YEAR

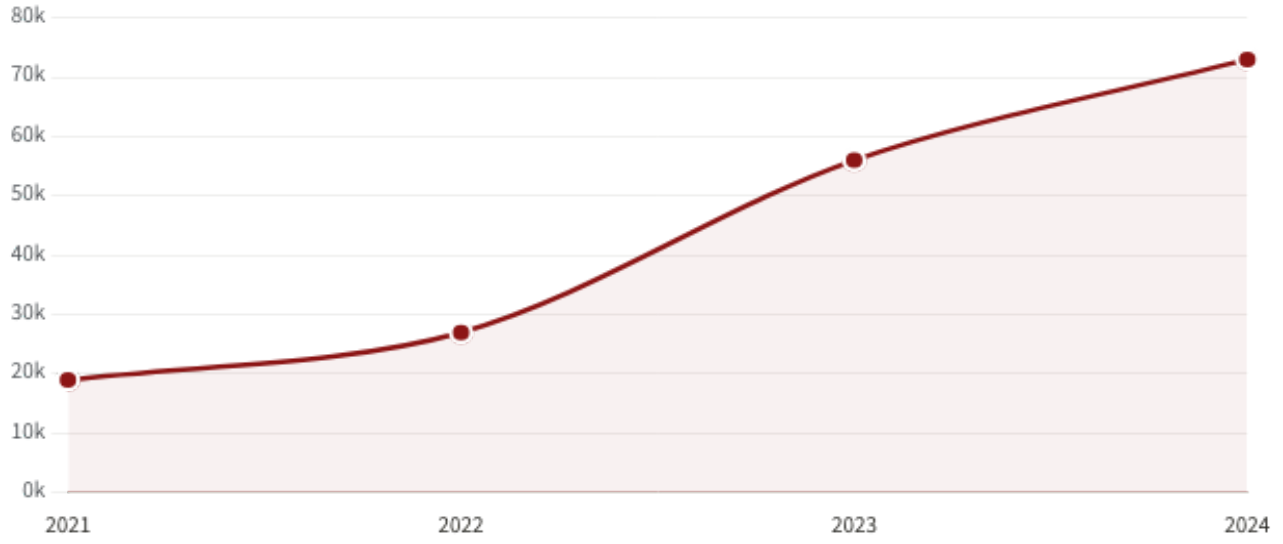


**READING** | The trajectory is clearly exponential, but the base is still small. Brazil is three to four years behind the US in relative volume, and sector concentration remains narrow.

Sector concentration is part of the story. In 2024, Information and Communication required AI in **4.2%** of its postings (more than triple the national average). Professional, Scientific and Technical Activities stood at **1.5%**. Finance and Insurance, around **3.5%**. Manufacturing, Healthcare and Education remain below the average, despite the narrative suggesting otherwise. *Where AI is discussed, it is not yet practiced.*

> *"Brazil's demand for AI skills is surging, quadrupling in just three years through 2024."*
>
> PWC AI JOBS BAROMETER · BRAZIL ANALYSIS 2025

## Theoretical capability × Real-world use

This pattern is not uniquely Brazilian. The paper *Labor Market Impacts of AI*, by Massenkoff and McCrory (Anthropic, March 2026), measures the distance between what AI is theoretically able to do and what it actually does today. The metric combines three sources: the O*NET database of American occupational tasks, the theoretical *β-scores* from Eloundou et al. (the LLM's technical capability per task) and real-world Claude use in professional contexts. The result is a **massive gap across all occupational categories**.

In **Computer & Math**, 94% of tasks are theoretically exposed; only **33%** appear in real-world use. In **Business & Finance**, 94% versus 28%. In **Legal**, 89% versus 15%. In **Engineering**, 85% versus 12%. About **30% of American workers** have *zero* coverage in Anthropic's data: cooks, mechanics, bartenders, and other occupations whose raw material is not digital. At the other end, professionals in high-exposure occupations earn **47% more**, are nearly **four times more likely to hold a graduate degree** (17.4% vs. 4.5%) and are **16 points more likely to be women**. AI has arrived, but only in a narrow slice of work.

# The capability exists. The use does not yet.

THEORETICAL COVERAGE × OBSERVED COVERAGE BY O*NET CATEGORY. ANTHROPIC, MARCH 2026

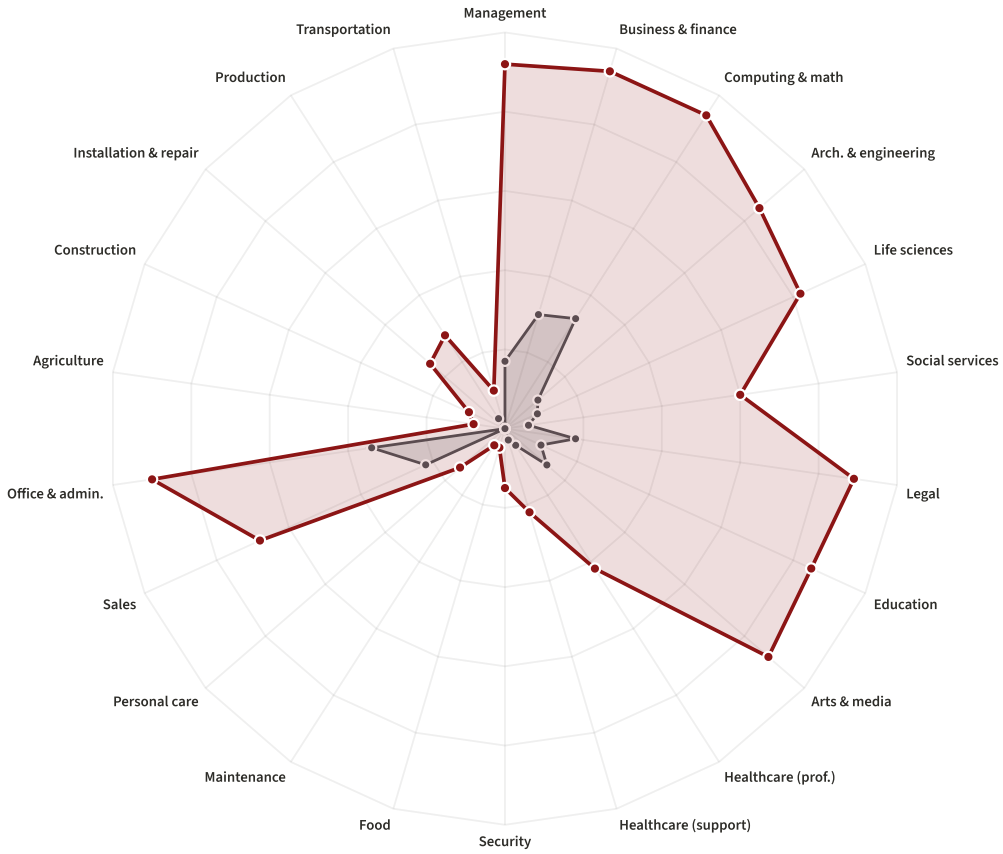


**READING** | The outer area (cardinal) shows what LLMs *could* automate by O*NET category; the inner area (gray) shows what already appears in real-world Claude use in professional contexts. The gap reaches 60 to 80 percentage points across white-collar categories such as Management, Legal, Education and Engineering. *Fluency is shallow not because of technical limits, but because of adoption.*

**Source:** Reproduction of Figure 2 from Massenkoff & McCrory (Anthropic), *"Labor Market Impacts of AI: A New Measure and Early Evidence"* (March 5, 2026). Theoretical capability = β-scores Eloundou et al. (2023, 0 to 1 scale); observed coverage = time-weighted average of O*NET tasks in each category that appear in real-world Claude use in professional contexts (full weight for automation, half weight for augmentation). Values anchored on *Computer & Math* (94% theoretical / 33% observed) and *Office & Admin* (90% theoretical), explicitly cited in the paper; other categories estimated visually from the original figure.

## At the top of exposure, but demand is not falling

Anthropic's *observed exposure* ranking (a metric combining theoretical capability with real-world use) places **programmers** (74.5%), **customer service representatives** (70.1%) and **data entry operators** (67.1%) at the top. Then come medical records specialists, market research analysts, wholesale sales reps, financial analysts and software QA. The composition is informative. *These are the professions whose raw material was already digitized.* AI did not create that risk; it only activated it.

### Top 10 occupations by observed AI exposure

% OF OCCUPATIONAL TASKS COVERED BY REAL-WORLD CLAUDE USE

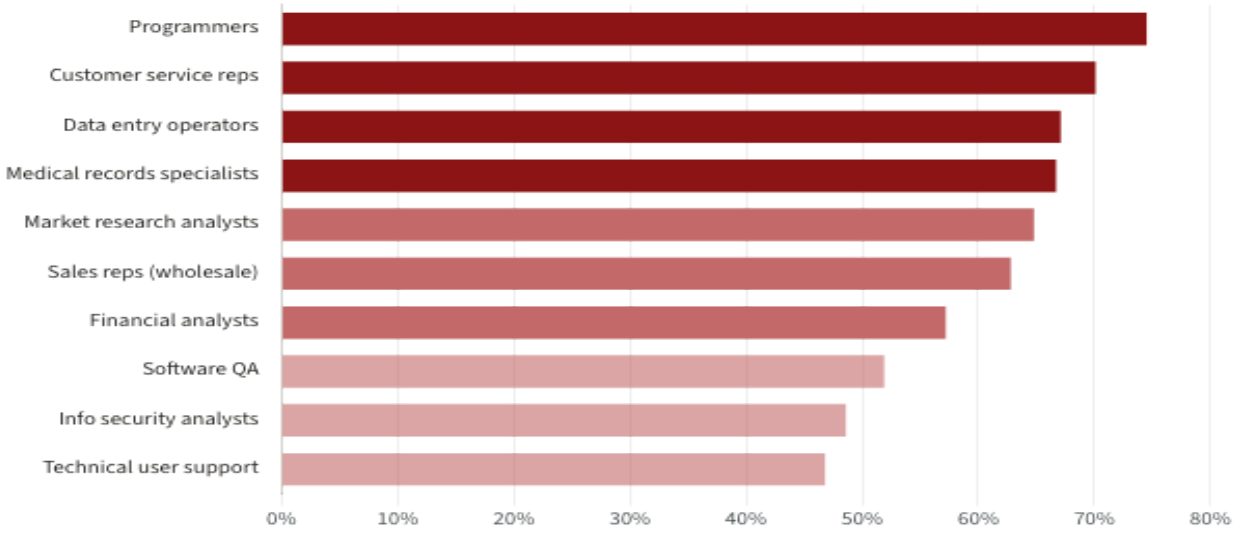


**READING** | None of the professions on this list is being eliminated. But all of them will see their productivity frontier redrawn, and the quality floor for what counts as "competent" will rise fast.

Software engineering is the most provocative example. Programmers sit at the **top of theoretical AI exposure** (74.5%), and yet global demand for engineers keeps rising. TrueUp's dashboard, cited by Lenny Rachitsky in March 2026, shows **67,665 open engineering positions at global tech companies**, up **78% from the 2024 trough**. The curve fell with the launch of ChatGPT, hit its low in early 2024 and has been rising since. In 2026, even after the launch of Opus 4.5 and Claude Code, it is the highest number of postings in three years. *Theoretical exposure is not the same thing as role extinction.*

## Global demand for software engineers despite highest AI exposure

OPEN ENGINEERING POSITIONS AT GLOBAL TECH COMPANIES. 2023 TO MARCH 2026

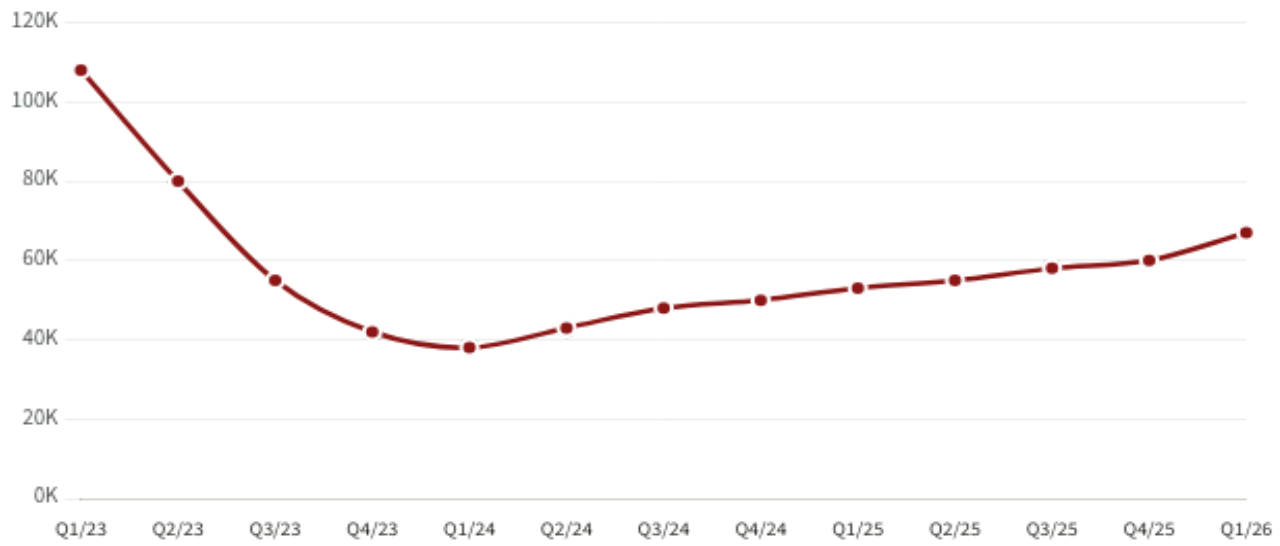


**READING** | The #1 profession in theoretical AI exposure saw demand fall with ChatGPT, but it has been climbing since 2024. In March 2026, there are 78% more openings than at the trough, even after Claude Code and Opus 4.5. **AI IS NOT CLOSING THE DOOR ON ENGINEERING. IT IS CHANGING WHAT ENGINEERING IS EXPECTED TO DELIVER.**

**EXPECTATION × REALITY · SYNTHESIS**

The four angles tell the same story. **Brazilians are mobile-first; companies are still procurement-first.** Demand for AI *skills* quadrupled, but represents just over 1% of total postings. The model's technical capability covers 94% of tasks in Computer & Math; real-world use covers 33%. Occupations with the highest theoretical exposure remain in high demand.

CHAPTER 03

# Shallow fluency

*Why high adoption is not turning into high productivity*

---

There is a wide distance between *using* AI and *operating* AI. The interviews in this study show that most Brazilian professionals, including C-level executives, are still firmly in the first group. The reasons are structural.

The most revealing evidence is not in the quantitative research, but in the language. In interviews with HR directors and MBAs, the phrase that comes up most often to describe current use is some variation of *"superficial"*, *"basic"* or *"email drafting"*. An HR director at a company in the pharmaceutical sector put it bluntly. **The use is ChatGPT for common tasks, not for strategic performance improvement.** The difference is not semantic. It is about ROI.

The quantitative data confirms the interview pattern. In a survey of **1,466 professionals** conducted by Brazilian career outlet *The News: Better Work*, **only 3% of respondents said they were satisfied with the way they use AI at work today**. The dissatisfaction is not with the tool. It is with the level of proficiency they have reached on their own.

> *"AI is daily use, but fluency is shallow."*
>
> SYNTHESIS OF INTERVIEWS WITH STANFORD GSB NETWORK TALENT · 2026

## The AI adoption cycle cannot keep up with the technology

A recurring pattern in the interviews explains the distance between use and fluency. Highly qualified professionals (some with backgrounds in *Big Tech*, strategy consulting, or executive leadership) describe the same dilemma. They realize they could automate a workflow, but the setup time for the automation seems to exceed the gain. The phrase captured is nearly identical across three separate interviews: *"I'll just do it manually"*.

On top of that is a gray zone between using AI as an individual and using AI as an employee. The same professional who uses ChatGPT naturally outside work often does not know what they can put into a tool inside the company, which data is sensitive, who approves an automation, or which platform is allowed. Extra caution around sensitive data is rational, but it has a cost: procurement and enterprise sales cycles move in months, while the tools change in weeks. The user, however fluent individually, ends up without resources, direction or institutional cover to use AI as a productivity lever.

The mismatch is measurable. The **METR Time Horizon** measures the duration of software tasks that agents can complete with 50% reliability. In four years, the curve went from *seconds* (GPT-3.5, March 2022) to *hours* of equivalent human work. **Technical capability doubles every few months; the enterprise cycle (procurement, governance, process change) is still measured in quarters and years.** That difference in clock speed is what produces shallow fluency: the technology arrives before the operating model.

## Agent capability accelerates. The enterprise cycle does not keep up.

METR TIME HORIZON · 50% SUCCESS ON SOFTWARE TASKS, MEASURED IN MINUTES OF EQUIVALENT HUMAN WORK

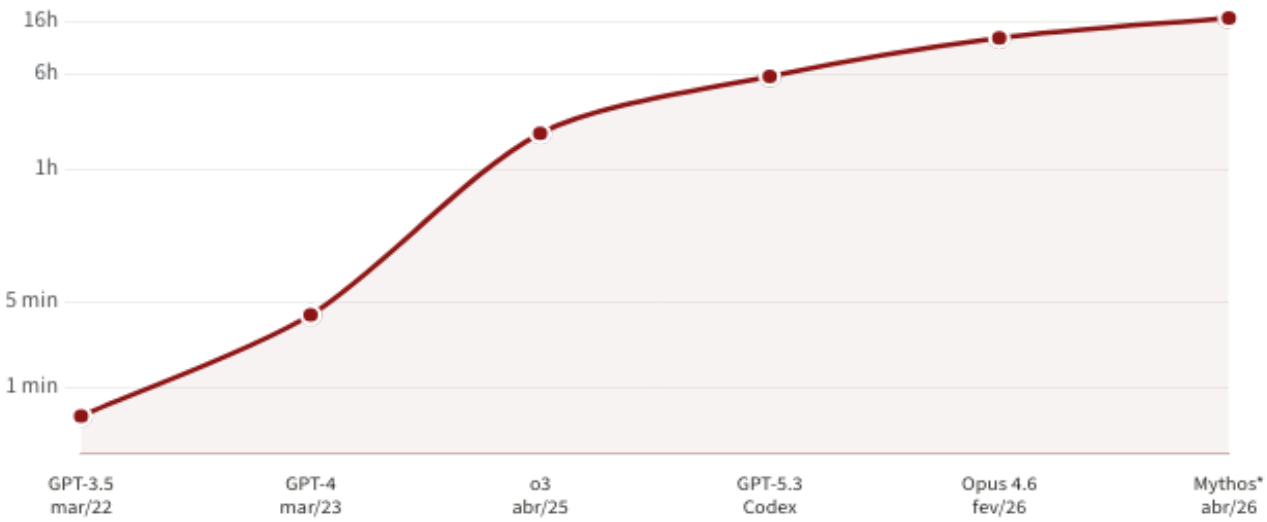


**READING** | METR does not measure job automation; it measures the duration of software tasks agents can complete with 50% reliability. The curve went from seconds/minutes to hours. *The asterisk matters: METR itself warns that measurements above 16h are still unreliable in the current suite.* Even with that caveat, the signal is clear: the ROI bottleneck has moved from model capability to workflow selection, verification, data, governance and procurement. All of those are *enterprise variables*, not technical ones.

In the Brazilian public sector, the same pattern appears as self-teaching at scale. The *Public Sector AI Adoption Index 2026* shows that **63%** of Brazilian public servants started using AI at work in the past year, almost always on their own initiative; **52%** describe their learning as entirely self-taught. That is a cultural strength, but it is also an institutional risk: when the learning curve grows from the bottom up, governance, security and ROI measurement arrive late.

The cognitive cost of thinking like an automator (choosing a tool, connecting APIs, debugging prompts, validating outputs) is real. And it is exactly that cost that separates those who use AI from those who operate AI. **What is missing is not technical training. It is a mindset shift around risk and testing.** The phrase comes directly from the HR director at a company in the pharmaceutical sector, and it echoes across almost every interview.

> *"Technical training is insufficient without a mindset shift around risk and testing."*

HR DIRECTOR · COMPANY IN THE PHARMACEUTICAL SECTOR (BRAZIL)

## What external data confirms

The BCG study with 480 consultants and 44 data scientists is the best experimental confirmation of this distance. When non-technical consultants received GenAI to perform *data cleaning* in Python, they jumped from **37%** of the data scientist benchmark to **86%** (a 49 percentage point gain). The result depended on one condition: some level of code experience. *"Code experience is a key success factor for workers using GenAI,"* BCG concludes.

### Data cleaning performance (% of data scientist benchmark)

BCG / BOSTON UNIVERSITY / OPENAI STUDY · 2024

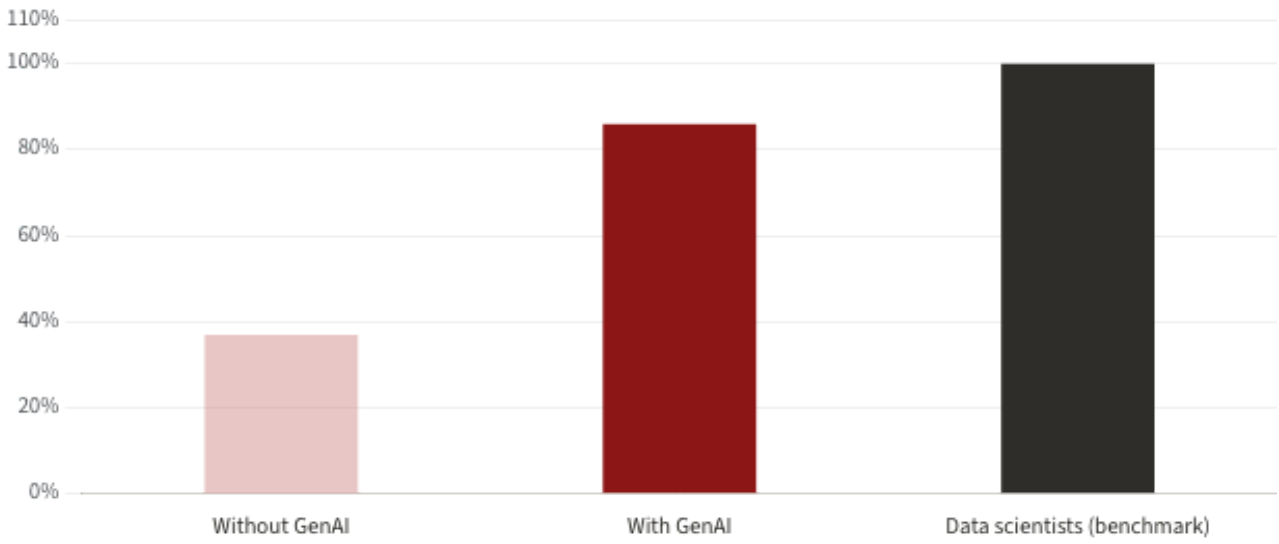


**READING** | The gain works like an exoskeleton. Those who already had a skeleton (a technical mindset) gained 49pp. Those who did not gained less.

**AI AMPLIFIES EXISTING CAPABILITIES; IT DOES NOT CREATE CAPABILITIES FROM SCRATCH.**

Another implication appears in the speed at which the *skills* themselves change. PwC global data shows that the skills required in AI-exposed jobs are changing **66% faster** than in others, and that rate is 2.5× higher than the previous year's. In Brazil, the effect appears in a specific way. Professionals in the quartile most exposed to AI register a **+1.6** *net skill change*, compared with

**+0.7** for the least exposed. Put simply: **those close to AI are relearning constantly. Those far from it stay frozen.**

> *"AI fluency will be critical for employability by 2026."*
>
> HR DIRECTOR · COMPANY IN THE FINANCIAL SECTOR (BRAZIL)

There is a direct analogy with digital literacy in the 1990s. Back then, knowing how to use email and Word was the minimum bar to sit at the table, not a differentiator. Knowing how to use AI is becoming that condition. The difference is speed. Digital literacy had a decade to spread. AI fluency has two or three years.

> SHALLOW FLUENCY · SYNTHESIS
>
> Brazil has an adoption problem that has not yet become a productivity gain. **The cost of learning to operate AI (not just use it) is the next frontier of competition.** Companies that recognize this and invest in protected pilots, with clear guardrails and permission to experiment, will see ROI jumps in the next 18 months.

CHAPTER 04

# Three portraits

*What Brazilian companies are doing, through three anonymized snapshots*

To move beyond the macro trend, the report interviewed HR directors in three segments of the Brazilian market. The portraits that follow show *real initiatives*, with scale, partnerships and metrics. Names were removed; the substance was preserved.

Click on each portrait to expand the detail. Together, the three form a spectrum that runs from the *scale operator in transformation* to the *regulatory fortress*.

## Company in the financial sector | SCALE OPERATOR +

## Company in the pharmaceutical sector | PRAGMATIC +

## Company in the healthcare sector | CAUTIOUS · COMPLIANCE-FIRST +

THREE PORTRAITS · SYNTHESIS

**There is no single right model. There is a right model for each sector.** The pattern across all three: *the companies that advance the most tend to be the ones that are most intentional about governance, access and training.*

CHAPTER 05

# The ladder without a first rung

*Why the real risk is the entry door closing*

---

There is a recurring narrative about AI and work: "it is going to eliminate jobs." The latest data show that is *not* what is happening, at least not yet. What is happening is subtler, and probably more serious. **AI is absorbing the tasks that historically trained junior professionals.**

A first piece of evidence comes from Anthropic Nowcasting in March 2026. Cross-referencing data from the American Current Population Survey with the exposure metric built from real-world Claude use, the authors reach a dual conclusion. **Unemployment in exposed occupations has not risen systematically since the launch of ChatGPT.** The aggregate effect is statistically indistinguishable from zero. But there is a significant hidden effect. **Hiring of young workers aged 22 to 25 in exposed occupations fell by around 14%** after 2022.

The **"canaries in the coal mine"** metaphor comes from another reference: the paper *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*, by Erik Brynjolfsson, Bharat Chandar and Ruyu Chen (November 2025), supported by the Stanford Digital Economy Lab. Using ADP payroll administrative data, the authors show a relative drop of **16%** in employment of workers aged 22 to 25 in AI-exposed occupations, after controlling for firm-level shocks, while employment of more experienced professionals remained stable. For software developers aged 22 to 25, the drop was **nearly 20%** from the late-2022 peak. The signal is not mass layoffs. It is that the first rung is no longer bringing people in.

The first equivalent signal in Brazilian primary data came in 2025. In an independent analysis by Elisa Pereira published in October, the occupations most exposed to AI in Brazil registered an average decline of **5.4%** in hiring in the first three quarters of 2025, relative to what was projected by the 2023 to 2024 historical average adjusted for the growth of low-exposure occupations. The effect is heterogeneous. **Collections fell 21.6%** (about 10 thousand fewer hires), **administrative support fell 5.4%** (~29 thousand), and accounting assistants dropped 3.5%. On the other hand, HR analysts grew 8.5% and active telemarketing 8.3%. It is the first national signal that the entry door has begun to close here too, in specific occupations, before appearing in the aggregate.

## Employment of young workers 22 to 25. AI-exposed vs. non-exposed occupations.

RELATIVE EMPLOYMENT INDEX (2018 = 100), WORKERS AGED 22 TO 25, US

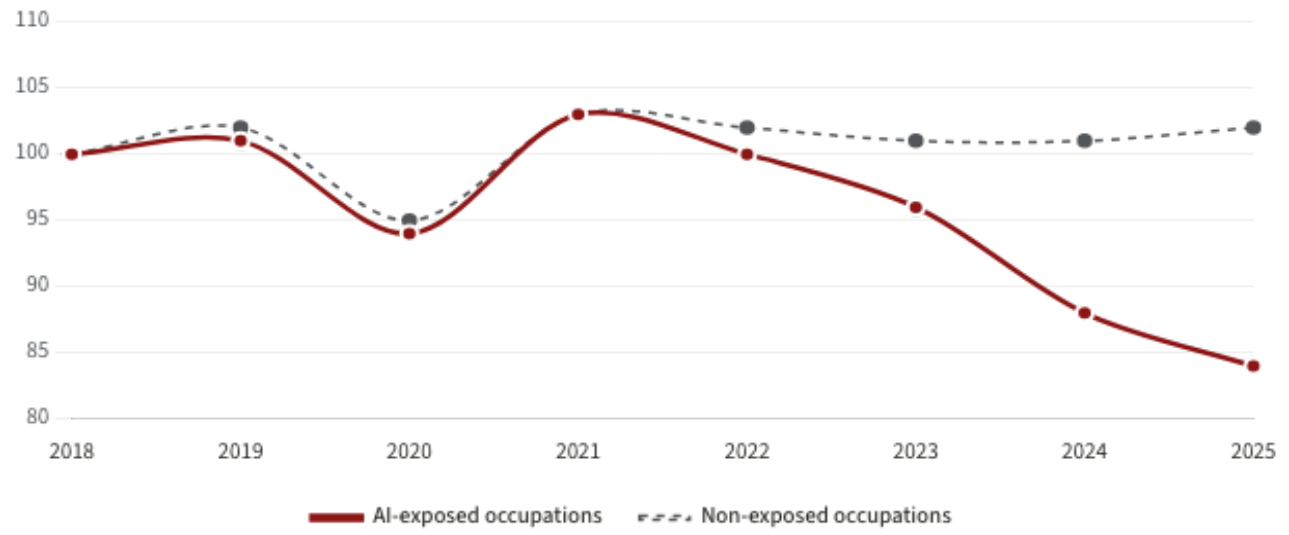


**READING** | The lines move together through 2022. From late 2023, exposed occupations diverge downward. *Young workers are not being fired. They are simply not being hired.*

**Source:** stylized illustration built from Brynjolfsson, Chandar & Chen, *"Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence"* (Stanford Digital Economy Lab, November 2025), with a ~16% relative drop in employment of workers aged 22 to 25 in AI-exposed occupations, and Anthropic, *"Labor Market Impacts of AI"* / Nowcasting (March 2026), with a ~14% drop in the job-finding rate for that same group. The annual values plotted here are a visual representation of the relative magnitude reported in those studies, not original datapoints.

Economist Impact, in the report *The AI Glass Floor*, describes the dynamic in one sentence: *"AI may not eliminate jobs, but it can block the door of entry into them"*. The intuition is direct. The tasks a junior analyst used to do to learn the trade (building spreadsheets, doing background research, writing first drafts, processing simple tickets) are precisely the most automatable. Without those tasks, there is no learning ramp. Without a ramp, there is no hiring. Without hiring, there is no next generation of senior professionals.

> *"Juniors are, increasingly, being left outside, looking for places to develop corporate skills."*
>
> MATT BEANE · MIT, CITED IN "THE AI GLASS FLOOR" (ECONOMIST IMPACT)

## Why Brazil is at greater risk

Three structural factors make the ladder without a first rung sharper in Brazil than in the US. **First, demographics**: Brazil is already in transition (median age of 35 according to the 2022 IBGE Census), but still concentrates **around 35 to 40 million people between 18 and 29 years old**. Proportionally, that age group weighs more here than in the US: around **17 to 20%** of the Brazilian population versus **~15%** of the American population. Every percentage point of decline in the *job-finding rate* applied to this cohort equals **hundreds of thousands of young people** fewer entering the formal market.

**Second, informality.** The floor was already fragile. Professionals who historically entered the market through formal internships and migrated to CLT employment now compete with platform freelancers and self-employed professionals who use AI to scale. The result is margin compression and reduced need for juniors within formal structures.

**Third, university.** As heard in the interview with the HR director at a company in the healthcare sector: *"universities are not teaching AI in the curriculum, and companies are not training what makes a good analyst in the AI era"*. The Brazilian education system is three to five years behind in practical training. Graduates in applied social sciences are entering the market without operational AI fluency, and losing out to peers who learned in a bootcamp or through self-teaching.

### Demand growth by occupation in Brazil (2021 to 2024)

AVERAGE % GROWTH IN JOB POSTINGS, BY AI EXPOSURE LEVEL

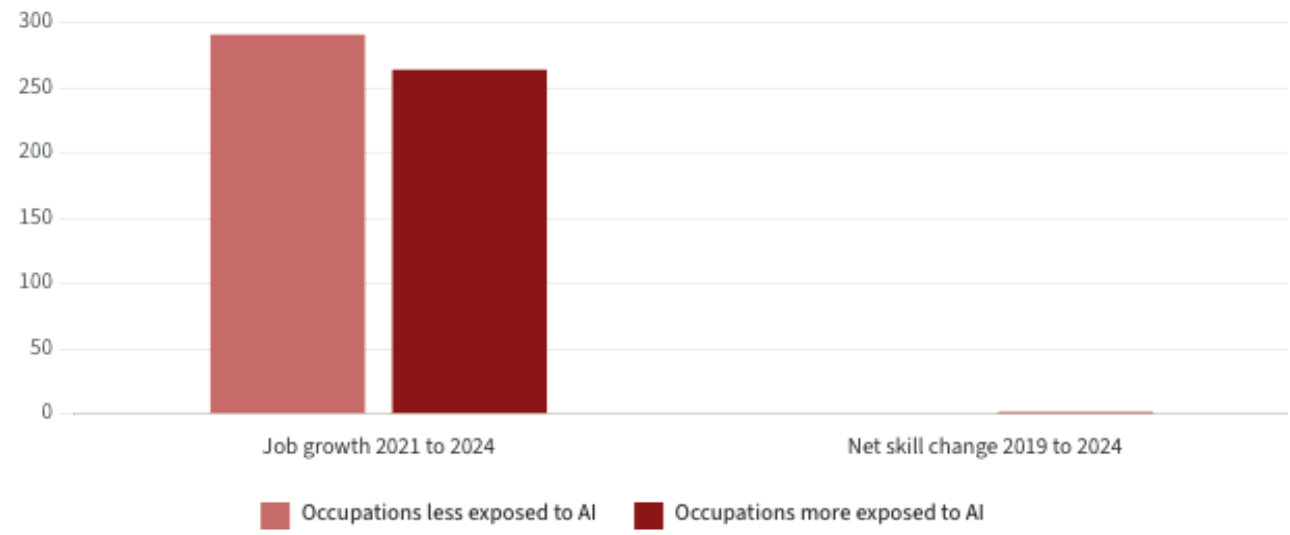


**READING** | The occupations *less* exposed to AI grew 27 percentage points more than the most exposed.

**THE CLOSER TO AI, THE LOWER THE GROWTH IN POSTINGS AND THE GREATER THE REQUIRED SKILL CHANGE.**

## The combination that hits hard

The data connect in an uncomfortable way. In Brazil, demand for AI *skills* quadrupled; but the most exposed occupations are growing less in number of postings (264% vs. 291% for the least exposed) and require 135% more skill adaptation. Abroad, the wage *premium* for those who master AI more than doubled in a year (from 25% to 56%). And Anthropic shows that workers in the most exposed quartile earn, on average, **47% more** than those with zero exposure, are **4× more likely** to hold a graduate degree, and are **less unionized**.

Put simply: *those who master AI are being well paid and retained. Those who do not yet have a door through which to learn will not get in.*

# The pipeline that eats itself

The effect is not a sequence of isolated mistakes. It is a single composite dynamic, looping back on itself. Seven steps reinforce one another.

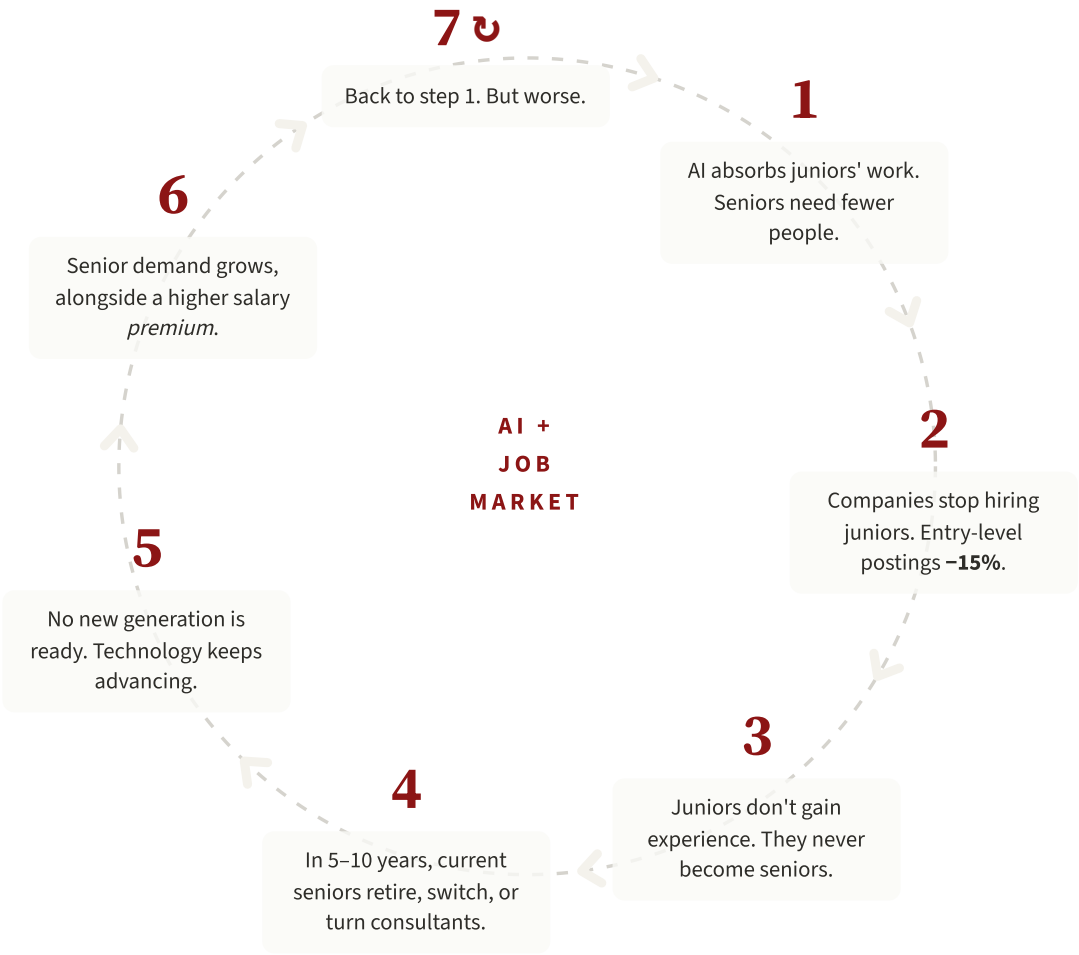


CYCLE MODELED FROM DATA FROM ANTHROPIC, PWC AND ECONOMIST IMPACT

LADDER WITHOUT A FIRST RUNG · SYNTHESIS

Brazil's short-term AI risk is not mass unemployment. It is **a generation of young professionals who may never learn.** The tasks that used to train them are disappearing, universities have not filled the vacuum, and companies have not yet built alternative training tracks.

CHAPTER 06

# Paradoxes

*Five tensions that every strategy needs to resolve*

Consulting reports like to present AI as a one-way vector. The interviews show something else. The terrain is full of tensions, and each one is both an obstacle and an opportunity. Mapping these paradoxes is more useful than listing best practices.

## 01 High investment, superficial use

**100% of interviewed companies already have at least one AI tool in use**, with partnerships signed with Microsoft, Fundação Dom Cabral and executive education platforms. Even so, they admit real-world use remains "ChatGPT for basic tasks." ROI does not come from buying the tool. *There is no clarity on how to measure impact, adoption, or each employee's individual fluency.* Without those three layers, investment becomes a cost line, not a productivity vector.

## 02 Leadership is not an example of AI productivity

The discourse of generalized anxiety about AI (about automation, substitution, the future of work) often hides a transfer of responsibility: *it is not clear whose job it is to train the workforce in AI within each sector.* Without an owner, capability building gets suspended between HR, IT, business units and external vendors. No one is accountable for the result.

---

## 03 HR wants to be strategic. It is not at the AI table.

HR directors describe a loss of relevance after the pandemic and a desire to return to the CEO's table. According to one interviewee, **about 80% of HR teams do not understand their own end-to-end journey**. The paradox is that *much of the AI adoption journey is, in practice, upskilling*, the natural territory of HR. But HR is not at the center of this discussion inside companies, and AI decisions end up happening without the team that should be leading them.

---

## 04 AI is a priority, but the adoption machine is slow

The AI agenda is among the most relevant for leadership, HR and technology. But the speed of adoption, upskilling and procurement remains a bottleneck. Tools change in weeks; enterprise purchases, compliance approval, internal data integration and guardrail design move in months. Between executive enthusiasm and real productivity sits an organizational queue.

---

## 05 Self-teaching accelerates adoption, but leaves governance blind

Brazil learns fast on its own. In the public sector, **67%** of public servants say their AI knowledge is entirely or mostly self-taught; in the interviews, talent tests tools before a corporate playbook even exists. That creates cultural speed, but also *shadow AI*: the company does not know who is using it, in which workflow, with what data, or to what quality standard.

---

PARADOXES · SYNTHESIS

The paradoxes do not point to a single solution. They show where adoption gets stuck: in the distance between tool and process, between enthusiasm and permission, between individual fluency and institutional governance. The opportunity lies less in selling "more AI" and more in reducing those frictions in concrete ways.

CHAPTER 07

# Opportunities & challenges

*Where to build, and what needs to be unlocked first*

The opportunities below show where AI can become operational capability. The challenges show the locks preventing that capability from scaling.

OPPORTUNITIES

## Where to look in the next 18 months

### Top-down and bottom-up strategy

AI adoption is multidirectional. Managers who operate AI fluently increase their own productivity and free up time to spend with their teams — a use case Daniela Amodei described as *"AI coach for leaders"*. Better-guided team members and ICs with access to the same tools amplify their output. Organizational capability emerges from combining both layers, not from choosing one over the other.

### Intentional evaluation of AI fluency in the hiring process

The goal is not to ban AI use in challenges and tests, but to validate the candidate's capability and fluency with the tools and routines they will operate after being hired. Treating AI as an evaluated variable (and not as contamination to be blocked) repositions recruiting: less measurement of prior knowledge, more measurement of the ability to operate at the current frontier.

### Automation governance

An operating model for deciding which workflows can be automated, who approves, what data can be used, how the output is verified and how the gain is measured. It can become a product, consulting offer, internal capability or platform layer. The point is less to sell control than to give people safe permission to use AI.

### Smart measurement

Metrics for fluency and productivity need to go beyond "active licenses" or "prompts sent." Companies need to know which teams are using AI in relevant workflows, at what quality level, how much time they save and with what risk. Without that layer, adoption becomes narrative.

---

CHALLENGES

## From macro to micro

### Outdated university

The Brazilian curriculum is three to five years behind in practical AI training. Without rapid reform, the alternative falls between the private bootcamp ecosystem (still expensive and fragmented) and self-teaching. Companies will play an increasingly essential role in the adoption curve, and building an environment that encourages experimentation and curiosity under the constraints of compliance remains a challenge.

### Cross-functional adoption models are limited

The AI agenda inside companies is treated as a cross-functional project: it passes through HR, IT, Compliance and business units in one-off pilots, isolated training and executive education partnerships, but rarely as a continuous track. Without that continuity, each new wave of investment starts from scratch, and the company's learning curve remains concentrated in individuals who upskill on their own.

### Compliance vs. speed

In regulated sectors, caution is rational: clinical, financial and customer data cannot flow through just any tool. The problem appears when compliance operates only as a binary veto or approval. Without safe test environments, clear data classification and risk-approved workflows, the company is forced to choose between official delay and informal off-radar use.

**Regional inequality**

Adoption is concentrated in urban centers in the Southeast. The interior of the North and Northeast operates with insufficient digital infrastructure. *Without a regional bridge, AI amplifies existing inequality.*

> *"AI should reduce work pressure by taking care of routine tasks, allowing humans to focus on interaction and care."*
>
> HR DIRECTOR · COMPANY IN THE HEALTHCARE SECTOR (BRAZIL)

OPPORTUNITIES · SYNTHESIS

The market does not need another generic list of "AI use cases." It needs mechanisms that move adoption from the individual into the operation: capable managers, faster onboarding, governance that unblocks instead of blocking, and measurement that shows where there is real gain. Without that, the company buys the tool but keeps running transformation on intuition.

CHAPTER 08

# Final provocation

*What is still under control*

---

The conclusion of this report is not that Brazilian companies need to "use more AI." That is already happening, even if unevenly. The conclusion is more uncomfortable: individual use has advanced faster than the institutional capacity to guide, measure and absorb that use.

The problem, therefore, is not choosing between enthusiasm and caution. It is designing the bridge between the two. In large companies, that bridge runs through three concrete decisions.

**First: name the operational owner.** AI cannot keep circulating between HR, IT, Compliance and business units as everyone's topic and no one's responsibility. The company needs someone with a mandate to define guardrails, prioritize workflows, measure results and resolve the conflict between risk and speed.

**Second: pick a few real workflows.** Adoption is not proven by active licenses, training completed or number of prompts. It is proven when an important process becomes faster, cheaper, more reliable or more auditable. Three well-chosen workflows say more about maturity than fifty scattered pilots.

**Third: start from the business objective.** AI adoption is never an end goal. It is a means to reduce cycle time, increase conversion, decrease error, accelerate onboarding, improve quality or free up team capacity. If a workflow cannot say which business metric it wants to move, it is still a tech experiment, not an operational strategy.

## The risk is not running out of AI. It is having AI everywhere and transformation nowhere.

Brazil has digital users, productivity pressure and talent learning on its own. What is still missing is turning that scattered movement into organizational capability: a clear way to decide what to automate, who can use it, with what data, under which metric and with what accountability.

BERNARDO PRECHT · FELIPE LOURENÇO · MAY 2026

CHAPTER 09

# Methodology

*How the report was built*

This report is more useful when its construction is clear: which sources were cross-referenced, what type of inference was made and what is out of scope. This section documents the research in six layers, from framing to acknowledged limitations.

## Framing and period

Research conducted over eight months, closing in May 2026. The focus is the intersection between the Brazilian labor market and the global AI movement, with a dual lens: **the workforce entering**

**the market** (Stanford MBAs, recent graduates, career transitions) and **the leaders who recruit, train and make decisions about that workforce** (HR directors, C-levels, board members). The framing is intentional: it avoids both the bias of research only with decision-makers and the bias of research only with talent.

## Primary sources

More than 50 hours distributed across in-depth interviews with HR managers and directors, C-levels and Stanford MBAs; unstructured conversations with talent and ecosystem operators; observations and discussions in MBA program classes; and analysis of internal materials gathered throughout the research. Three Brazilian company portraits were built from formal interviews with the HR leadership of each company, with subsequent confirmation of numbers and initiatives against the original transcript.

## Secondary sources

Review of reports, academic papers, public datasets, specialized newsletters and external materials on AI, work, enterprise adoption, HR Tech, productivity, education and the Brazilian market. The main references cited in the report:


Pereira, Graylin & Brynjolfsson, "The Enterprise AI Playbook: Lessons from 51 Successful Deployments" (Stanford Digital Economy Lab, April 2026)
Brynjolfsson, Chandar & Chen, "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence" (November 2025)
MIT NANDA, "The GenAI Divide: State of AI in Business 2025"
McKinsey & Company, "Superagency in the Workplace: Empowering people to unlock AI's full potential" (January 2025)
Tamkin, Handa, Durmus, Ganguli et al. (Anthropic), "Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations" (February 2025, arXiv:2503.04761)
Massenkoff & McCrory (Anthropic), "Labor Market Impacts of AI: A New Measure and Early Evidence" (March 5, 2026)
METR, "Task-Completion Time Horizons of Frontier AI Models" (May 2026)
PwC AI Jobs Barometer Brazil 2025 and PwC AI Jobs Barometer Global 2025
OECD / BCG / INSEAD, "The Adoption of Artificial Intelligence in Firms: New Evidence for Policymaking" (2025)
INSEAD Knowledge, "Brazil's Digital and AI Talent Landscape: Leapfrogging or Left Behind?"

DataReportal / GWI, "Digital 2026 Global Overview Report"
Public Sector AI Adoption Index 2026 (Public First / Center for Data Innovation, via Kiteworks)
BCG / Boston University / OpenAI, "GenAI Doesn't Just Increase Productivity. It Expands Capabilities" (2024)
Economist Impact, "The AI Glass Floor" (2025)
Pereira, Elisa C. M., "Some AI-Exposed Occupations Already Show 5.4% Hiring Decline in Brazil" (Medium, October 2025)
*The News: Better Work* (thejobs.com.br), survey of 1,466 Brazilian professionals on satisfaction with AI use at work (2026)
Drake Star, "Global HR Tech Report Q3 2025"
TrueUp Engineering Jobs Dashboard, via Lenny Rachitsky, "State of the Product Job Market" (March 2026)
Business Insider, "Meta Is Tracking Employee AI Use"

## Triangulation

Most of the quantitative data on occupational exposure, hiring contraction and technical capability comes from American research (Anthropic, Stanford, METR, BCG/BU/OpenAI). Brazil does not yet have equivalent series at the same scale of academic production. The report's approach was to cross those findings with Brazilian data whenever available (PwC AI Jobs Barometer Brazil, OECD/BCG/INSEAD São Paulo, INSEAD Knowledge Brazil, Pereira's CAGED analysis) and with qualitative evidence from interviews to verify whether the pattern observed abroad was reproduced locally. When Brazilian data did not exist, the international source is presented as directional evidence, not as local measurement.

## Data treatment and attribution

All quantification was referenced to the original source. When a chart in the report visually reproduces a figure from a paper whose authors did not publish the underlying numerical table (the case of Figure 2 from Massenkoff & McCrory, chapter 02), the values were anchored on the explicit points cited in the body of the paper, with the remainder estimated by visual inspection. The chart source notes that limitation. The charts about youth employment in exposed occupations (chapter 05) are stylized illustrations of the relative magnitude reported by the papers, not original datapoints; this is stated in the chart source.

## Acknowledged limitations

Three limitations the reader should keep in mind. **First:** the three company portraits are intentional and narrative, not a statistically representative sample of the Brazilian market; they serve to illuminate patterns, not to quantify them. **Second:** no proprietary sample survey was conducted, which limits the report's ability to produce new Brazilian estimates; the work depends on third-party surveys (PwC, OECD/BCG/INSEAD, *The News: Better Work*) for quantitative data on Brazil. **Third:** the speed of the AI market makes any conclusion sensitive to the time window. The data in this report reflects the state of the art through May 2026; periodic revisions will be necessary.

CHAPTER 10

# The authors

*About the people writing this report*

---

## Bernardo Precht

SURGEON · FOUNDER
STANFORD GSB SLOAN FELLOW '26

Surgeon and entrepreneur. He founded *Além da Medicina* during medical residency, built the company between shifts, and turned a mentorship business into a scalable software platform for medical education and career development.

The company was acquired by Afya, Brazil's largest healthcare ecosystem, where he led digital products. He was recognized on the Forbes Under 30 list.

He is currently a Sloan Fellow in the MSx program at the Stanford Graduate School of Business (class of 2026), where he deepens his work in leadership, innovation and entrepreneurship.

# Felipe Lourenço

**FOUNDER · 2× EXITS**
**STANFORD GSB SLOAN FELLOW '25**

Founder with two SaaS *exits* in Brazil. As Founder & CEO of *iClinic*, he built the country's #1 clinical SaaS, backed by SoftBank and acquired by Afya (NASDAQ:AFYA), where he led the digital health division. Previously, he co-founded *Stored*, a platform acquired by Spiti/XP Inc. (NASDAQ:XP).

He graduated from USP and holds a master's degree from the Stanford Graduate School of Business as a Sloan Fellow. He currently works as an angel investor in AI, SaaS and healthcare in Brazil and the Valley, an independent board member nominated by VC and private equity funds, and an advisor to Silicon Valley startups expanding into LatAm. He also serves as AI Transformation advisor at Hospital Albert Einstein, among other institutions.