albert

The AI Cookbook

2026

Methodology Analysis Framework

Foreword

WHY THIS BOOK, AND WHO IS IT FOR?

Last year, we published The SWP Cookbook. Our starting point was simple: the real sticking point in Strategic Workforce Planning isn't technical or technological — it's the methodology. Everyone brought different assumptions and expectations to the planning exercise, and it seemed useful to lay down a framework.

With AI, we're seeing exactly the same thing, only worse. The noise is deafening. Not a week goes by without a new 'revolution,' a study announcing the disappearance of millions of jobs, or a vendor promising to replace half your workforce. And faced with this cacophony, SWP practitioners — that's you — find themselves on the front line, expected to answer a question nobody has a reliable answer to: what is the actual impact of AI on our workforce and skills?

Through our work as a SWP software vendor, AI touches us at Albert in three ways:

- Its use in our day-to-day production workflows
- Its impact on our clients' headcounts and skills
- Its integration into our service offering

These three angles give us a perspective that few players have: that of practitioners who use AI daily, who support large organizations in their planning, and who observe the often-considerable gap between the rhetoric and the reality on the ground.

This guide grew out of that experience. Like The SWP Cookbook, there's no required reading order. Feel free to jump straight to the sections that interest you:

- **Chapter 1** demystifies AI jargon and gives a fact-based overview of use cases by function — what works, what sometimes works, and what doesn't work yet
- **Chapter 2** analyzes real data on AI's employment impact (spoiler: it's not entire professions disappearing — it's an age group getting quietly squeezed out)
- **Chapter 3** lays out our thesis: why we must stop thinking in terms of tasks and start thinking in terms of objectives, and why the future is about doing "more with more" (not "more with less").
- **Chapter 4** translates all of this into SWP methodology: how to factor in AI as a business driver, what to track, and how to adapt your action plans

Whether you're a SWP lead, CHRO, HRBP, or your responsibilities sit on the business and finance side, you'll find in this guide a fact-based resource to help you navigate the ambient fog. As ever with us: our ambition isn't to predict the future, but to plan for it.

Even when we're being a little cheeky, we have a sense of nuance. Facts over emotion, subtlety over simplism. This guide is written in that spirit. Skip the paragraphs where you learn nothing — no one will be offended.

Vincent Barat - CEO - Albert

About us

VINCENT BARAT - CEO OF ALBERT - AUTHOR

Vincent is an engineer who graduated from CentraleSupélec ‘08 and the University of Wollongong. He co-founded the consulting firm Akoya in 2010 where he was in charge of commercial development in the luxury, retail & consumer goods sectors.
As the CEO of Albert since 2021, Vincent defines and leads Albert’scommercial efforts, thanks to his strong connections built with CHROs during his time at Akoya. He also oversees the company's overall development that is guided by the vision he instills.

BARBARA LOPEZ - CMO OF ALBERT - DESIGN & PUBLISHING

Barbara, who graduated from INSEEC Paris with a Master's degree in Communications in 2013, began her career as Marketing Manager at Nextperf (acquired by Rakuten). There, she piloted the marketing strategy for France and the DACH zone.
In 2018, she joins ARMIS as Head of Marketing & Communication, where she deploys lead generation and branding strategies, while structuring the marketing team into competency clusters. In 2023, Barbara joins Albert to strengthen Marketing & Communication initiatives in close collaboration with Vincent and the sales team.

Table of contents

Terms & use cases

WORDS BEFORE WOUNDS

The topic has become so ubiquitous it's frankly exhausting. You know the one — AI, and more specifically Gen AI, the ultimate buzzword. For three years now, ever since the public launch of ChatGPT, the media, conferences, trade shows, and a certain breed of 'experts' and opportunistic entrepreneurs have been milking it for all it's worth. New models from Anthropic, OpenAI, or Google; acquisition rumors and hyper-scale data center projects cooled with aircraft turbines; China's low-cost DeepSeek crashing Nvidia's stock price; European sovereignty worries anchored around our national champion Mistral; agentic AI with Manus; Grok's virtual girlfriends... The relentless churn of news makes the whole field completely illegible.

Those who have followed us for a while know we have a deep aversion to opportunism. That's why AI has never been a marketing platform for us. But with its 'commodification,' we think the topic has now reached an inflection point that calls for some clarity.

But before we get to the wounds of AI, let's talk about the words. The goal isn't to do an exhaustive dive into the history of AI — there's plenty of better material out there for that. But it does seem essential to democratize the jargon, so you don't get bamboozled by silicon quacks.

So here's a slightly unconventional little AI dictionary to help you find your bearings. Except it's not in alphabetical order. And it's not THAT unconventional. You get the idea.

A (not so) unconventional little AI dictionary

AI vs. Gen AI

The term 'artificial intelligence' is extremely vague. 'Traditional' AI has been working away around us for a long time already — in medical research, recommendation algorithms, fraud detection, video games, high-frequency trading. These AIs are highly specialized at a single task. Gen AI is different: it's more generalist, and it 'produces' knowledge. Or at least appears to, because in reality it just generates content. There's a vast ecosystem of specialized AIs: for sound, images, text, voice, video, and some that can mobilize multiple 'specialisms' to respond to a user's

LLM (Large Language Model)

The most spectacular exploits and hallucinations of Gen AI are usually textual. In those cases, they're powered by Large Language Models. Remember T9 predictive text for writing SMS on your Nokia 3310? It's roughly the same thing. Except that instead of guessing the character or word you're trying to type, the model predicts the next hundred or thousand words, based on a very large training dataset. That's also why it's notoriously bad at answering the question 'how many R's are in strawberry.'

Training and bias

Algorithms are built on a training mechanism, supervised or unsupervised. The way the model is trained and the source material influence the outputs — hence the famous biases. Ask two AIs to picture two corporate executives: one generates two men, the other generates a man and a woman. Statistically, the first is simply reproducing a historical bias. The second is introducing an inclusive one. This also raises intellectual property questions (content creators never gave any AI permission to impersonate them) and security issues (a pharma lab asking confidential questions to a shared AI risks inadvertently training the model on its own secrets).

Publishers, models and products

These three are often confused. **Publishers** are companies, either general-purpose (Meta, Google, X) or specialized (OpenAI, Anthropic, Mistral). They develop **models** (GPT-4o, Claude Sonnet, DeepSeek R1, Llama), which end up in **products**: ChatGPT, Claude, Le Chat, Grok. Some models are open (downloadable by anyone), others are closed. There are even app stores for them (French unicorn Hugging Face lists hundreds). The honest truth is that nobody today knows exactly what a query costs, or what it should cost for the underlying business model to be sustainable.

Token

The unit of currency in the AI economy. When you send a request to an LLM, it gets broken down into tokens (word fragments, typically 3–4 characters). The longer your request and the more detailed the response, the more tokens you consume, and the more it costs. Some queries use a few dozen ("Hey, how's it going?"), others consume hundreds of thousands ("Rewrite War and Peace using gender-neutral language"). It's a bit like mobile data on your phone plan: you don't really know how much you're using, but the bill eventually arrives.

RAG (Retrieval Augmented Generation)

When you ask ChatGPT a question, it answers based on its general training — i.e., everything it has 'read' before a certain cutoff date. RAG adds a step: before answering, the AI first searches for information in a specific document base you provide (your security policy, your collective agreements, your product documentation).
This is what makes custom AI assistants useful in enterprise settings: instead of answering 'in general,' they answer in your context. It also reduces (though doesn't eliminate) hallucinations.

Fine-tuning

Fine-tuning means retraining an existing model on your own data so it behaves in a specific way. It's heavier and more expensive than RAG, and often confused with it. In practice, RAG is sufficient for the vast majority of enterprise use cases. Fine-tuning makes sense when you need the model to adopt a style, vocabulary or reasoning pattern very specific to your domain. It's the heavy artillery — rarely the right first move.

Reasoning

Recent models (5.3 from OpenAI, R1 from DeepSeek) have introduced the ability to break down a problem into pieces (Chain of Thought) and to backtrack in their predictions to produce more coherent responses. This can be genuinely unsettling — you really do get the impression the AI is thinking. But this is absolutely not thinking in any cognitive sense. The perception is fueled by a trick: the AI displays text as it's responding. To top it all off, some companies base their entire communications strategy on fear marketing: 'We tested our next model internally and, honestly, we scared ourselves. By the way, the new model will only be available to paying customers. Subscribe now, xoxo.'

Agentic AI

The hot thing right now. Imagine an AI assistant that doesn't just suggest a meeting time — it sends the invite, books the room, and adjusts the slot if someone declines. The model no longer passively responds; it acts autonomously to achieve a goal. The current limitation is reliability. Because if a chatbot hallucinates, fine. But if it's an agent that hallucinates and then acts on it, that's where things get dangerous.

AGI

Artificial General Intelligence. The ultimate quest: an autonomous AI so powerful across all domains that it surpasses human capabilities. Today, it's a theoretical concept that has become a marketing argument.

What AI can (actually) do in a business context

They often say rugby is the only sport that needs everyone: the small ones, the tall ones, the heavy ones, the fast ones, the idiots and the crafty ones. Gen AI users are the same: there are skeptics, the wary, the enthusiasts, the naive, and the brilliant. And you need to be a bit of all of them to use it effectively.

78% of companies use AI in at least one area in 2025 (Stanford AI Index). Enterprise spending on Gen AI tripled in a year, reaching $37 billion (Menlo Ventures). But behind these aggregated figures, the reality is more nuanced than it appears: only 6% of companies are 'high performers' with measurable impact on EBIT (McKinsey), and 42% have abandoned most of their AI initiatives in 2025, up from 17% in 2024. Nearly one in two projects.

Adoption looks massive, but real impact is still very concentrated. To cut through the noise, we cross-referenced the main market studies to build a function-by-function snapshot of what works, what sometimes works, and what doesn't work yet.

AI Maturity Matrix by Function

Function	Works	Depends	Not yet
Support functions (HR, legal, finance, compliance)	Assisted drafting, document summarization, compliance questionnaires	Contract analysis, automated financial reporting, CV screening	Arbitration (payroll, disputes, performance reviews)
Marketing & communications	Ideation, content repurposing, campaign strategy support	Brand content production, SEO	Autonomous content creation
Customer relations & after-sales	Ticket triage and routing, automated FAQs	Chatbots for simple queries, sentiment analysis	Complex requests, individual cases
IT & software development	Code completion, documentation, debugging	Simple software architecture, pull request review	Security, complex architectures
Operations & supply chain	Demand forecasting, logistics optimization	Real-time supply chain monitoring, production planning	Autonomous operational steering agents
Sales & business development	Prospect context, CRM enrichment, competitive intelligence	Communications personalization, lead scoring	Autonomous negotiation, automated closing
R&D / Innovation	Literature synthesis, hypothesis generation, rapid prototyping	Molecular discovery (pharma), generative design (industry)	Autonomous R&D
Executive & strategy	Meeting summaries, board preparation, strategic monitoring	Strategic scenario planning, financial modeling	Strategic decision-making

What this matrix tells us

- **The 'Works' column is dominated by assistance tasks**: drafting, summarization, research, completion. AI excels when it augments a human who stays in control. Enterprise users save an average of 40–60 minutes per day (OpenAI), and professionals can reclaim 26–36% of their time on repetitive tasks (BCG).
- **The 'Depends' column reflects a reliability problem**: the moment AI has to deliver production-quality output — branded content, a complex client diagnosis, a security review — human oversight remains indispensable. 77% of companies express concerns about hallucinations.
- **The 'Not yet' column has one common denominator:** autonomy. The moment you take the human out of the loop, the failure rate explodes. 42% of companies abandoned their AI initiatives in 2025. And 47% of users have already made a significant decision based on hallucinated content.

For SWP Pracitioners

This matrix is a conversation tool with your business leads. When a department head tells you 'AI is going to reduce my headcount needs by 30%,' ask them which column their use case sits in. If it's green, gains are probably real but modest (5–10%). If it's orange, gains depend on the supervisory effort. If it's red, they're probably confusing a proof of concept with reality.

"AI doesn't matter": what are the limits?

In 2003, Nicholas Carr wrote a now-famous article arguing that 'IT doesn't matter.' His thesis: when a technology commodifies, it ceases to be a competitive advantage.
We're making the same bet for Gen AI. All Gen AI services will consolidate around pre-existing provider players, and you'd do well to stay agnostic. Nothing is stable in this space: there will be meteoric rises and bankruptcies, consolidations, partnerships. Taking a long-term view means not locking yourself in with a single vendor — treating the AI layer as a commodity.
Here are a few of our reservations, drawn from our observations.

Technical reservations

- **Inconsistency**: the same inputs should produce the same outputs. But two identical prompts don't necessarily produce the same results, due to the probabilistic nature of these models.
- **Approximation**: when it comes to language, using synonyms isn't a big deal. But when you're delivering a decision-support product based on numerical data, even a minor approximation can have serious consequences.
- **Hallucinations**: the rate varies from 3% to 40% depending on the model and task. Recent models do better, but none are 100% reliable. In an SWP context, a hallucination can distort a 5-year headcount projection.

Security and economic reservations

- **Sovereignty**: many large European groups ban certain technologies or models. Deep integration of OpenAI (US risk) or DeepSeek (China risk) is likely to create contractual and regulatory headaches.
- **Regulation**: the framework is shifting rapidly. The EU has passed the AI Act; GDPR is of course already in play; and as Anita Lettink often points out, the purpose of payroll software or HRIS isn't just functional — it's also bound up with legal requirements.
- **The real cost**: 70–85% of AI initiatives fail to achieve expected results. As noted above, the abandonment rate jumped from 17% to 42% in a single year.
- **The existing stack**: companies are already paying for their own tools. Satya Nadella declared that SaaS was going to die, replaced by generalist AIs. Maybe someday — but other studies suggest we're nowhere near ditching graphical interfaces entirely, which will always be more efficient for ordering a pizza on Deliveroo than typing a request into a chat window.

In a nutshell

- Gen AI is the visible tip of a larger technological wave. Behind the jargon (LLM, tokens, agents, AGI), the mechanism is simpler than it looks: **it's a supercharged T9.**
- Use cases with demonstrated ROI are concentrated in assistance (drafting, summarization, coding, IT). The moment you move to autonomy, the failure rate explodes.
- **Adoption is massive (78%); real impact is concentrated (6% of high performers).** 42% of companies abandoned their AI initiatives in 2025.
- **The right posture is technological agnosticism:** treat AI as an interchangeable commodity, and focus your energy on the human decisions it informs.

The real impact on employment

MUCH ADO ABOUT NOTHING? WHAT THE NUMBERS ACTUALLY SAY

The great fog

What's your reaction when we quote the following study: *'Of 702 occupations representing 100% of US employment, 47% are going to disappear'*?

We can picture you looking weary, perhaps a little irritated, at the thousands of Cassandras who bombard you with this kind of estimate (the famous 'according to a study by the University of Wherever, Wyoming'). Your openness to the 'Impact of AI' topic has no doubt been sorely tested by its omnipresence at HR conferences and seminars. Put simply: the topic is trendy and the doom-laden prophecies are plentiful.

Yet this study isn't recent. It's over 10 years old. And above all, it's not talking about artificial intelligence — it's talking about... digitalization.

Let's take a step back: for 30 years, since Jeremy Rifkin's *The End of Work*, every technological wave has been an opportunity to frighten us with a dystopian society where technology supplants us and buries entire industrial sectors. Y2K, Big Data, digitalization, self-driving cars... good grief, even 3D printers were supposed to kill manufacturing, and blockchain was going to take down the banks! **Every time, our inner homo sapiens — addicted to fiction — plays tricks on us and leads us to wildly overestimate the speed of change in the world around us.**

So if we look back at what should properly be called a bullshit study — 10 years on — what was the actual impact of digitalization? For occupations considered at extreme risk (around 171 roles, 30% of employment), they shrank by 2%. For those considered at high risk (149 roles, 17% of employment), they actually grew by 14%. Hardly the upheaval everyone predicted.

And that's the bullshit study recipe: use emphatic language, and above all give no precise time horizon (that way you're guaranteed to always be right eventually) — or make it so distant that nobody will ever go back and check.
Aren't we doing the same thing with generative AI?

The truth is, nobody knows. Our clients — mostly large corporations that can hardly be accused of lacking foresight and proactivity (they use Albert, after all) — are all testing hypotheses. But none of them are brimming with certainty. Even in scale-ups, some radical visions are questionable. Klarna, for example, dramatically cut its workforce in 2024 and replaced its customer support with AI. Reviews became so scathing they reversed course and ended up rehiring humans. According to the CEO, *'nothing will be as valuable as people.'* No shit, Sherlock...

That said, this time something is different. Not in scale — predictions of mass destruction remain as exaggerated as ever — but in the location of the impact. And that's where the numbers get interesting.

AI isn't killing jobs. It's killing a generation.

Forget for a moment the big studies that 'score' occupations by their AI exposure. Let's look at what employment data is actually telling us. Spoiler: it's not entire sectors pulling back. It's an age group being silently squeezed out.

France: the Insee signals

A note from Insee (the French National Statistical Institute) published in March 2026 paints an unprecedented picture. In the IT and information services sector, salaried employment fell by around 3% between end-2023 and end-2025. So far, nothing spectacular.
But the breakdown by age tells a very different story: **this decline is almost exclusively driven by the under-30s**, whose negative contribution reaches -3.8 percentage points. The 30–54 age bracket, by contrast, shows slight growth (+1.4 points).

The painful number: in Q4 2025, employment of under-30s in IT recorded a **-7.4%** year-on-year fall. For all private salaried employment, that decline is just -0.7%. And all the while, value-added in the sector keeps growing. Companies are producing more, with fewer young people.

Insee is cautious in its conclusions (employment is broadly slowing in France, and IT may be experiencing a cyclical slowdown), but the authors note that *'the reversal is abrupt and coincides with the arrival of generative AI,'* and that *'employment has not fallen due to a lack of activity.'* An important distinction.

United States: Stanford delivers the diagnosis

The Stanford Digital Economy Lab (led by Erik Brynjolfsson, the world's foremost reference in digital economics) published a study in late 2025 based on the payroll records of millions of American workers, across tens of thousands of companies, between 2021 and mid-2025. This is heavyweight research — not a LinkedIn poll.

Their conclusion: in the occupations most exposed to AI, the employment rate of 22–25-year-olds has dropped by around **16%** since the arrival of generative AI at the end of 2022. In software development and call centers, the decline reaches **20%** for entry-level profiles.

And for more experienced workers? Their employment in those same roles has actually increased by **6–9%**. Companies are hiring more 35–49-year-olds in the most AI-exposed jobs. These are the profiles capable of overseeing and directing the tools — not being steamrolled by them.

Everywhere else, same story

The Federal Reserve Bank of Dallas corroborates these findings: the share of employment of 20–24-year-olds in the most AI-exposed occupations fell from 16.4% to 15.5% between November 2022 and September 2025. For other age brackets, employment is stable or rising.

Revelio Labs data shows that job postings for entry-level positions in the United States plunged **35%** between January 2023 and June 2025. The WEF notes a 29% year-on-year decline, while projecting a net creation of 78 million jobs by 2030 (good old Schumpeter, creative destruction, all that jazz). Because yes, you read that right: jobs aren't disappearing — but the entry doors are closing.

In France specifically, a BSI study (2025) reveals that 26% of French executives have already cut or eliminated early-career positions due to AI-related efficiency gains. In the financial sector, that proportion rises to 50%. And a Deel/IDC study from late 2025 indicates that **67% of French companies plan to reduce their hiring of junior profiles over the next three years**. Put that against the 28% of French companies that have an AI training program — and you begin to grasp the scale of the disconnect: we want the benefits of AI, but we're not training people for the disruption.

Key data snapshot

Source	Finding	Key figure
Insee (March 2026)	Decline in employment of under-30s in IT in France	**-7.4% / year**
Stanford (2025)	Drop in employment of 22-25-year-olds in AI-exposed roles (US)	**-16%**
Stanford (2025)	Rise in employment of 35-49-year-olds in those same roles	**+6 to 9%**
BSI (2025)	French executives who have cut junior positions due to AI	**26%**
BSI (2025)	Financial firms that have cut junior positions	**50%**
Deel / IDC (2025)	French companies planning to reduce junior hiring within 3 years	**67%**
Revelio Labs (2025)	Decline in entry-level job postings in the US (Jan 2023–Jun 2025)	**-35%**
WEF (2025)	Projected net job creation worldwide by 2030	**+78 million**

The mechanism: no layoff announcements, just silence

One important point: the adjustment isn't happening through waves of redundancies. It's happening through a silent drying-up of entry-level hiring. Economists call it 'low hiring, low firing.' Companies aren't laying people off en masse — they're simply not replacing junior roles as they become vacant. When a position opens, it's filled by a more experienced profile, or absorbed by AI.

That's what makes the phenomenon particularly insidious. There are no trumpeted restructuring announcements — **just a quietly narrowing range of job postings.**

A ticking time bomb

This is where the SWP practitioner needs to make their voice heard — because what's happening before our eyes goes beyond a simple cyclical adjustment.

The implicit deal of junior employment is dead

For decades, the labor market has run on an implicit deal: the young graduate accepts unglamorous execution tasks in exchange for learning through immersion. They observe seniors, build skills through practice, make mistakes on low-stakes issues before being trusted with the things that matter. It's apprenticeship, corporate-style.

That deal is collapsing. The execution tasks — writing reports, first-level data analysis, building simple features, customer support — are precisely the ones AI automates best. The career ladder is still there, but the first three rungs have been sawn off.

Matt Beane (MIT): *'The way you make a senior employee is not through school. t's by doing the job alongside someone who knows more, and you learn by doing. And that's where the bulk of our skill comes from.'*

Today's juniors won't become tomorrow's experts

This is where SWP thinking comes in. Project yourself 5–10 years out. If companies massively stop hiring juniors today:

- **Who will replace the 45-year-old experts leaving in a decade?** Profiles who will never have had the 10 years of experience needed — because they never got a foot on the first rung.
- **Who will transmit the tacit know-how, the expertise that can't be taught in a classroom but is passed on through mentorship?** No one — because the transmission chain will be broken.
- **Who will have the capacity to judge whether AI-generated output is actually correct?** Not the Sunday vibe coders. Not the casual users. Professionals forged through practice, who have seen enough real cases to tell a plausible result from a correct one.

This is a demographic time bomb. The effects won't be felt this year or the next. They'll manifest in 5–10 years, when companies go looking for experienced profiles and find none — because they will have collectively dried up the talent pool.

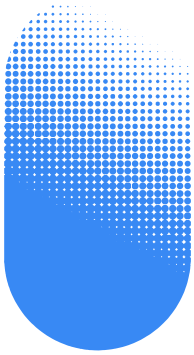

Story Time

Recall the SWP Cookbook: in 2013, an industrial company sets up an early retirement scheme to cut costs. 500 people leave.
The company then wins major contracts and needs to hire urgently in an already-depleted labor market. Six months later, 40% of production fails to meet quality standards.

They had to bring the retirees back and pay them twice over. A perfectly avoidable fiasco with scenario-based projections. Replace 'early retirement scheme' with 'end of junior hiring,' and you have exactly the same movie playing out again.

The paradox for the SWP practitioner

The paradox is brutal: companies using AI to reduce their junior headcount today are creating a future skills gap that they'll have to fill later — probably at a much higher cost.

It's a textbook case for SWP. If you only model a 2–3-year horizon, everything looks fine: you save headcount, you gain productivity. But extend the horizon to 7–10 years, and the available skills curve collapses. And the 42% of global employers anticipating a talent shortage between 2025 and 2030 (WEF, Future of Jobs 2025) may have the right diagnosis — just the wrong culprit.

What this means for your projections

In your SWP exercises, this data changes the picture:

- **Your demographic assumptions need to differentiate junior roles.** Your renewal models must account for this to remain intentional
- **Your medium-term gap analyses are probably underestimating the risk.** If your business needs stay stable and your demographic projections rely on 'normal' workforce renewal, you'll mechanically discover a senior skills gap in 5 years that you hadn't anticipated.
- **Your action plans need to start now.** The question isn't 'how many juniors are we going to hire this year' — it's 'what will our expertise stock look like in 10 years if we reduce junior hiring by 30% over the next 3 years?' Do the math. It stings.

In a nutshell

- **AI is not obliterating entire sectors. Its impact is falling disproportionately on young workers and entry-level roles.**
- 22–25-year-olds are seeing employment drop 16% in exposed roles (Stanford), while 35–49-year-olds grow 6–9%. **AI isn't erasing jobs — it's erasing juniors from jobs.**
- **The mechanism is silent:** no mass layoffs, just a drying-up of junior hiring. In France, employment of under-30s in IT fell 7.4% year-on-year while value-added grew (Insee, Q4 2025).
- **Today's juniors are tomorrow's experts.** Cutting the entry pipeline now means creating a senior skills shortage in 5–10 years. It's the same movie as the early retirement schemes — we just don't see it coming.
- **For the SWP practitioner, this is a full-blown business driver** — one that needs to be modeled, scenario-planned, and tracked with the same rigor as any other strategic driver.

From tasks to outcomes

DOING MORE WITH MORE

The Jevons Paradox, or why developers aren't dead

In 1865, a British economist named William Stanley Jevons published *The Coal Question*. His thesis was counterintuitive: the improvements James Watt had made to the steam engine — making coal more efficient per unit of work — were not going to reduce coal consumption. They were going to increase it. Considerably. Why? Because lower costs made steam economically viable for an infinite number of new uses (factories, locomotives, ships) that hadn't previously been feasible.

160 years later, the same paradox is playing out before our eyes with software development.

The numbers that contradict the narrative

If AI were replacing developers, we should be seeing a contraction in development activity. The same headcount would produce the same quantity of code (or more), and overall demand would shrink. But what's actually happening is the exact opposite.

GitHub's 2025 Octoverse Report is unambiguous:

- **986 million commits** (i.e. blocks of code) pushed in 2025, up 25% year-on-year. Nearly a billion. With a B.
- **43.2 million pull requests** (i.e code modification requests) merged per month on average (+23%).
- **36 million new developers** joined the platform in a year (more than one per second). GitHub now hosts 180 million developers.
- **121 million new repositories created** — an all-time record.
- The US Bureau of Labor Statistics projects **15% growth in developer employment** by 2034. Government statisticians, not known for sensationalism, see more developers — not fewer.

If AI were killing the profession, those curves should be pointing downward. Instead they're exploding upward. This is classic Jevons: **when the cost of producing software falls, you don't produce the same amount of software for less. You produce vastly more software.**

The vibe coder vs. the truly augmented dev

Collins Dictionary named 'vibe coding' its word of the year in 2025. The promise is seductive: describe what you want in plain language, and the AI builds the application for you. Designers build their own tools. Marketers prototype landing pages. Product managers create internal dashboards.

At Albert, we tried it: Lovable or Base44 can generate a complete application from a simple description of your need. It's impressive, and brutally effective if you're a non-technical Muggle looking to quickly validate a proof of concept.

However...

With 500,000 lines of code in Albert, our developers are writing something closer to War and Peace every day. For quickly testing an idea, AI advances are tremendous — but for scaling, following sound development practices, securing and making an application robust, it's a different story. Just look at the quality of code generated by a vibe coder: not robust, not secure, great for a proof of concept, catastrophic for production. NB: Nothing that came out of a Lovable build was ever commercially viable by our standards...

GitHub's figures confirm this indirectly: comments on commits **fell 27%** while commits grew 25%. We're producing more code, but reviewing less of it. Vulnerabilities of the 'Broken Access Control' type (essentially: some users can access content they're not supposed to see) exploded 172% year-on-year, partly because AI-generated code creates endpoints that 'work' but skip authentication checks. 29% of Python code generated by Copilot contains security flaws. In short: vibe coding produces code that looks correct but often isn't. And unlike a boiler explosion in the steam age, AI errors don't make noise. At least — not until credit card numbers end up on the dark web!

Story Time

What does a good developer using AI look like in 2025? They carefully document instructions in CLAUDE.md or .cursorrules files to frame the AI's behavior. They architect their instances (you can run 7 or 8 in parallel), having some monitored by others — which review code quality, run continuous improvement loops, write tests, review pull requests... They spend less time writing code line by line, and far more time specifying, orchestrating, validating, and guaranteeing.

A profession supposedly made accessible to everyone has actually become hyper-technical and complex again. Just differently.

This is exactly the Jevons pattern: the tool is more accessible, more people use it, the total quantity of software explodes, and the need for expertise grows rather than shrinks. The developer pool is broadening at the bottom (the vibe coders, the casual POC-ers), but the top of the pyramid — architects, security specialists, agent orchestrators — is more sought-after than ever.

Ownership over task substitution

Now let's talk about methodology. Consulting firms (themselves heavily disrupted by AI, incidentally) have a dominant approach for assessing AI's impact on roles. It involves:

- Breaking down each role into elementary tasks
- Scoring each task as substitutable, augmentable, or unaffected by AI
- Inferring a percentage of role exposure

It's beautiful in theory. Intellectually elegant, even. But it has a fundamental problem: **substitutable does not mean desirable.**

The task-based approach trap

Let’s take an example. A financial analyst spends 40% of their time compiling data, 30% producing reports, 20% presenting conclusions, and 10% challenging assumptions. The task-based approach concludes that 70% of the role is 'automatable.' You infer a proportional headcount reduction.

Except this reading misses several things:

- Compiling data is also when the analyst spots an anomaly, a weak signal, a problem that no prompt would have known to ask about
- Producing the report is what structures the thinking and lets the employee take real ownership of their conclusions

Those 'subordinate' tasks are precisely where learning happens (back to the previous chapter).

And above all: **just because you can automate a task doesn't mean you should.** The decision depends on context, culture, regulation, and acceptable risk levels.

Business leaders still have free will over their organizational decisions, after all.

What matters is ownership of the output

Our conviction is simple: rather than obsessing over mapping tasks (means), it's more productive to analyze results (ends).

You own your results at work — whether you use AI or not. You are responsible for the output. If AI helps you produce a report in 2 hours instead of 8, great. If the report is wrong because AI hallucinated and you didn't check, that's your problem. The technology used is irrelevant. The quality and reliability of what comes out is not.

This completely changes the SWP lens:

- The task-based approach leads to substitution calculations: *'30% of the role is automatable, let's reduce by 30%.'* It's mechanical, it slides well in a presentation, and it's generally wrong in practice.
- The ownership approach leads to different questions: *'What results do we need? Who is accountable for them? How can AI improve these results?'* Less elegant, but far more realistic.

For SWP Practioners

When modeling AI's impact on a role, don't start from tasks. Start from expected outcomes. Ask operational leads: 'What results do you need, and in what quantities?' It's the same logic as identifying business drivers in classical SWP methodology (see The SWP Cookbook, Chapter 2).

More with more.

There's a dominant narrative pushed by AI solution vendors that boils down to three words: 'do more with less.' Reduce headcount, gain productivity, cut costs.

Our conviction is different. What we observe with our clients — and what macro data confirms — is more of a **'more with more'** dynamic. AI doesn't replace people, even though as we've seen, it restricts employment access for juniors. It amplifies the production capacity of those already there, creating new needs, new projects, new ambitions that weren't economically viable before.

Without betraying any confidences but staying specific, our clients are working with two types of hypotheses:

- **Productivity gains** across all roles involving case management. But no mass redundancies. We're typically talking 5–10% gains in the most aggressive scenarios. It's about absorbing more workload rather than shrinking the operation. In regulated environments especially, human oversight is still required by law.
- **Skills evolution**. Everyone is going to have to 'get on with it.' It's the digitalization playbook, repeated — no different. And it's not without risk: as with digitalization, some people will struggle or refuse to adapt and fall behind. But unlike the digitalization wave, another danger lurks (apparently particularly acute in working-class populations, according to recent studies): a kind of credulity and absence of critical thinking, leading users to take everything Gen AI outputs at face value.

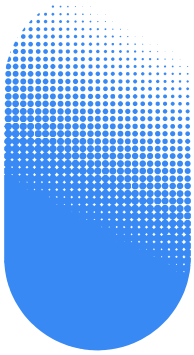

It's a safe bet that the headcount changes driven by AI will be nothing compared to the skills transformation it induces. And that's where the gaps risk widening: between those who know how to use AI to amplify their role, and those who struggle to acquire the new skills. **That's what we need to anticipate.**

Remember the Big Data era: we all suddenly had to become data scientists. Ten years later, the profession is more scarce than ever — while having splintered into a dozen ultra-technical specialisms nobody had anticipated. History has this annoying tendency to stutter.

A new era of productivity

To return to the doomsayers: claiming flatly that one in two jobs is going to disappear is a dangerous shortcut — even irresponsible, because it manipulates our emotions. The substitution scenario — this Luddite 'Great Replacement' — is not coming tomorrow.

What is coming tomorrow is a profound reshaping of how we work. SWP, in its agile and iterative form, is precisely designed to accompany this reshaping. **Just as you don't look after your teeth only when you visit the dentist, but every day by brushing them — you don't take care of your HR hygiene occasionally, but continuously.**

Used this way, SWP becomes a genuine GPS for employment: mapping out several possible routes that update as we move along the road. Planning and agility are not opposites, and combining them is what gets you where you need to go.

In a nutshell

- **The Jevons Paradox is at work:** when the cost of software development falls, the total amount of software produced explodes (986 million commits in 2025, +25%). Demand for developers is growing, not shrinking.
- **Vibe coding democratizes prototyping** but doesn't replace software engineering. Security vulnerabilities in AI-generated code are exploding (+172%). The developer role hasn't become accessible to all — it's just become hyper-technical in a different way.
- **The task-based approach is a trap:** substitutable doesn't mean desirable. What matters is ownership of results. Think in outcomes, not FTE × tasks.
- **The 'more with less' thesis is seductive but false.** What we observe is 'more with more': modest productivity gains (5–10%) and above all a skills transformation that needs to be anticipated.
- **Agile, iterative SWP is the ideal tool for navigating this transformation:** continuously revising assumptions, scenario planning, and not confusing speed with haste.

Integrating AI into your SWP

SO CONCRETELY, IN MY SWP?

The three preceding chapters have laid out the diagnosis: AI is not destroying entire sectors, it's quietly pushing young people out of entry-level roles, and the right framework is outcome ownership rather than task substitution. That leaves the question every SWP practitioner is asking themselves: **concretely, how do I model this?**

This chapter offers five methodology cheat sheets, in the spirit of The SWP Cookbook. But let's start with an illustration.

The generational gap: an illustrative case

Take a fictional company — a 10,000-person international industrial group. Its age pyramid at the start of 2026 looks like this:

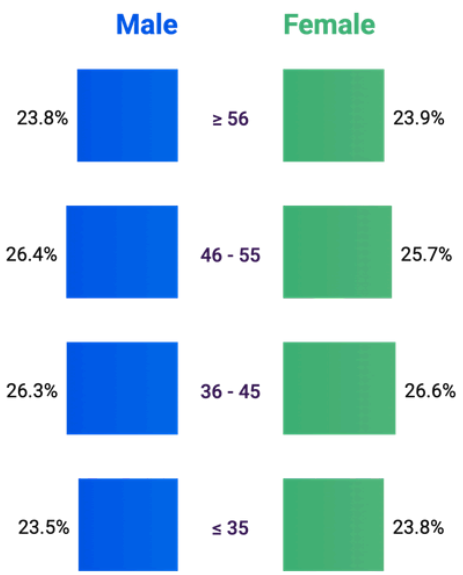

With the following grade distribution (schematic, since it's fictional):

Now imagine that, on top of its usual challenges, the company adopts a specific demand trajectory for junior roles:

- **Scenario A (steady state)**: our needs don't change; we aim to replace departures
- **Scenario B (moderate)**: AI absorbs some junior tasks (roughly half)
- **Scenario C (aggressive)**: AI is seen as a substitute for junior profiles — hiring is frozen. Scenario C's curve mirrors the demographic evolution curve.

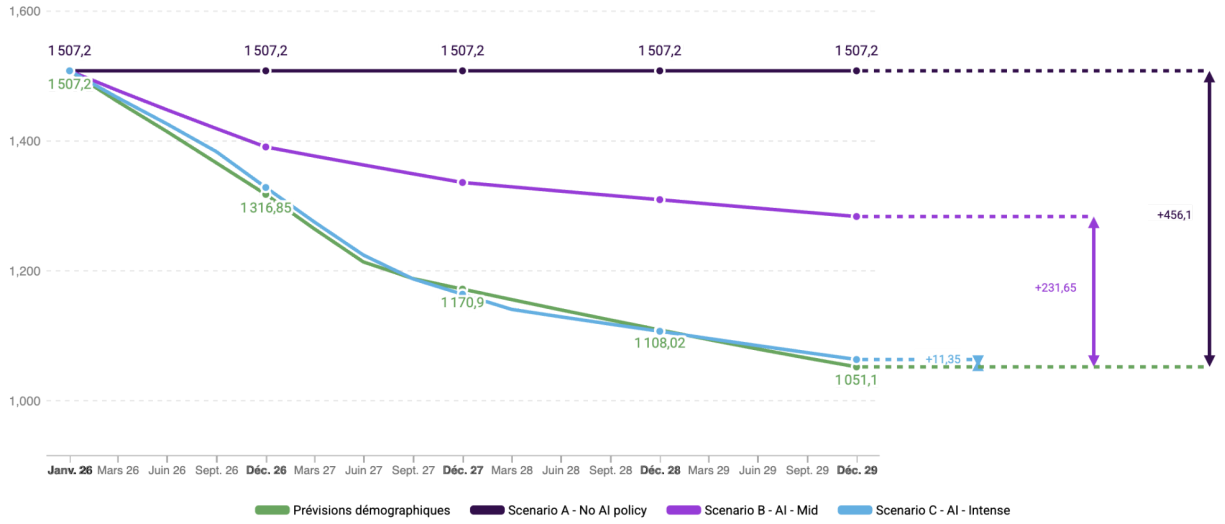

Projection result to 2032

On the surface, at the company level and with equivalent business needs, we do indeed accomplish just as much with less. And overall, the optimization looks fairly measured.

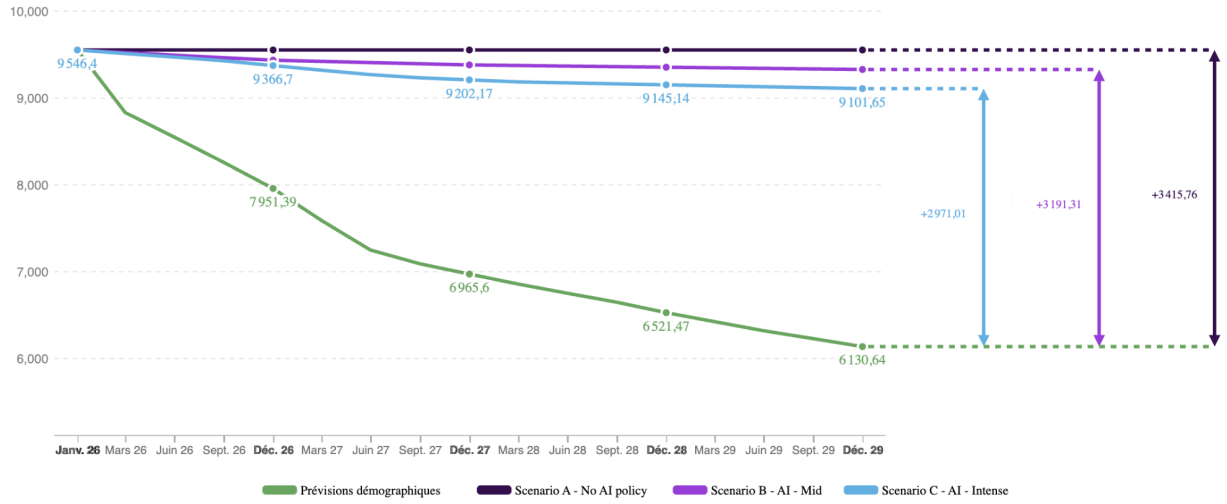

But the slow poison of this approach becomes very apparent when you zoom in on the confirmed professionals category:

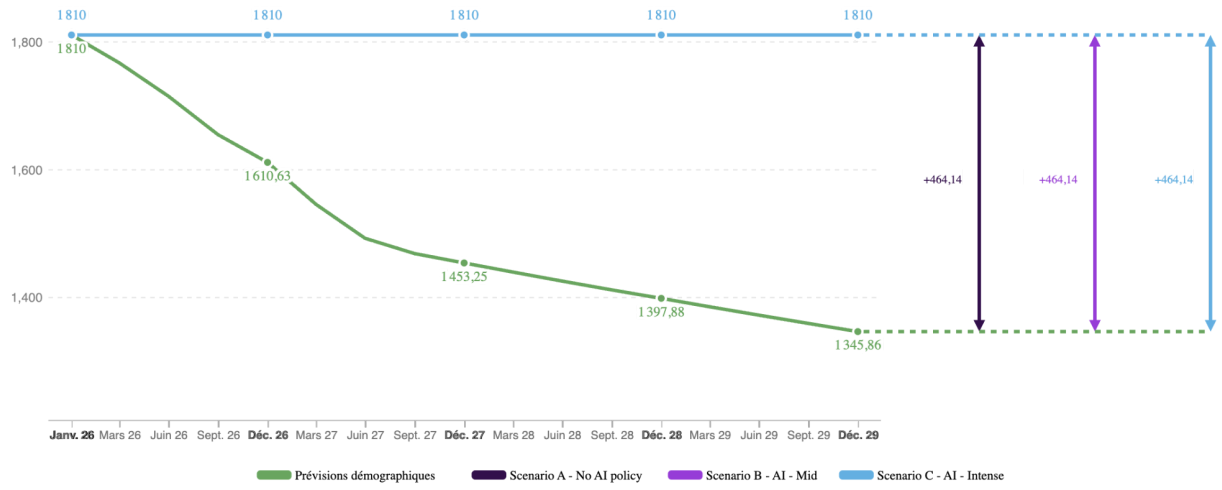

This category is depleting, and is no longer being fed from the rung below. Over time, this population segment is doomed to disappear without active recruitment. And of course, recruiting confirmed profiles is harder than recruiting juniors, generating a shortage that will be difficult to reverse given the inertia of developing new hires.

The key point

Even if Scenario B (moderate) seems reasonable in the short term and flies under the radar in an annual budget review, it is already creating significant gaps in 3 years in the pool of future experts. This is exactly the type of weak signal SWP is designed to detect.

Cheat sheet number 1

IDENTIFYING AI-IMPACTED BUSINESS OUTCOMES

Challenges

The dominant approach — breaking roles down into tasks, scoring the automation of each task, inferring a headcount impact — slides well in a presentation and is generally wrong in practice. Substitutable doesn't mean desirable. Start from outcomes, not means.

Methodology

For each business scope, identify the expected outcomes (outputs): deliverables, objectives, performance indicators. Then ask three questions:

- **Volume**: will AI change the volume of expected outcomes? (for example: more applications to maintain, more content to produce, more cases to process)
- **Method**: will AI change how those outcomes are produced? (an augmented developer produces differently — not necessarily less)
- **Skill**: will AI change the level or nature of skills required to produce those outcomes? (more orchestration, less execution)

Objective

In your SWP approach, this thinking translates into defining business hypotheses within your drivers. You define your unit of measure (revenue, production volumes, number of applications, cases processed), calibrate volume evolution against it, and apply a non-linear elasticity to capture the efficiency AI brings. Infinitely richer than a blunt '-5% productivity' applied across an entire job family.

Example

IT manages 200 applications. The strategic plan calls for managing 280 by 2028 (+40%). With AI, each developer can manage 20% more. Result: instead of recruiting 160 additional people, you recruit 115. That's 'more with more,' not 'the same thing with less.'

Cheat sheet number 2

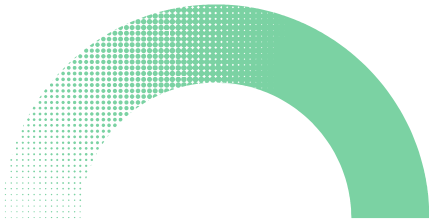

BUILDING YOUR AI SCENARIOS

Challenges

AI should be treated like any other business driver — with its variants, integrated into scenarios. Not as a global parameter applied uniformly.

Methodology

Build three intensity levels, following the criteria for a relevant driver (see The SWP Cookbook Chapter 2):

- **Conservative scenario**: AI is used in assistance mode. 5% productivity gains on exposed functions. No headcount reductions — time gained is reallocated to higher value-added tasks. This is the scenario most frequently observed with our clients today.
- **Moderate scenario**: AI is integrated into workflows. 10–15% gains. Reduction in entry-level hiring, upskilling of existing profiles. Recomposition of certain teams.
- **Ambitious scenario**: AI is deployed at scale with agents. Some functions are fundamentally redesigned. 20% gains. High risk of a generational gap in the medium term.

Objective

Each scenario is modeled as a distinct driver, with its start date, duration, and differentiated impact by job family. The point of scenario planning is precisely to be able to compare trajectories and revise them as reality clarifies. Remember: SWP is a GPS, not a prophecy.

Cheat sheet number 3

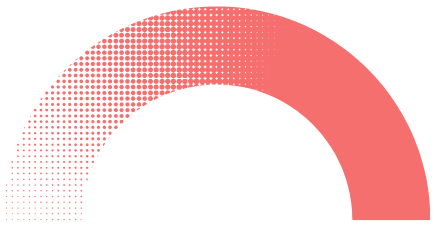

MODELING THE DRYING-UP OF THE JUNIOR PIPELINE

Challenges

Classic renewal models are built on exit assumptions (turnover, retirements, mobility). The AI impact introduces a new parameter: entry assumptions. If you reduce your junior needs by 30% over 3 years, your pool of experienced profiles in 5 years collapses.

Objective

In your SWP tool, model demand assumptions differentiated by age bracket and job family. The demographic projection incorporates these flows into the headcount evolution calculation, making the 'gap' visible well before it materializes.

Methodology

Three steps:

1. **Map your current cohorts** by job family and age bracket. What is their 3-year retention rate?
2. **Scenario-plan the junior reduction**, linked to your AI scenarios (Cheat sheet 2). Even a modest reduction (20–30%) has considerable cumulative effects.
3. **Project the impact on the senior pool** over 3–5 years. That's where the gap appears. Cross-reference this projection with your business needs at the same horizon: the gap is your early warning signal.

Cheat sheet number 4

INTEGRATING AI SKILLS INTO THE COMPETENCY FRAMEWORK

Challenges

Everyone is going to have to 'get on with it' — or so they said about digitalization. We're running the same playbook for AI. But as with digitalization, the skills to integrate are not the ones you'd expect. The reflex is to think 'technical' (knowing how to code an agent, fine-tune a model). In reality, the most critical skill is far more mundane: knowing how to assess whether what AI is telling you is actually true.

Methodology

Rather than creating an absolute AI competency framework, map AI skills onto your existing job families — as you did (or should have done) for digitalization. Concretely, for each job family, ask three questions:

- **What interaction skills?** Writing an effective prompt, providing context, iterating on a result. This is the baseline — the 'knowing how to use a computer' of 20 years ago. Applies to 100% of employees eventually.
- **What judgment skills?** Verifying an output, detecting a hallucination, cross-referencing with other sources, knowing the regulatory limits (AI Act, GDPR, confidentiality). This is the most underestimated and most important skill. Without it, your employees take everything the machine produces at face value (and 47% already have, don't forget).
- **What integration skills?** Setting up an assistant on an internal document base, automating a workflow, assessing the relevance of an AI tool for a given use case. This applies to business leads, managers, and obviously technical profiles (who will go all the way to agent orchestration and security).

Don't aim to be exhaustive. Remember The SWP Cookbook: a good competency framework doesn't exceed a few hundred skills at multinational scale. For AI, three skills per job family (one per question above) is enough to start. You'll refine as you go.

Objective

These skills integrate into your existing framework like any other competency. Flag them as 'emerging' or 'critical,' and attach target levels by job family. The gap analysis that follows allows you to size the training effort — and above all, highlights the chasm between the 67% of companies planning to cut junior hiring and the 28% with an AI training program. Spoiler: it's often the same ones.

Cheat sheet number 5

ADAPTING YOUR ACTION PLANS WITH THE SEVEN B'S

Challenges

The SWP Cookbook presented the 6B method for closing gaps (Build, Borrow, Bridge, Buy, Bind, Bounce). These levers remain perfectly valid in the face of AI: you'll need to train massively (Build), attract augmented profiles (Buy), retain those who have mastered AI (Bind), create bridges from roles whose content is evolving toward augmented positions (Bridge), and manage transitions when necessary (Bounce). We won't revisit these levers in detail — they're already covered in The SWP Cookbook.

That said, AI invites us to add a seventh B: **Boost**. And it's the one worth dwelling on.

Methodology

The Boost is a genuine organizational transformation.

When an AI-augmented employee can produce 20% more than before, the question isn't 'do I cut 20% of headcount?' The question is 'what do I do with that extra capacity?' And the answer involves org design choices:

- **Scope expansion**. An HRBP covering 200 employees can, with AI assistance for report drafting, data summarization, and interview preparation, cover 280. That doesn't mean reducing the number of HRBPs. It can mean redeploying some of them onto higher-value missions (transformation support, strategic talent management) that nobody had the bandwidth to cover before.
- **Team recomposition**. In a team of 10 analysts who spent 60% of their time compiling data, the Boost doesn't mechanically lead to 6 analysts. It can lead to keeping 10 analysts who now spend 60% of their time actually analyzing (rather than compiling) — producing insights the company simply didn't have before. That's 'more with more' in its most concrete form.

- **Management ratios**. If your employees become more autonomous thanks to AI (because they have access to synthesized information, preliminary analyses, generated documentation), span of control can evolve. A manager overseeing 8 people can oversee 12 — not because you're cutting managers, but because the nature of management changes: less execution oversight, more strategic steering and skills development.
- **Organizational performance**. The Boost, well thought out, isn't measured in FTE reduction. It's measured in team output: revenue per employee, cases processed, service quality, time-to-market, innovation capacity. This is a shift in the SWP lens: instead of reasoning solely in 'how many people do I need,' you reason in 'what production capacity do I need, and how do I organize my teams to achieve it.'

Objective

The Boost is the lever that concretely translates the 'more with more' thesis. In your SWP approach, it's modeled through business drivers (see Cheat sheet 1): expected production volume increases, elasticity reflects the capacity gain AI brings, and the resulting headcount trajectory is neither a mechanical freeze nor a budget cut — it's an intentional recomposition. It's also the lever that requires the most dialogue with operational leads, because it involves organizational choices that HR can't make alone.

In a nutshell

- **The generational gap is a real and modelable risk.** Even a moderate scenario (-30% junior hiring) halves the pool of future seniors within 7 years.
- **Start from business outcomes, not tasks.** Define your AI drivers in terms of production volume and elasticity — not a blanket productivity percentage.
- **Scenario-plan AI like any other driver.** Three intensity levels, revised regularly. SWP is a GPS, not a prophecy.
- **Model age cohorts.** The drying-up of junior employment is the major weak signal of this decade.
- **The 7 Bs are your toolkit.** The Boost completes the classic arsenal. It's the lever that embodies 'more with more.'

Conclusion

AI, AND THEN WHAT?

In the 2010s, companies paid seniors to leave and hired cheaper younger replacements. Early retirement schemes, collective voluntary redundancies, PDVs — the whole arsenal was mobilized to rejuvenate headcounts and reduce payroll. We know how that ended: a hemorrhage of expertise, emergency recalls of retirees, and recruitment costs that exploded.

In the 2020s, we're doing exactly the reverse. We're keeping the seniors (they're expensive, but at least they know their stuff), and stopping junior hiring in the hope that AI will do the job for them.

The irony is delicious. The outcome will be the same.

In 10 years, companies that turned off the junior hiring tap will be desperately searching for experienced profiles. They won't find them, because nobody will have had the chance to become one. They'll discover that AI still can't judge the relevance of its own outputs, that a vibe coder doesn't replace a software architect, and that mentorship has no algorithmic equivalent. They'll end up paying a premium — recruiting seniors at compensation levels their salary grids never anticipated, or outsourcing skills they should have cultivated in-house.

This future is entirely avoidable with SWP. Because AI doesn't change the need to plan. It makes it even more urgent — because the consequences of inaction won't be visible this year, or the next. They'll manifest in a decade, when it's too late to rebuild what was allowed to erode.

Our biases invite us to look at the AI wave with a mixture of fear and fascination. The debate crystallizes around the potential destruction of jobs — all the more unsettling because it could threaten 'intellectual' professions that were usually considered safe. But this debate distracts us from what should be at the heart of our concerns: rather than asking how many jobs AI will destroy, perhaps we'd do better to worry about how many will not be created.

Your move, workforce planners!

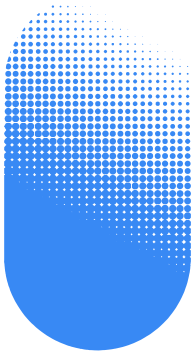

About Albert

WHO ARE WE?

Albert is revolutionizing Strategic Workforce Planning (SWP) for enterprises by eliminating guesswork and complexity. Our AI powered SaaS platform simplifies SWP, empowering businesses to plan and adapt with confidence. Unlike traditional spreadsheets or BI tools, Albert seamlessly handles the intricacies of data science and real-time collaboration, helping organizations manage business transformations and optimize their workforce with ease.
Albert is the first challenger to traditional human resources planning practices. Albert was founded in 2021 by Vincent Barat, an engineer who graduated from CentraleSupélec with a ten-year background in SWP, Jérôme Soulard, who graduated from Dauphine with a strong background in HR consulting and Cyril Siman, our CTO with 20 years of experience in Software Engineering. Our AI based SaaS Platform enables companies to plan their human resource needs, which account for 60% of their expenses. Albert currently has 15 employees and has successfully gained around twenty clients, present in 22 countries, most of which are part of the SBF 120.

Sources

- McKinsey (2025) *The State of AI 2025.*
- BCG (2025) *AI Radar 2025.*
- Menlo Ventures (2025) *The State of Generative AI in the Enterprise.*
- Deloitte (2025) *The state of Generative AI in the enterprise.*
- OpenAI (2025) *The state of enterprise AI.*
- Google Cloud (2025) *The ROI of AI 2025.*
- MIT / NANDA (2025) *The GenAI Divide: State of AI in Business 2025.*
- Insee (2026) *Note de conjoncture (Adjerad & Vermersch).*
- Stanford Digital Economy Lab (Brynjolfsson, Chandar and Chen) (2025) *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.*
- Federal Reserve Bank of Dallas (2026) *Young workers' employment drops in occupations with high AI exposure.*
- BSI (2025) *Evolving Together: AI, automation and building the skilled workforce of the future.*
- Deel / IDC (2025) *AI at Work: The Role of AI in the Global Workforce 2025.*
- Revelio Labs (2025) *Is AI responsible for the rise in entry-level unemployment?*
- World Economic Forum (2025) *The Future of Jobs Report 2025.*
- PwC (2025) *The Fearless Future: 2025 Global AI Jobs Barometer.*
- CBS News (2025) *New study sheds light on what kinds of workers are losing jobs to AI.*
- France Info (2026) *Les jeunes ont plus de mal à trouver du travail dans la programmation.*
- Xerfi, Passet, O. (2023) *D'Internet à ChatGPT : 30 ans de fadaises sur la fin du travail.*
- GitHub (2025) *Octoverse 2025 Report.*
- U.S. Bureau of Labor Statistics (2025) *Occupational Outlook Handbook, Software Developers, Quality Assurance Analysts, and Testers.*
- METR (2025) *Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity.*
- Collins Dictionary (2025) *Word of the Year 2025: "vibe coding".*
- Faros (2025) *The AI Productivity paradox.*
- Jevons, W.S. (1865) *The Coal Question.*

Thank you!

OF HAMMERS, NAILS, AND KITTENS

People often ask me whether I'm for or against AI. The question doesn't really mean anything: it's like asking whether I'm for or against a hammer. It depends on what you do with it. For driving in a nail, I'm pretty much for. For bludgeoning kittens, less so.

What matters isn't the tool but the use. And it's on the use that we have a real responsibility: the responsibility of being opinionated, a stance we're willing to stand behind, rather than just shrugging and letting AI happen to us. That's what this guide has tried to argue.

The world keeps telling us, keynote after keynote, that "AI is going to change everything," that "AI is going to replace this," that "AI is going to wipe out that." It's not true. AI is going to do exactly nothing. YOU are going to decide what to do with it. Just as you decided what to do with digital twenty years ago, with the cloud ten years ago, and with big data five years ago. Technology has never had an agenda of its own: humans give it one. So it's on you to stay intentional about what you decide.

And while we're at it, let's not forget something that's becoming almost blasphemous to say out loud: just because a task can be automated doesn't mean it should be. We all know someone who genuinely loves ironing their laundry. There's real pleasure to be had in the simple, repetitive task, and a real intellectual breath in doing things "by hand," even when an agent could perfectly well take care of them. We can't presume to know what someone finds rewarding in the craft of a gesture well done, or what feels almost meditative to them. What is a productivity gain for one person is sometimes a quiet loss for another.

As always with us, we'd rather bring some nuance to the conversation. It's less sexy, and it does work against us at times. But it's probably also a little more honest.

And that, in the end, is what I'd like you to take away from this guide: AI is not going to write your story. You will. Either with it, alongside it, or against it. But you will write it.

albert

Linkedin
@getalbert

Email Address:
hello@albertapp.com

Website:
www.albertapp.com

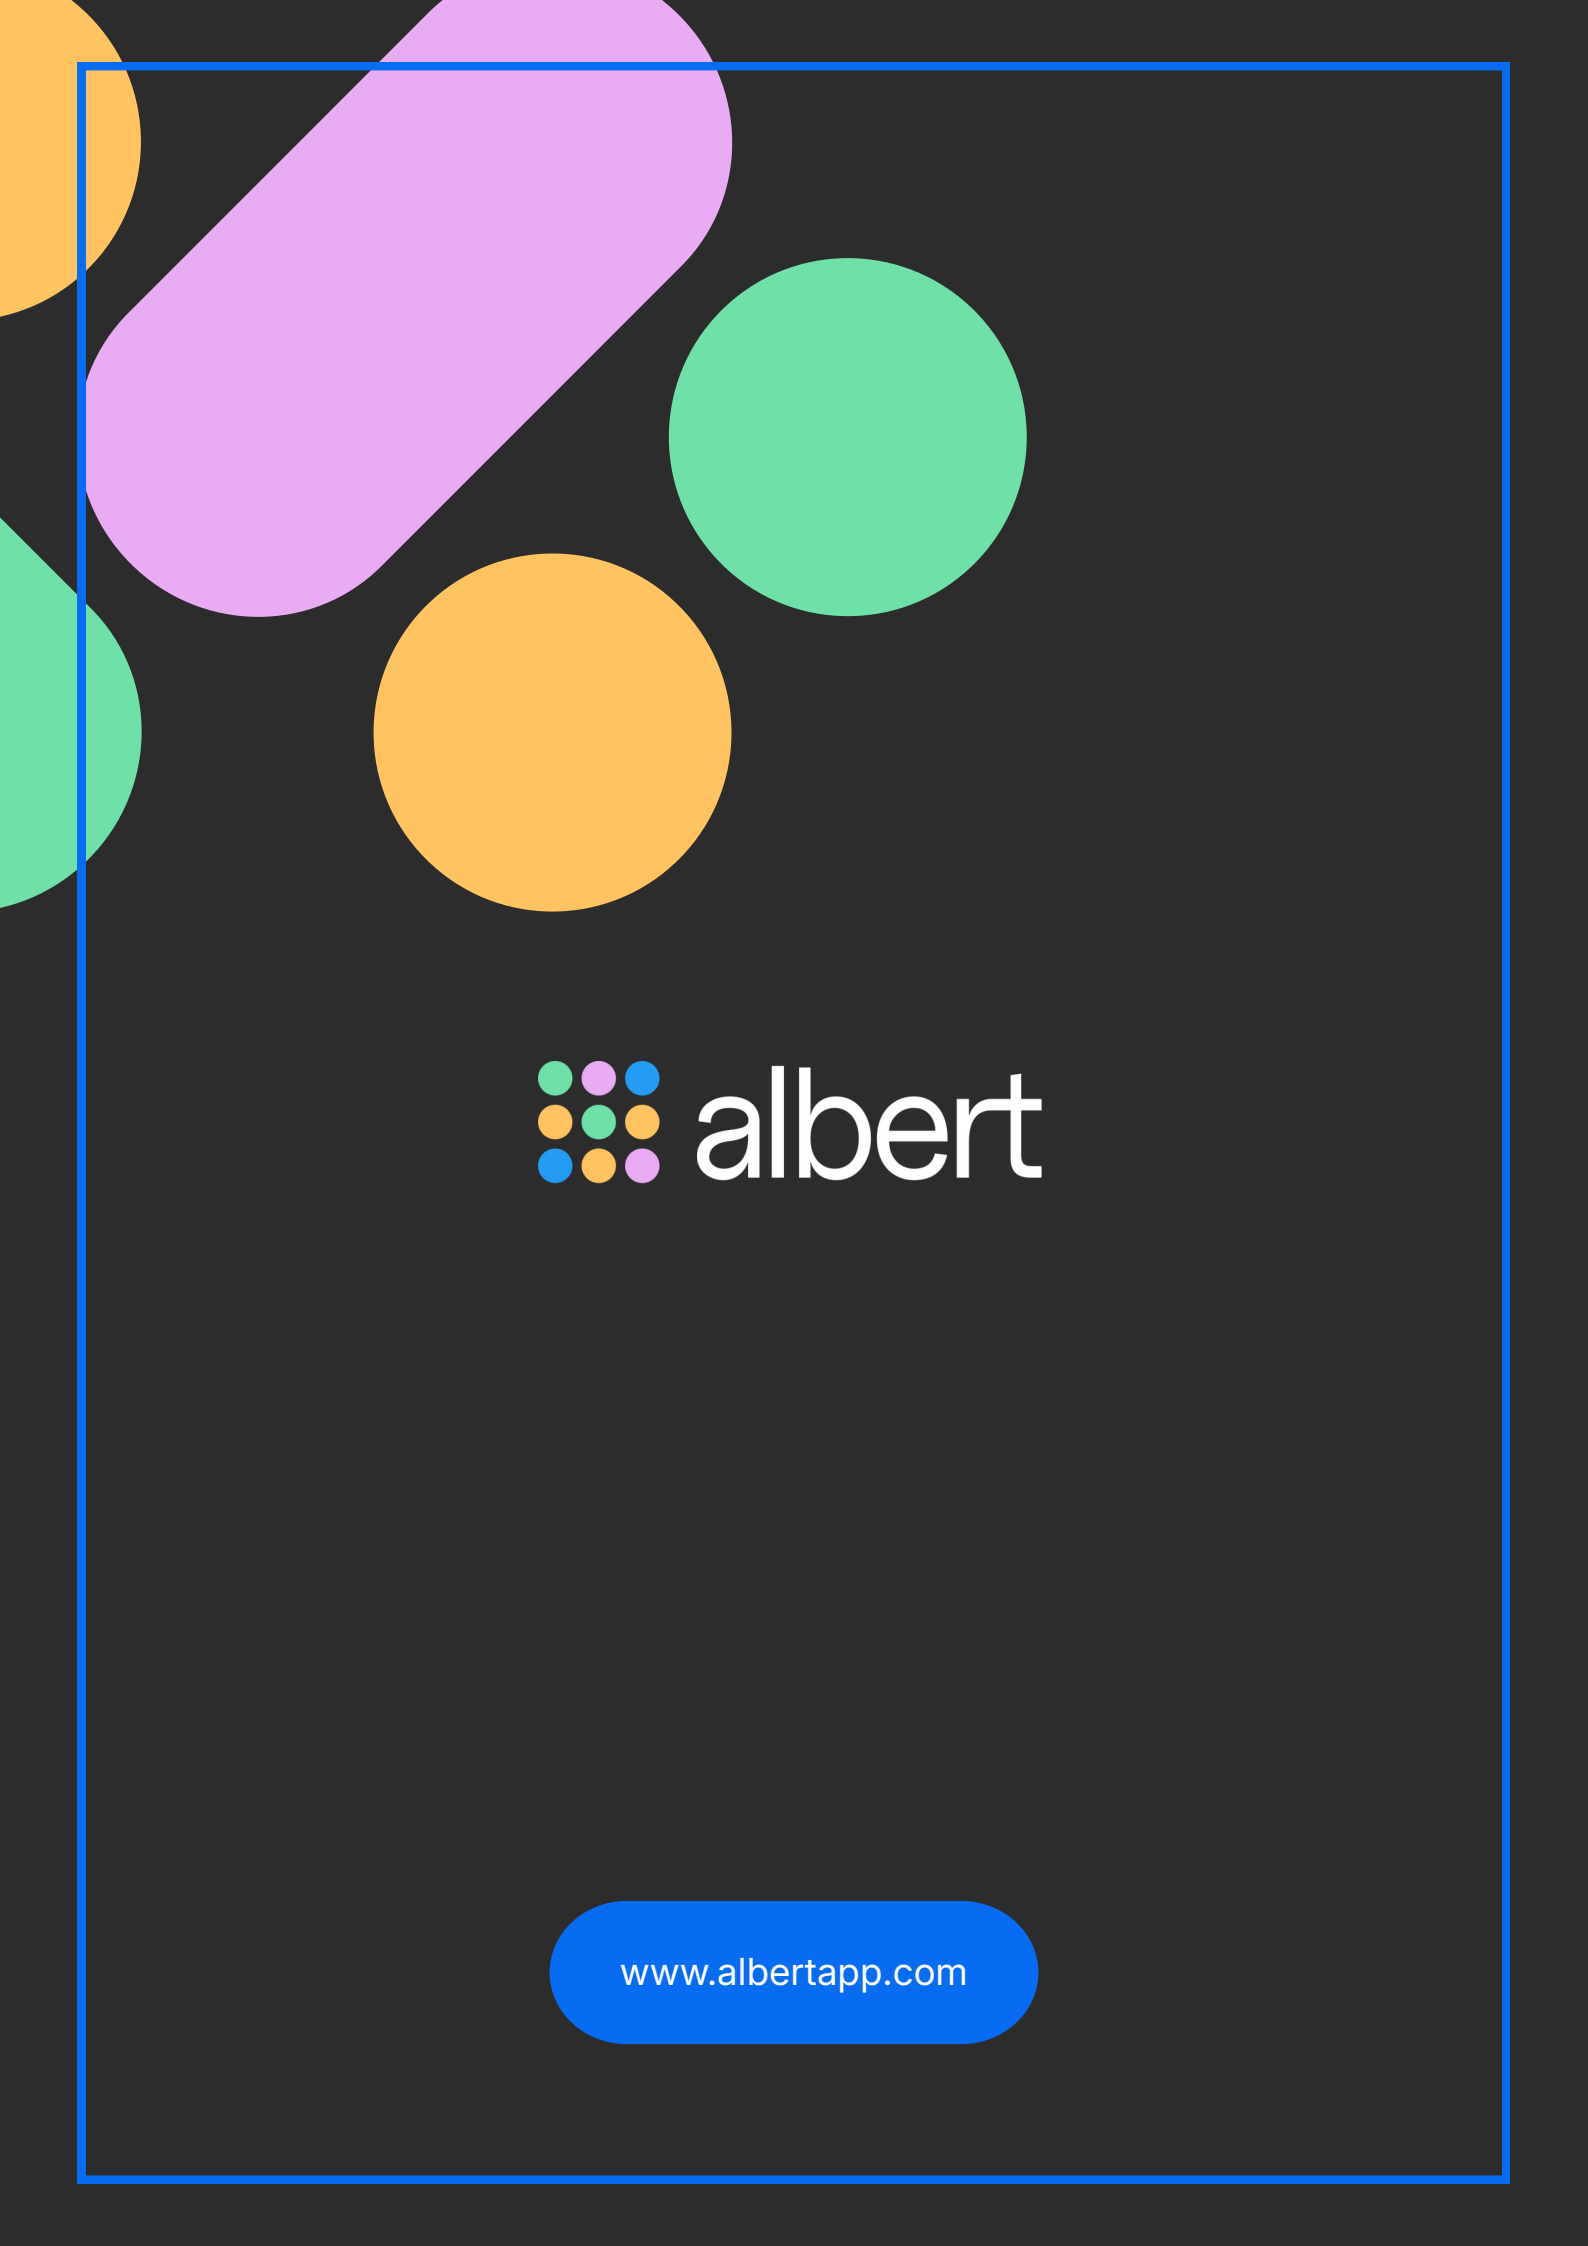
albert
www.albertapp.com