95% of AI Pilots Never Reach Production. The 5% That Make It Share One Habit.

An Article of the Swiss Association of MBAs

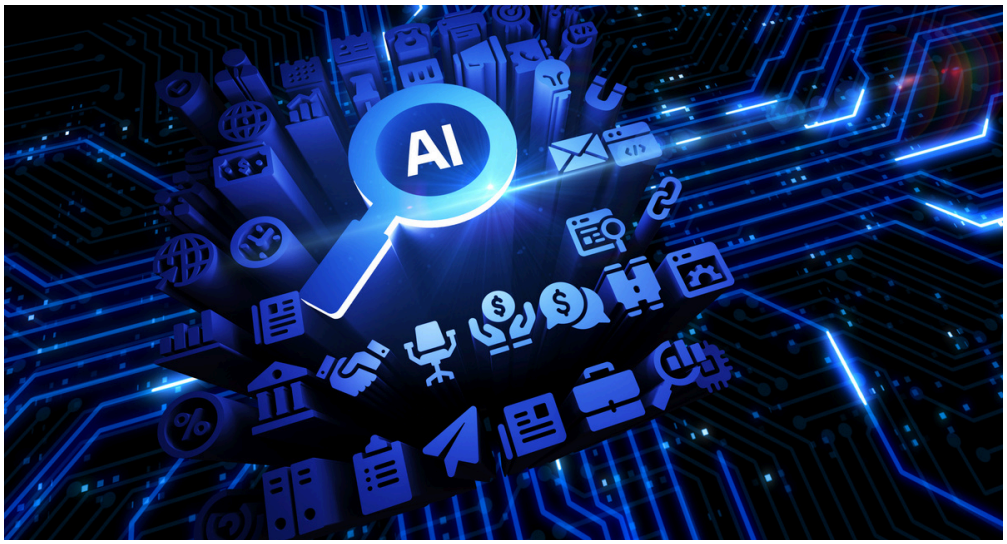

Salama Belghali
Author

You have read the number. Somewhere between 70 and 80 percent of corporate AI projects never leave the pilot phase. It is the kind of statistic that makes a leadership team nod, file it under “good to know,” and move on.

At the first AI Roundtable of the SwissMBAs AI Council, the number got worse. “From our point of view, it is 95 percent,” said Urs Bucher, Chief Marketing Officer at Cognizant Moment. “Five percent survive and get into production.”

That correction set the tone for the hour. I hosted on behalf of the AI Council, and I had invited two people who build this technology for a living rather than talk about it. Urs sits where enterprise pilots either scale or quietly die. Rupert Breheny spent sixteen years as an engineer at Google and now runs the advisory Cobalt AI. He calls himself a skeptical engineer and means it as a compliment.

What follows is not the headline version. It is what they told a room of Swiss leaders about the AI decisions that cannot wait.

The wall is not technical. It is a project.

The first myth both speakers dismantled is that pilots fail because the technology is not ready. It almost never is the technology.

“In the end, this is a full-fledged project,” Urs said. “It is not only AI. It is a whole framework you need to build for the project to succeed in a working (enterprise) environment.” His three rules sound unglamorous, which is the point. Invite compliance into the very first meeting, not the last. Map the whole environment the pilot has to connect to, the systems feeding it and the systems displaying its output. And train everyone, because handing your people a Copilot licence is not the same as teaching them to use it.

Rupert framed the same truth through the human side. “Culture eats strategy for breakfast,” he said. Every one of these projects needs several teams to agree. The sales lead promises the solution in a month. The engineer says it will take three times as long and cost twice as much. Legal says they will not touch it. Shipping something that keeps all three happy is the actual work. “It is a marathon, not a sprint. We have all heard the starting gun, but we have to pace ourselves.

The proof sat in one of Urs’s own engagements: a marketing team inside a major German airline group. Today 75 percent of that team uses the tool, and they use 70 percent of its features. The projected saving for the year is half a million. What moved it from pilot to production was not a better model. It was compliance on board early, a small proof of concept killed after two weeks and rebuilt with the infrastructure team in the room, and the discipline to run it like any other project.

Data is the part everyone underestimates

Ask a room what the hardest part of an AI project is and most will say the model. Both speakers said the data.

“The underestimation was how much time it takes to get the right data in,” Rupert said. For the heritage and luxury brands he works with, the output has to carry the brand’s DNA, which means gathering archives that are often large, rough, and badly labelled. Garbage in, garbage out, as the saying goes. And the first output is rarely the one you ship. “Everybody can be an art director now,” Urs added, quoting one of Switzerland’s most awarded designers, “but a mediocre designer using AI still produces mediocre results, only faster.” The new skill is taste: knowing which output to reject.

The sharpest exchange of the session came from the floor. If the data is the bottleneck, asked one participant, why wait to clean it the old way? Why not use AI to structure the data itself? Urs answered without a pause. “Of course. You clean up your data, but you use AI to clean it. Anything else would be foolish.” The lesson is not “fix your data, then start.” It is “start, and let the tools help you fix the data as you go.

Speed, control, cost: pick two

The platform question came down to a trilemma. You want speed, which points you to the hyperscalers. You want control, which points you to building and running your own. You want low cost, which points you to the open models coming out of China at a fraction of the price. You cannot have all three. Every leader has to decide which one they are willing to give up.

Three practical signals stood out.

The bill is real. “How many tokens did you burn today?” is Urs’s favourite question to his developers, because there is no free lunch. Frontier models are priced below cost today to win the market. As that pricing normalises, both speakers expect AI costs to rise by two to five times over the next couple of years. Budget for it now, and cap it in any multi-year contract.

Do not put all your eggs in one basket. Rupert pointed to outages and geopolitics in the same breath: cloud regions go down, and the lines being drawn through semiconductors, token access, and cross-border data are only getting harder. Resilience means more than one option.
Sovereignty has a Swiss answer, and it is a short one. Asked which sovereign option runs in a Swiss data centre, Urs did not hesitate: “There is exactly one, by Swisscom. Full stop.” It will not be the most advanced model every week. If sovereignty is your hard constraint, that may be the trade you accept.

The risk hiding in plain sight: shadow AI

The most useful warning was about your best people. The ones who have made themselves dramatically more productive with AI at home are the ones most likely to bring consumer tools into work, where they may feed confidential information into a system that uses it for training. “You accrue the benefit personally,” Rupert said, “but the corporation takes on the risk.” Shadow AI is not a fringe problem, and it reaches into senior management. The answer is not to forbid the enthusiasm. It is to give people safe, compliant tools that do what they already know AI can do.

What to decide before the end of June

I asked both speakers for the one decision a leader should make this quarter to land on the right side of the hype. Neither said buy a platform.

Rupert said: dive in and create something. Prompt engineering is table stakes now. Treat these tools as conversational advisers rather than one-shot machines. And pick the lowest-technology solution that solves the task, because it will be cheaper, better tested, and less fragile than the fashionable one. “Do not walk before you run. But jump in, because your competitors already have.”

Urs reached for a German word: begreifen, which means both to grasp something with your hands and to understand it. Get your hands dirty. Set up one small agent to do one small thing. Then look honestly at where your organisation sits on the AI literacy curve, and make a plan to climb one step. Being a first mover in your industry starts there.

Five questions to bring to your next AI pilot

1. Is compliance in the room from the first meeting?
2. Have we mapped every system this needs to connect to?
3. Who owns data quality, and can AI help us improve it as we go?
4. Have we modelled the token cost at two to five times today’s price?
5. Are we training everyone who will touch this, not only the team that built it?

The AI Roundtable is a new format from the SwissMBAs AI Council, built to give our members the view from inside the work rather than the view from the headlines. My thanks to Urs Bucher and Rupert Breheny for an hour of exactly that. The decisions are already on the table. The only real mistake is leaving them there.

About the Author

Salama Belghali is the founder of the AI Career Accelerator. She helps high-impact leaders and senior professionals navigate the shift from corporate roles to independent authority using AI leverage.

Linkedin: https://www.linkedin.com/in/salamabelghali/