How companies are actually using AI

Findings on how they adopt and govern it

The AI Chat interview series · Oli Woodman, Beema Solutions · July 2026

In short

Over the past couple of months I've sat down with senior and founder-level people across a dozen industries to talk about how they're really using AI at work, not how the internet says they are. The short version: the adoption is real, but almost all of it still happens inside a chat window.

Three things came through. Where companies actually are with adoption, how they think about regulation, and a quieter one running under both about what happens when you hand off the thinking as well as the task.

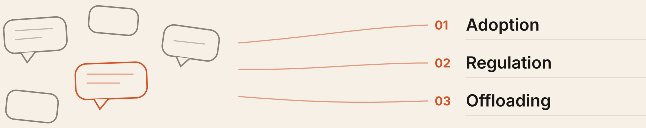

1. Adoption — it still stops at the chat window

Almost nobody is anywhere near the agentic AI the internet keeps promising. Most companies have bought a batch of Copilot or ChatGPT licences, handed them to a small group of people, and stopped there, using them to research, summarise and draft. Very few have taken the next step of letting AI run a task on its own, and fewer still have automated a genuine slice of a workflow. One advisor puts most of the companies he works with on *"step one or two of a ten-step journey"*; another, across roughly twenty-five companies, hasn't seen a single one hand a real chunk of work to AI and let it run. The wins that do exist are quieter than the headlines suggest, a monthly process that took four hours now takes forty-five minutes, which is real but a long way from the whole roles we keep being told are already gone.

Part of the reason is that people are trying to skip the queue. As one consultant put it, *"they want to drive the Ferrari before they've read the driving theory test"*, asking for autonomous agents before anyone's been taught to prompt or the AI's been connected to their own email and files. What you see on LinkedIn is either hype or the top one percent, and it sets an expectation the ground can't meet yet.

Which is why ROI barely comes up, and when it does nobody can really measure it. The gains are scattered minutes spread across a lot of people, and the licences are shared, so no one can cleanly say which team or project the benefit belongs to. At this stage that's fine. Arguing about the return on a monthly seat licence is a distraction; the cost is small, and the token-cost anxiety you read about is a big-company problem that simply doesn't bite here. ROI becomes the real conversation later, when someone is spending serious money building agents and has to justify it. Almost nobody is there yet.

2. Regulation — the rules aren't the problem

This was the part that surprised me most, in a good way. The people I spoke to in regulated sectors weren't overwhelmed by regulation at all. The rules that already exist, GDPR, DORA, the UK conduct regime, were described as sufficient for AI and, in one person's words, *"actually quite reasonable"*. Nobody was waiting on new AI law, and the EU AI Act barely came up. They understand where the line sits, too: you can't tell a regulator *"the AI gave me the answer"*, so they keep it away from the decision itself and use it to do the work around one. That's a grown-up read of the situation rather than a fearful one.

Where it gets shakier isn't the regulation, it's understanding what actually happens to the data, and that cuts both ways. Some are too cautious: IT blocks a tool it doesn't understand, so people go around it and use AI on personal accounts, which is far riskier than the thing that got blocked. Others aren't cautious enough. One example stayed with me. An advisor has AI transcribe customer calls; a customer rings to change their bank details, and the AI quietly writes their name, address and account number into a document sitting on SharePoint, outside every control the core banking system would have applied. Same data, no protection, visible to people who could never see it in the system it came from. And underneath all of it a subtler trap: ticking *"don't train on my data"* is not the same as your data not being processed, and very few people know the difference.

3. Offloading — handing off the judgment, not just the work

Running under every one of these conversations was a quieter worry, and it's the one I keep coming back to. When you hand a task to AI, it's easy to hand over the judgment that used to come with it, and the judgment was usually the point.

You can already see it in small ways, people turning up to a meeting with work they'll disown the second it's questioned (*"oh, Copilot did that bit"*). And in the long run the same people raised a bigger version of it: today's senior people can smell when an answer is off because they built that instinct the hard way, over years without AI. Stop training juniors on the bet that AI now fills the gap, and in twenty years there's nobody left who can look at a plausible answer and say *"that doesn't quite stack up, and I can't tell you why"*. One regulator put it most plainly about his own work: *"the intellectual exercise required to perform my work is vital to perform my work. I don't see outsourcing this as a win, even if it means a drop in productivity."*

What this changed for me

First, thank you. These conversations were generous with time and honesty, and this write-up only exists because of them.

They've also sharpened what I want to do next. The gap I keep seeing isn't a missing tool, it's capability. Most companies already own more than they've learned to use. So the consultancy I'm building is about putting that capability inside organisations and their people, rather than bolting on another software solution for them to maintain, or another external dependency.

I'll keep running these interviews on a rolling basis, so if you'd like to take part, please reach out. And if any of this landed, or you'd disagree with it, I'd love to keep the conversation going.