

THE SHIFT THIS MONTH

## The death of the *dashboard*.

Agentic AI is ending the dashboard as we knew it. But the surface that replaces it isn't a chatbot. It's a hybrid workspace where humans and AI operate side by side, and most fintechs are about to get this transition badly wrong.

## WHAT SHIFTED

# The dashboard was built for a human operator. That operator is no longer alone.

**F**or two decades, the dashboard has been the center of gravity in financial software. A human operator lands on it, comprehends the status of things, and acts. That is why dashboards filled up with visualizations, decision support and action controls. But the whole metaphor rests on one assumption: that a human is the primary operator. In 2026, agentic AI is quietly dissolving that assumption, and taking the dashboard with it.

The market is crossing from generative AI that answers to agentic AI that acts. Financial institutions are moving from AI "assistants" to AI "operators", systems that approve, reconcile, flag and transact through multi-step workflows. When software takes over part of the operating, the surface designed for a human-only operator no longer fits the way work actually happens.

The lazy conclusion, the one most of the industry is reaching for, is that the dashboard simply gets replaced by a chatbot. Bolt on a conversational layer, let users type what they want, declare the interface "AI-native." This is wrong, and expensively so. It mistakes a change in who operates the system for a change in how you talk to it.

The truth is more demanding. We sit at the very beginning of an AI-to-trust continuum. Where a product

sits on that continuum determines the division of labour between human and AI. Some tasks the AI now performs invisibly in the back end, removed from the operator's view entirely. Some it performs, but a human must review before it is allowed to act. And some remain human-primary, with AI only assisting. All three coexist, and the mix shifts as trust and capability grow.

A surface that has to hold all three modes at once is not a dashboard, and it is certainly not a chatbot. It is a workspace: a place that supports a human operator, a human reviewer of an AI operator, and an autonomous AI, and keeps evolving as the product moves down the continuum. The dashboard isn't being replaced by conversation. It is maturing into something that can carry a hybrid way of working.

**The dashboard isn't dying into a chatbot. It's evolving into a workspace, one that holds human and AI operators on the same surface.**

## WHAT THE FRONTIER GOT WRONG

# Klarna ran the experiment. It is now the cautionary tale.

No company tested the "AI replaces the human operator" thesis more publicly than Klarna. The early numbers were spectacular, which is exactly why the reversal matters.

## ~700+

Agents' worth of work the AI assistant was said to handle (later cited at 853)

## ~40%

Headcount reduction as hiring froze and attrition ran

## U-turn

By mid-2025, rehiring humans after quality fell

For a while, the metrics agreed with the strategy. Two-thirds of customer conversations routed through the bot, resolution times dropped sharply, labour costs fell, and the efficiency story compounded into the IPO narrative. On a dashboard of operational KPIs, it looked like a win.

Then the part the KPIs didn't show arrived. Customers hit robotic responses, inflexible scripts and the loop of repeating themselves to a human after the bot failed. The AI could answer questions. It could not handle nuance, refunds or loyalty. The CEO conceded the company had focused too much on efficiency and cost, that the result was lower quality, and that this was not sustainable. Klarna began rehiring people.

The instructive nuance: Klarna did not abandon AI. It rejected full replacement. The correction was to a hybrid model, AI for high-volume routine queries, humans for escalation, empathy and judgment. The industry pattern now reads the same way across case after case. Hybrid human-AI consistently beats full automation on both cost and satisfaction.

This is the cost of mistaking the continuum for a switch. Klarna leapt to the autonomous end for a class of work, emotionally complex service, that had not earned that level of trust. The error wasn't using AI. It was collapsing a continuum into a single jump, and designing the experience as if the human operator could simply be removed.

**The mistake wasn't AI. It was treating a trust continuum as a switch, and removing the human before the trust was there.**


## THE FUTURRIS READ

# Design for the continuum, not the destination

The surface that survives this transition is one built to move along the trust continuum. It supports human-primary work today and absorbs more autonomy as trust is earned, all without a redesign at every step.

## THE TRUST CONTINUUM

Trust in AI deepens over time →



**One workspace supports all three modes and shifts right as trust grows**

The dashboard-to-workspace evolution along the AI-to-trust continuum.

Here is the read only a behavior-engineering firm will give you. The design problem is no longer "how do we visualize status for a human operator." It is "how does the role of the human operator evolve as AI takes on more of the work, and what must the surface support at each stage of that evolution." A product stuck at one point on the continuum, or one that pretends the whole journey collapses into a chatbot, will be rebuilt within a year. The workspace is the metaphor because it is the only one that holds a human operator, a reviewer of an AI operator, and an autonomous agent on the same surface, and keeps evolving as more tasks move rightward.

## WHERE WE LAND

**"The dashboard isn't being replaced by a chatbot. It's growing up into a workspace for humans and AI to operate together."**

Don't ask "where do we add the AI?" Ask "where are we on the trust continuum, and how does the human operator's role evolve from here?" Klarna answered the first question and skipped the second. The surface you build is the answer to the second, and it is a product-design decision, not a model one.

## ABOUT THIS BRIEF

# Futurris Frontier is a monthly read on where financial product experience is heading.

One structural shift, the evidence behind it, and the read only a behavior-engineering firm would give you. Built for product leaders who can't afford to be six months behind. No vendor hype, no feature lists. Just the signal and what to do about it.

Futurris is a product design and behavior engineering firm working with payments, credit, banking, broking and wealth companies. We design products around how people actually behave, and increasingly, around how the human operator's role evolves as work is shared with AI.

---

## CONTINUE THE CONVERSATION

### Speak to us about transforming your digital experience →

We work with product and business teams on activation, retention, feature adoption and journey redesign, from diagnosis through to shipped experience.



*"The dashboard isn't dying into a chatbot. It's growing up into a workspace where humans and AI operate together, and that boundary is a design decision, not a model one."*

**Amber Krishan**

FOUNDER, FUTURRIS