


AI in Banking

From Deployment to **Measurable Value**



TITLE	AI in Banking
SUBTITLE	From Deployment to Measurable Value
AUTHOR	Polina Fomenkova
FOREWORD	Tytus Golas
EDITION	First edition, 2026
PUBLISHER	Tidio for Tidio Agents
RIGHTS	© Tidio, 2026. All rights reserved.

ISBN 978-83-68606-32-4

CONTENTS

AI IN BANKING: FROM DEPLOYMENT TO MEASURABLE VALUE


Foreword	6
1. Executive Summary	7
2. How Banks Are Putting AI to Work	9
3. What Banks Are Doing With Chatbots	12
4. Inside the Bank: Where the Evidence Is Strongest	15
5. From Copilots to Agents	16
6. The United States	18
7. Europe and the United Kingdom	20
8. Poland	23
9. Asia-Pacific	26
10. Who Controls the AI Supply Chain	28
11. An Operating Checklist for the Next Twelve Months	30
Sources	34
About the author	42

MORE SALES, FEWER TICKETS.

The leading AI that resolves customer issues on its own

Tidio Lyro is the AI agent for customer service. It recommends products like an experienced shopping assistant, captures leads, closes sales, and resolves customer issues end to end. No army of engineers required.

Tidio Lyro's average resolution rate vs. the industry average



Market data: Contentsquare; Tidio Lyro's resolution rate represents the average across all customers.

Resolve 72% of customer queries with an AI agent that acts, not just chats. Tidio Lyro gets to work immediately, using your product catalog, knowledge base, and conversation history. Issues are handed off to people only when necessary.

AI in Banking: From Deployment to Measurable Value

Foreword

Banks have used chatbots for years, but the pre-ChatGPT generation rarely felt like a conversation. Most systems followed decision trees, matched keywords, and returned scripted answers. They could deflect simple questions, but they often failed as soon as a customer used unexpected language, combined several issues, or needed the system to understand context. Too often, the customer's real task became finding a way past the bot.

Generative conversational AI changes what is technically possible. In only the past few years, systems have become much better at understanding natural language, working across a bank's knowledge, and maintaining context through a conversation. That makes a more useful service model possible: routine help can be immediate, while human expertise becomes easier to reach when a situation is complex, sensitive, or consequential.

The technology is ready to do substantially more, but banking is a difficult proving ground. A poor answer from a retail chatbot is frustrating; a poor answer about a payment, disputed transaction, credit product, or financial hardship can cause real harm. Banks therefore face a tighter set of obligations than most companies around privacy, explainability, conduct, security, and accountability. Those obligations should shape implementation, not become an excuse to delay it indefinitely.

Customer expectations have changed faster than most banking service models. People are becoming accustomed to asking a question in natural language and receiving a useful answer immediately. They will not distinguish between a bank's chatbot, search function, help center, and mobile application. They will judge the whole experience by whether it understands the request, uses the right information, completes the next step safely, and knows when to involve a person.

This is where newer systems differ from traditional frequently asked question bots. A useful banking assistant should work from bank-controlled knowledge, recognize the relevant customer context, guide onboarding, explain verification requirements, provide

payment updates, and answer questions about products or card limits. The evidence in this report shows that these capabilities are no longer hypothetical. The goal is not unlimited freedom for the model but a bounded role, verified information, clear permissions, and a reliable route to a human advisor.

My strongest view is that containment alone is the wrong measure of success. A conversation that never reaches an employee can still leave the customer confused, generate a repeat contact, or become a complaint. Banks should measure resolution, satisfaction, repeat contact, escalation quality, and cost together. If automation improves one while damaging the others, it has not improved the service.

I also see competitive advantage moving away from access to any particular model. Models will continue to improve and become more interchangeable. The advantage that lasts will come from the knowledge connected to them, the workflows built around them, the quality of evaluation, and the discipline with which a bank limits and monitors their actions.

Banks do not need to automate everything to move decisively. They should choose a narrow, high-volume service problem, establish a measurable quality baseline, and expand only when the evidence supports it. In banking, trust is not a constraint on innovation; it is the condition that allows innovation to scale.

Tytus Golas Founder and Chief Executive Officer, Tidio

1. Executive Summary

Deployment is widespread, comprehension is not. The latest published United Kingdom (UK) supervisory survey found 75% of financial firms using artificial intelligence (AI) and another 10% planning to; 46% reported only partial understanding of their systems, while just 2% of use cases ran fully autonomously.¹

Published in November 2024 and republished by the Financial Conduct Authority (FCA) in December 2025, it remains the last verified reading while a new edition is run in 2026.^{1,2}

Evidence is strongest inside the bank and thinnest at the customer boundary. Banco Bilbao Vizcaya Argentaria (BBVA) has taken ChatGPT Enterprise past 100,000 employees, JPMorgan Chase reaches roughly 250,000, and Citi trained 175,000 employees across 80 markets.^{3,4,5,6} Customer assistants also operate at scale—Bank of America’s Erica reached 24.6 million active users in the second quarter of 2026—but banks disclose little about containment, cost, or service quality.⁷

Disclosure, not capability, now separates the leaders. Only 38% of AI use cases announced by the world’s 50 largest banks in the first quarter of 2026 reported any outcome, and 15 of those 50 publish an aggregate return figure.⁸ DBS Bank—originally established as the Development Bank of Singapore—has committed to a single annual economic-value number since 2022. It reported roughly S\$1 billion for 2025 from data analytics and AI/ML in its annual report, benchmarked against control groups rather than simply asserted.^{9,10} Most of its peers publish nothing comparable.

Regulatory obligations are becoming more specific—but not uniform. European Union (EU) customer-transparency duties took effect on 2 August 2026 while high-risk obligations were deferred; Germany’s Federal Financial Supervisory Authority (BaFin, from its German name *Bundesanstalt für Finanzdienstleistungsaufsicht*) now enforces the chatbot duty.^{11,12} United States (US) agencies excluded generative and agentic AI from the scope of revised model-risk guidance, leaving them subject to banks’ broader governance frameworks.¹³ The European Central Bank (ECB) requires significant institutions to submit AI-enabled cyber-threat action plans by 31 October 2026.¹⁴ The difference is one of regulatory architecture and timing, not a simple split between stricter and looser markets.

Chatbots are becoming service orchestrators. Bank of America’s Erica supports everyday financial management and proactive guidance; PKO Bank Polski’s IKO voice assistant can execute transfers, change limits, and repay card debt; and Bank Millennium’s Milla was designed to answer account questions and execute defined transactions.^{7,15,16,17} The opportunity is broader than deflection, but the control requirement rises sharply once a bot can move money or influence a regulated decision.

The conclusion is implementation. Banks have credible uses for all four layers of AI: predictive models, customer-facing conversational systems, employee copilots, and agentic workflow execution. Their different risks call for controls proportionate to what each system can see, decide, communicate, and execute, not a blanket limit on adoption.

The report first distinguishes the four technology layers, then examines customer service, internal deployment, and agents. Regional chapters explain the applicable rules before the final sections bring the evidence together into management lessons and an operating checklist.

2. How Banks Are Putting AI to Work

“AI in banking” has become a phrase that obscures more than it describes. Four distinct technologies sit behind it, with materially different uses and controls, and most large banks run all four simultaneously—sometimes inside a single customer journey. Implemented well, they form a practical progression: predictive models improve decisions, conversational systems improve access to service, copilots improve employee work, and agents execute bounded workflows.

2.1 Predictive Models

Fraud scoring, transaction monitoring, credit decisioning, and market-risk models have operated for years. Many are probabilistic rather than deterministic, so the distinction that holds is not that they are inherently stable or explainable but that banks have longer-established validation, monitoring, documentation, and model-risk processes for them. Their principal risks are bias, drift, and unexplained decisions; formal validation and outcome monitoring make deployment possible. They carry much of the sector’s established AI value and sit mostly outside this report’s focus.

Example. OCBC Bank reports that AI supports about six million decisions each day, ranging from personalized credit-card offers to

fraud detection.¹⁸ This illustrates the scale of predictive AI already embedded in routine banking rather than introduced as a separate customer product.

2.2 Customer-Facing Conversational Systems

Chatbots and voicebots answer questions, retrieve information, guide onboarding, route service cases, and—when connected to authenticated banking rails—execute bounded actions. They are the most visible layer and carry immediate conduct, privacy, fraud, and reputational exposure because model output reaches the customer directly. Clear escalation rules and transaction confirmation, backed by continuous quality review, allow banks to use this layer safely rather than avoid it.

Examples. Bank of America’s Erica supports account servicing, money movement, bill payment, financial insights, and connection to specialists.^{7,19} PKO Bank Polski’s IKO voice assistant goes further into bounded transactions: customers can use it to make transfers, change limits, and repay card debt.¹⁵

2.3 Employee Copilots

Retrieval, drafting, summarization, meeting support, and coding tools serve relationship managers, analysts, developers, and service staff. Human review and access controls reduce external harm, but they do not make the layer automatically safe: confidential-data leakage, hallucinated policy, automation bias, and inadequate verification remain material risks. Because employees can verify outputs before acting, this is often the most practical entry point for generative AI.

Example. BBVA expanded ChatGPT Enterprise to roughly 120,000 employees across 25 countries. In one documented workflow in Peru, more than 3,000 employees use the system, and the bank and its model provider report that average query-handling time fell from about 7.5 minutes to about one minute.^{3,6}

2.4 Agentic Workflow Execution

Agentic systems plan and complete multi-step work across enterprise tools within a defined permission set. Excess authority is the principal risk, alongside cascading errors and weak traceability. Permission limits, continuous monitoring, complete logs, and a tested stop mechanism are prerequisites for deployment. Publicly announced agentic applications represented 31% of new banking AI use cases tracked by Evident in the first quarter of 2026, up from 15% in the previous quarter; the dataset measures announced use cases, not the full population of deployments.⁸

Example. PKO Bank Polski has deployed an internal agent that classifies and supports responses to incoming customer email. The agent has its own digital identity inside the bank, while employees retain decision authority—a concrete example of bounded agency rather than uncontrolled automation.²⁰

The supervisory problem is that generative and agentic systems do not fit neatly into controls designed for traditional quantitative models. Outputs are probabilistic and can be fabricated, and they may change after a model or orchestration update. In April 2026, US banking agencies expressly excluded generative and agentic AI from revised model-risk guidance while stating that banks' broader governance should determine appropriate controls.¹³ The exclusion defines the guidance's scope; it says nothing about whether these systems are governable or lower risk.

Across acquisition and service, credit, payments and operations, financial crime, and internal support, no current independent estimate of realized sector-wide value was identified for this report. The stronger evidence is narrower: banks tracked by Evident announced nearly 50% more AI use cases during 2025, yet only 38% of use cases announced in the first quarter of 2026 reported any outcome, and only 15 of the 50 banks in the dataset disclosed an aggregate return figure.⁸ That evidence supports workflow-level business cases with explicit baselines; it does not support a precise claim for the value of AI across global banking.

OUTCOME DISCLOSURE

38%

of use cases announced in the first quarter of 2026 reported any outcome, and 15 of the 50 banks disclosed an aggregate return

3. What Banks Are Doing With Chatbots

The customer-facing deployments examined for this report are a small and self-selected sample because banks disclose successful programs more readily than unsuccessful ones. Even within that limitation, effective designs share a recognizable structure. Informational and self-service traffic can be fully automated. Transactional requests can be automated with confirmation steps and constrained action scopes. Complaints, hardship, disputes, fraud, and bereavement should remain human-led, with AI supporting the advisor. This division allows banks to automate substantially without treating every customer need as equivalent.

3.1 Current Deployments

Bank of America's Erica supports account servicing, money movement, bill payment, spending and balance insights, proactive financial guidance, and connection to specialists. The bank discloses adoption and interaction volume, but not an Erica-specific series for containment or cost per contact.^{7,19}

PKO Bank Polski's IKO voice assistant supports transfers, balances and transaction history, phone top-ups, payments, limit changes, and card-debt repayment. The cited public account reports more than one million conversations but does not provide a current containment or satisfaction series.¹⁵

Bank Millennium's Milla answers defined account questions and executes defined transaction commands. Launch materials from December 2018 describe roughly 800 question types and 400 transaction commands; those figures have not been publicly refreshed since, and this report treats the "first in Poland" description as the bank's launch claim rather than an independently established market fact.^{16,17}

Erica reached 20.6 million users and nearly 700 million interactions in 2025; in the second quarter of 2026, active users rose 23% year over year to 24.6 million and interactions reached roughly 200 million.^{7,21} Nubank reports resolving about 55% of first-tier contacts without

ASSISTANT REACH

24.6 million

active users in the second quarter of 2026, up 23% year over year

escalation across more than two million monthly chats; the figures come from its model provider.²²

Klarna is often described as having replaced 700 customer-service employees with AI and then hired them back. The documented sequence is different. In February 2024, Klarna said its new assistant handled two-thirds of customer-service chats and performed work **equivalent** to 700 full-time agents; it did not say that 700 employees had been dismissed when the assistant launched.²³ The company had broadly frozen hiring in late 2023. By August 2024, its chief executive described headcount falling mainly through natural attrition, with recruitment continuing only for selected roles.²⁴ Klarna's filing with the US Securities and Exchange Commission records full-time employees declining from 5,527 at the end of 2022 to 4,352 in 2023 and 3,422 in 2024, and links the reduction to a strategy of lowering headcount and improving efficiency with AI.²⁵

In May 2025, the company began recruiting a limited group of flexible, remote customer-service workers so customers could always reach a person. Chief Executive Officer Sebastian Siemiatkowski said Klarna had focused too heavily on cost and that the resulting service quality was too low, while stressing that the company was giving AI more—not less—customer-service work.²⁶ This was therefore a correction to the service model, not a wholesale reversal of AI adoption or a restoration of the former workforce. Later disclosures said the assistant was doing the work of 853 full-time agents, had a net promoter score of 73, and had generated \$60 million in savings. They also showed third-quarter service and operations costs rising from \$42 million to \$50 million year over year.²⁷ Those figures do not establish that human recruitment cost more than the AI savings: the cost line includes broader service and operations spending, and Klarna has not published the cost of the hiring initiative separately. The useful lesson is narrower. Automation volume and modeled labor equivalence cannot substitute for direct measures of resolution quality, repeat contact, escalation, and access to a human.

Human-augmented service is the more defensible pattern. Bank of America pairs Erica with EricaAssist, deployed to more than 18,000 service representatives and surfacing answers in under three seconds.²⁸ Bank Millennium in Poland describes a deliberate two-

SAVINGS AND COST

\$50 million

in third-quarter service and operations costs, up from \$42 million a year earlier, while the same disclosures credited the assistant with \$60 million in savings

level design in which bots take simple queries with an advisor always one step away, while a second AI layer analyzes data and suggests options to the human; the bank's own conclusion is that a person equipped with these tools is both more helpful and more effective than either alone.²⁹

Graduated autonomy is the cheapest control available. Monzo's Ops Agent works across a support surface of more than 150 distinct customer intents and lifted resolution rates on transaction queries by 10 percentage points; the bank describes a staged rollout, beginning with a phase in which humans reviewed every outbound message before it reached a customer, then permission to speak to customers directly under quality assurance on a sample of conversations, with the sampling rate thinned only as confidence grew.³⁰ Reading that sequence as the transferable lesson, rather than the resolution-rate figure, is this report's inference and not a claim Monzo makes.

Disclosure duties now attach to this layer specifically, and one supervisor has put an enforcement machine behind them. The EU's Article 50 transparency requirements, which oblige firms to tell customers they are dealing with AI, have applied since 2 August 2026, the AI Act's general application date, and were untouched by the deferral of the high-risk regime.^{11,31} Germany shows what that becomes once a national financial supervisor is wired into it. The German implementing act entered into force on 29 July 2026, making BaFin the market surveillance authority for AI systems directly related to regulated financial activity. From 2 August 2026 BaFin enforces the transparency duty itself, so a chatbot in customer communication must identify itself as one, and the authority can impose fines for breaches.¹² High-risk supervision—creditworthiness assessment, insurance risk evaluation—follows on 2 December 2027, tracking the deferred EU timetable.¹² For any bank running a customer assistant in Germany, the compliance question stopped being theoretical a fortnight before the date of this report. Hong Kong's supervisor arrived at the same place by a different route two years earlier, with a dedicated circular on consumer protection in generative AI use in August 2024: dated, but not superseded.³²

The United States offers no equivalent, and that absence is a finding rather than a shortfall in the research. No federal consumer-side AI

disclosure rule and no comparable federal measurement of chatbot exposure was identified for this report. A US bank’s customer-facing assistant is governed through general prohibitions on unfair, deceptive, or abusive practices, and through the adverse-action reasoning duties described in section 6—not through any duty to say it is a machine.

4. Inside the Bank: Where the Evidence Is Strongest

Internal deployment carries a lower risk profile and, not coincidentally, produces better-documented results. The same archetypes recur across every market surveyed.

The enterprise assistant. JPMorgan’s large language model (LLM) Suite reaches roughly 250,000 employees, about half daily, and went from zero to 200,000 onboarded users in eight months—the firm’s own figure.⁴ The denominator excludes branch and call-center staff. On time saved, the bank says “several hours per week”; more specific ranges circulate in secondary coverage but could not be traced to a verifiable source for this report, and are omitted.⁴ Citi had reached roughly 140,000 employees across eight countries by the end of 2024 and expanded to 11 countries by April 2025, extended prompt training to 175,000 employees across 80 locations that autumn, and logged more than 6.5 million prompts across the year.^{5,33} Goldman Sachs went firmwide after a pilot of roughly 10,000.³⁴ Oversea-Chinese Banking Corporation (OCBC Bank) announced on 24 October 2023 that it would extend its assistant to its entire 30,000-strong workforce the following month.³⁵

Internal knowledge retrieval. PKO Bank Polski’s szukAI serves more than 16,000 employees searching the bank’s own documents, up from 11,000 nine months earlier, and is the largest internal retrieval deployment disclosed in the Polish market on the evidence reviewed here.³⁶

Developer assistance. Nearly 1,400 developers at PKO Bank Polski and roughly 40,000 at Citi use coding assistants or agents.^{5,36}

ENTERPRISE ASSISTANT

250,000

employees reached by
JPMorgan’s LLM Suite, about
half of them daily

Commonwealth Bank of Australia, unusually, disclosed that its Copilot rollout produced mixed results in utilization and efficacy.³⁷

Research and meeting summarization. Morgan Stanley’s assistant reached adoption above 98% of advisor teams, a figure carried on its model provider’s customer page, and was later joined by Debrief for client meetings.^{38,39}

Correspondence triage. PKO Bank Polski’s first agent classifies and supports inbound customer email under human decision authority.²⁰

Measurement problems run through all of it. First, hours-saved figures are often employee assessments rather than observed output. Second, most exclude verification time. OCBC reports that staff completed tasks roughly 50% faster even after including time spent checking model output.³⁵ Any vendor presenting a comparable claim should be asked whether the same is true of theirs.

The most precise disclosed productivity metric in this report comes from one BBVA workflow in Peru. More than 3,000 employees use it, and the bank and its model provider report that average query-handling time dropped from about 7.5 minutes to about one minute.³ The metric is specific and reproducible, although it remains bank- and vendor-reported rather than independently audited. A survey result such as “hours saved per week” measures perceived productivity and belongs in a weaker evidence category.

5. From Copilots to Agents

Agentic applications reached 31% of newly announced banking use cases in the first quarter of 2026, concentrated in product and service operations.⁸ Yet public evidence remains thin: at an earlier count, only 9 of the 50 largest banks had documented agents in pilot or production, and only Bank of New York Mellon (BNY), Capital One, and JPMorgan had disclosed supporting architecture.⁴⁰

Current deployments favor bounded autonomy. Wells Fargo has rolled out an agent platform to its workforce and, in August 2026, joined the Agentic AI Foundation as a gold member, arguing that open collaboration and interoperability are needed to make agents

VERIFIED PRODUCTIVITY

~50%

faster task completion at OCBC, measured after the time staff spent checking the model’s output

AGENTIC SHARE

31%

of newly announced banking use cases in the first quarter of 2026, concentrated in product and service operations

useful across platforms; Goldman Sachs has piloted agentic software development; and BNY has provided employees with a multi-agent platform.^{41,42,43} PKO Bank Polski gives its correspondence agent a digital identity inside the bank while retaining human decision-making, and plans 30 more agents by the end of 2026 across complaints, cybersecurity, transaction protection, and post-sale support.^{20,44}

Governance is fragmented. US model-risk guidance excludes generative and agentic AI but still leaves banks exposed to examination under general safety-and-soundness principles.¹³ Singapore's proposed guidelines expressly cover autonomous agents but are not yet final; India's Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI) recommendations assign accountability to the deploying institution and preserve human override, but have not been adopted as Reserve Bank of India (RBI) directions.⁴⁵ Designing a common control baseline is therefore cheaper than maintaining a separate architecture for every jurisdiction, even though each legal regime applies only within its own scope.

Payments expose an unresolved issue. The International Monetary Fund (IMF) argues that adaptive systems acting under delegated authority challenge laws written on the assumption that a human initiated the transaction, and asks whether Know Your Customer needs a "Know Your Agent" counterpart linking a transacting bot to a legal entity.⁴⁶ PKO Bank Polski's digital identity is an early operational answer, while the UK Treasury's adoption plan shows that the wider policy response is still forming.^{20,47} Hong Kong's sandbox has also tested AI-based checks on AI output, but no effectiveness data were identified.⁴⁸ The durable requirement is continuous evaluation: autonomy that cannot be monitored and stopped cannot be scaled safely.

The forms of AI described above can be deployed in every major banking market, but the route to deployment is not the same everywhere. The next chapters apply a common lens—what banks have implemented and what they can demonstrate, set against what supervisors require—to the United States, Europe and the United Kingdom, Poland, and Asia-Pacific. The comparison matters because

a control that is good practice in one market may be a legal duty in another, while elsewhere it is only a supervisory expectation or a proposal.

6. The United States

The United States provides the clearest example of institution-led governance. American banks are deploying all four forms of AI, but the federal framework has just been redrawn in a way that leaves generative and agentic systems formally outside model-risk guidance. Banks have considerable room to design their own controls, alongside continuing safety-and-soundness and consumer-protection obligations.

6.1 What the Largest Institutions Have Published

Beyond the enterprise assistants already described, the deployments below merit specific attention.

Bank of America offers the longest continuous public record among the deployments reviewed here, and its discipline lies in scoping. Erica was not positioned as a substitute for human service; it was expanded incrementally across eight years and paired with an internal tool for staff.²⁸ What its disclosure lacks is economic: volume without cost.

Wells Fargo contributes a useful privacy-preserving architecture. Speech is transcribed locally, scrubbed and tokenized internally, and only then sent through a narrow external-model call for intent and entity extraction.⁴⁹ The design exchanges some conversational richness for lower data exposure. Internally, the rollout of Google Agentspace began with more than 2,000 staff, and the bank announced generative and agentic tools for all 215,000 employees.^{42,50,51}

Morgan Stanley is the methodological outlier. Its most transferable contribution is the structured evaluation process used to govern deployment.³⁸

6.2 Revised Model-Risk Guidance and Its Scope

On 17 April 2026 the Federal Reserve, Office of the Comptroller of the Currency (OCC), and Federal Deposit Insurance Corporation (FDIC) issued Supervision and Regulation Letter 26-2 (SR 26-2), replacing SR 11-7 and superseding the 2021 interagency statement on model risk for Bank Secrecy Act and anti-money-laundering (BSA/AML) systems.¹³ It is expected to be most relevant above \$30 billion in assets, narrows the model definition by excluding simple arithmetic and deterministic rules, and replaces fixed validation cycles with a risk-based schedule tied to purpose, changes, and data limitations.¹³

Generative and agentic AI sit outside the document's scope, with banks directed to apply their own risk practices.¹³ SR 26-2 is guidance rather than a rule, but agencies retain authority over unsafe or unsound practice; the absence of a prescribed method is therefore not an absence of exposure.¹³ A promised request for information on AI had been announced but not issued at the report date.⁵²

Banks should build institution-specific controls now because older model-risk concepts strain when applied to dynamic, probabilistic, and autonomous systems.⁵³

Consumer protection obligations are unchanged whatever the enforcement climate. Creditors using AI must still give specific adverse-action reasons under the Equal Credit Opportunity Act (ECOA) and Regulation B regardless of model complexity.⁵⁴ The Federal Trade Commission (FTC) has pursued deceptive AI claims under its Operation AI Comply sweep, an exposure that sits in marketing copy as much as in models; the cited actions date from September 2024, and no comparable sweep was identified for the current administration.⁵⁵ At state level, fragmentation should be the planning assumption: Colorado's AI Act was repeatedly delayed and then substantially narrowed, stripping out its duty of care and deployer risk-program obligations in favor of narrower disclosure.⁵⁶

The US model is permissive in form but not consequence-free: banks can choose the architecture, yet they remain accountable for outcomes. Europe begins from the opposite direction, defining

more of the obligations in legislation and allocating enforcement to specific authorities.

7. Europe and the United Kingdom

Europe runs two regulatory philosophies simultaneously. The European Union is building an AI-specific legal framework, while the United Kingdom continues to apply technology-neutral financial rules. The practical difference for a bank operating in both is larger than the shared geography suggests, even when the underlying chatbot, copilot, or agent is the same.

7.1 Brussels: Prescriptive and Deferred, but Already Binding in Part

The AI Act took force on 1 August 2024, with prohibited practices applying from February 2025 and general-purpose model obligations from August 2025.⁵⁷ Delays to standards, guidance, and national supervision led to the Digital Omnibus.³¹

Regulation (EU) 2026/1744 took force on 27 July 2026.^{11,31} High-risk obligations for standalone Annex III systems, including natural-person creditworthiness and certain employment uses, moved to 2 December 2027; those embedded in regulated products moved to 2 August 2028.¹¹ Earlier generative systems received until 2 December 2026 to comply with content marking, and member states have until 2 August 2027 to operate sandboxes.¹¹

What did not move is precisely what bears on customer service. Article 50 transparency kept its original schedule and has applied since 2 August 2026, the AI Act's general application date.^{11,31}

Germany's implementing act took force on 29 July 2026 and made BaFin the market-surveillance authority for AI tied directly to regulated financial activity. From 2 August it can enforce transparency duties for systems interacting with people, including customer chatbots; high-risk creditworthiness and insurance requirements follow on 2 December 2027.¹² The same authority now

HIGH-RISK DEFERRAL

2 December 2027

when standalone Annex III obligations bite for creditworthiness and employment uses, with those embedded in regulated products following on 2 August 2028

examines both a bank's prudential position and whether its assistant identifies itself as AI.

Article 4 is separate: its AI-literacy duty has applied since 2 February 2025 and was changed from ensuring sufficient literacy to supporting its development—an obligation of effort rather than result.^{11,58} The Omnibus also added Annex XIV code AIH 0401 for agentic AI, but this is administrative nomenclature, not a legal definition or autonomy-based risk tier.^{11,59} Customer-facing disclosure was not relaxed; the high-risk rules were deferred, with transitional treatment for systems placed on the market before the new dates until they are substantially modified.¹¹

7.2 Frankfurt: Operational-Resilience Requirements

For most euro-area banks, immediate supervisory pressure comes through operational-resilience and information and communications technology (ICT) risk frameworks as well as the AI Act. The Digital Operational Resilience Act (DORA) is the EU framework governing financial-sector ICT resilience. A European analysis that stops at Brussels misses the ECB's role.

The European Banking Authority (EBA) found no significant contradictions between the AI Act and existing EU banking and payments law, describing the Act as complementary to a sectoral framework that already provides comprehensive risk management, and identifying no immediate need to introduce new EBA guidelines or revise existing ones. It did note that integrating the two frameworks will take some effort, and that the co-existence of prudential, conduct, and market-surveillance authorities makes supervisory cooperation important.⁶⁰

The ECB monitors AI through operational resilience and ICT risk. Across roughly 110 significant institutions, use cases rose sharply; in a 13-bank workshop, about half had dedicated AI policies or committees and none allowed post-deployment self-learning.⁶¹ In June 2026 the European Systemic Risk Board (ESRB) warned that frontier models can discover vulnerabilities and generate exploits

faster than traditional remediation cycles.^{62,63} The ECB consequently called on significant institutions to submit AI-enabled cyber-threat action plans to their Joint Supervisory Teams by 31 October 2026, with a horizontal analysis of the submissions to follow.^{14,64} Brussels deferred some build obligations while Frankfurt accelerated defense obligations.

7.3 London: The Technology-Neutral Approach Under Scrutiny

The UK declined AI-specific rules in favor of existing regimes, with Consumer Duty carrying much of the weight for customer-facing systems.⁶⁵ Its distinctive instrument is AI Live Testing, whose second cohort from April 2026 includes Barclays, Lloyds Banking Group through Scottish Widows, and UBS.⁶⁶ Testing inside the regulator's own perimeter converts supervisory uncertainty into documented dialogue before launch rather than after, a substantially different risk posture from the EU's conformity-assessment route.

In January 2026 the Treasury Committee found gaps in AI-specific stress testing, Critical Third Party designation, and published senior-manager accountability guidance, setting end-2026 deadlines for action.⁶⁷ His Majesty's Treasury (HM Treasury) responded in its July adoption plan with 10 recommendations, including a cross-regulator guidance hub, faster provider designation, and agentic-payments readiness.⁴⁷ It also reported that 26% of UK adults trust general-purpose AI for financial guidance, creating a strategic risk that banks lose the advice-discovery channel.⁴⁷

The latest published Bank of England/FCA survey found that one-third of use cases were third-party implementations and 55% involved some automated decision-making. Firms rated 16% as high materiality—meaning potentially significant business or customer impact—most commonly in general insurance, risk and compliance, and retail banking.^{1,68} Although 84% named an accountable person, most named at least three people or bodies, diluting ownership.¹ A fourth survey was underway in 2026; the 2024 figures should be revisited when it appears.²

NAMED OWNERSHIP

84%

of firms named an accountable person for AI, but most named at least three people or bodies

7.4 BBVA: The Anatomy of a Rollout That Scaled

BBVA expanded ChatGPT Enterprise from 3,300 licenses in 2024 to roughly 120,000 employees across 25 countries by December 2025.^{6,69} Users of the 11,000-license phase reported saving 2.8 hours weekly, and employees across the organization have since created more than 20,000 assistants, although the evidence is bank- and vendor-reported.^{3,69} The transferable decisions were secure access to displace shadow AI, frontline distribution, leader training, and a multi-model estate: almost the whole workforce holds licenses under global agreements with both OpenAI and Google, and by April 2026 the average employee was using ChatGPT on 12 days a month and Gemini on nearly nine.^{3,70}

NatWest put about 25 of more than 275 explored use cases into production, a useful measure of pipeline attrition.⁷¹ Commerzbank launched Ava, an app-based avatar assistant built on Microsoft's Azure AI toolkit.⁴⁰

These European rules define the regional baseline, but implementation still depends on national institutions, local banking structures, and language. Poland makes those differences visible. Its banks are technically ambitious and its customer interfaces must work in Polish, while its domestic supervisory architecture is still taking shape.

8. Poland

Poland is a practical test of how the European story works at national level. It warrants focused treatment for a concentrated banking sector with genuine technical ambition, a language constraint that larger European markets face less acutely, and a newly established national supervisory architecture.

8.1 What the Banks Have Built

PKO Bank Polski, Poland's largest bank, reports nearly 1,400 developers using a coding assistant, more than 16,000 employees

LICENSE SCALING

3,300

ChatGPT Enterprise licenses in 2024, reaching roughly 120,000 employees across 25 countries by December 2025

INTERNAL TOOLS

16,000

employees on PKO Bank Polski's szukAI, beside nearly 1,400 developers on a coding assistant

using its internal-search tool szukAI, 1,500 on Microsoft Copilot with a full rollout planned, and two in-house generative models.^{20,36} It also operates platforms for managing conventional and large-language models across their lifecycle—commonly called MLOps and LLMOps—and co-developed the Polish RoBERTa 8k model and a related safety model.^{20,72}

In August 2026 PKO deployed an agent that classifies and supports inbound customer email. It runs inside the bank, has its own digital identity, and leaves decisions to employees.²⁰ The projected 70% improvement in response time had not yet been measured at publication.²⁰

On the customer side, the bank's IKO voice assistant has run since May 2020, executing transfers, checking balances and history, topping up phones, making payments, changing limits, and repaying card debt, and has passed a million customer conversations. It was built deliberately to parse colloquial Polish, including instructions phrased in slang.¹⁵

Bank Millennium launched Milla in December 2018, presented at launch as the first Polish bank chatbot able both to answer questions and execute transactions, built around roughly 800 cataloged account questions and 400 transaction commands, and designed for Polish dialects and slang.^{16,17} Its 2025–2028 strategy adds a call-center assistant and round-the-clock automated chat while explicitly preserving direct human contact.⁷³ **mBank** and **Alior Bank** were among the earlier movers in contact-center voicebots and speech analytics; the cited account dates from January 2020, when Alior's voice assistant was live and mBank's voicebot and speech analytics were still planned.⁷⁴

8.2 The Language Question

Poland has two significant domestic open models. Bielik is developed by the SpeakLeash community with AGH University of Krakow; PLLuM comes from a consortium led by Wrocław University of Science and Technology and now numbering eight institutions, among them the NASK National Research Institute, the National

Information Processing Institute (OPI-PIB), and two Polish Academy of Sciences institutes.^{75,76} The working division is that PLLuM targets public administration and citizen communication—it includes an intelligent citizen assistant for public administration—while Bielik addresses commercial use.^{76,77}

Bielik and PLLuM address the Polish-language constraint with local deployment and stronger handling of native register, typos, and document types.⁷⁷ Benchmarks are task-dependent and not always comparable, and the fullest published comparison is Bielik’s own technical report, which is not a neutral referee: on Polish EQ-Bench it places Bielik-11B ahead of much larger open models, including the best PLLuM variant, while GPT-4o and GPT-4-turbo still rank above it.^{75,78} Banks should therefore use frontier models for broad reasoning under contractual protection and local models for narrow, high-volume work involving sensitive Polish data.

8.3 A Newly Established AI Supervisor

Poland’s Act on Artificial Intelligence Systems, Dz.U. 2026 poz. 1003, entered into force on 11 August 2026.^{79,80}

It creates the Commission for the Development and Security of Artificial Intelligence (KRiBSI, from the Polish **Komisja Rozwoju i Bezpieczeństwa Sztucznej Inteligencji**) as the national market-surveillance authority and EU contact point. The chair of the Polish Financial Supervision Authority (KNF, **Komisja Nadzoru Finansowego**) is represented in its collegial structure.^{80,81} The government timetable envisaged a functioning commission by November 2026 and a regulatory sandbox within four months.^{79,82}

The commencement is staged, and the staging is the part that matters operationally. The provisions on individual opinions and explanations, and the chapters on control and penalties, take effect three months after announcement, on 28 October 2026. Until that date KRiBSI conducts no inspections and imposes no penalties.^{80,83} Institutions have a defined and unusually legible window: the supervisor exists, its powers do not yet bite, and the questions it will ask are already published.

ENFORCEMENT WINDOW

28 October
2026

when Poland’s control and penalty chapters take effect, and until then KRiBSI runs no inspections and imposes no penalties

The Act offers pre-deployment opinions on classification and obligations, protects business secrets, and provides judicial review and an affected-person complaint channel.^{80,81}

KNF separately treats frontier AI as an operational-resilience issue under DORA. Its July 2026 recommendations call for stronger ICT controls, faster vulnerability handling, AI-enabled attack scenarios, supplier-concentration analysis, and incident-response cooperation.^{62,84,85} They are not binding rules, but KNF says implementation will be tested through supervision and inspection.⁸⁵

Poland shows how national supervision modifies a wider regional framework, and so do language and local models. Asia-Pacific offers a different comparison, where the most informative cases are distinguished less by a shared legal regime than by stronger measurement at leading banks and supervised experimentation across several markets.

9. Asia-Pacific

Two banks publish figures absent from most Western disclosures, while regulators in Singapore, India, Hong Kong, and Australia are using guidelines, recommendations, sandboxes, and governance reviews to shape deployment. The instruments differ, but the direction is consistent: enable adoption while making accountability and human control explicit.

9.1 Measurement as Competitive Practice

DBS Bank reports about S\$1 billion of 2025 economic value from data analytics and AI/ML across more than 430 use cases and 2,000 models, up from S\$180 million in 2022.^{10,86} It measures against control groups and separates revenue uplift, savings, and risk avoidance; that makes its bank-reported figure more credible than unsupported aggregate claims.⁹ Two-thirds of employees use DBS-GPT; CodeBuddy has delivered time savings of up to 20% on certain coding tasks for the bank's data professionals; and DBS Joy passed 20,000 corporate customers.^{10,87}

MEASURED VALUE

S\$1 billion

of 2025 economic value across more than 430 use cases and 2,000 models, up from S\$180 million in 2022

OCBC Bank reports up to 20% productivity uplift for relationship-manager support and about six million AI-assisted decisions daily, while simultaneously hiring 600 relationship managers.^{18,88} This is evidence for augmentation, not automatic substitution.

The two banks do not point the same way on employment, and the contrast belongs alongside the value figures. In February 2025 DBS said it would not renew about 4,000 temporary and contract roles over three years as AI took on more work, against roughly 1,000 new AI-related positions; its then chief executive put it bluntly: “For the first time, I’m struggling to create jobs. I’m struggling to say how I will repurpose people.”⁸⁹ Permanent staff were not affected, and the bank has committed to reskilling 13,000 employees. A bank can lead the sector on measuring AI value and still reduce headcount at the contingent edge of its workforce. Measurement discipline and workforce expansion are separate decisions, and DBS and OCBC have made them differently.

WORKFORCE MIX

4,000

temporary and contract roles not renewed over three years, against roughly 1,000 new AI-related positions

9.2 Singapore’s Supervisory Ladder

The Monetary Authority of Singapore (MAS) progressed from its Fairness, Ethics, Accountability and Transparency (FEAT) principles and the Veritas initiative to Project MindForge and proposed Guidelines for AI Risk Management covering traditional, generative, and agentic AI.^{45,90} The guidelines were not yet in force at publication; consultation closed in January 2026 and a twelve-month transition was proposed.^{45,91}

9.3 India, Hong Kong, and Australia

India. The Reserve Bank of India’s FREE-AI Committee proposed 26 recommendations, including accountability regardless of autonomy, human override, board-approved policies, and vendor audit rights.⁹² These remained recommendations rather than adopted RBI directions at publication.

Hong Kong. The Hong Kong Monetary Authority’s (HKMA) second generative-AI sandbox cohort selected 27 use cases from 20 banks;

Sandbox++ later expanded supervised testing across banking, securities, wealth, insurance, and stored-value facilities.^{48,93}

Australia. The Australian Securities and Investments Commission (ASIC) reviewed 23 licensees and 624 use cases. Nearly half had no policy addressing consumer fairness or bias, and only 43% had policies referencing disclosure of AI use to affected consumers; vendor oversight and output monitoring were also weak.^{94,95} The data describe systems in place by December 2023, so they establish a governance baseline rather than a current adoption rate.

Across the regional chapters, the regulatory approaches differ more than the technology stacks do. Banks repeatedly depend on the same cloud platforms, model providers, data suppliers, and specialist applications. That common dependency is the next management question: not only what AI a bank uses, but who ultimately controls the infrastructure on which it runs.

10. Who Controls the AI Supply Chain

The supplier picture has changed in a way that carries direct governance consequences, and it has changed in two directions at once.

The UK supervisory survey quantifies concentration across cloud, model, and data suppliers. The top three providers account for 73% of reported cloud provision, 44% of model provision, and 33% of data provision.¹ These are combined top-three shares within each category, not the share held by any single provider. They show a narrow supplier bench; they do not measure any individual bank's exposure to one company.

Yet the application layer is fragmenting. Anthropic was the most referenced vendor in banking AI use cases announced in the first quarter of 2026, and specialists beyond the hyperscalers accounted for 68% of announced use cases.⁸ The tracker counts publicly announced use cases and pairs them with a private outcomes benchmark; reading the announcement count as a measure of disclosure as much as of deployment is this report's inference rather

DISCLOSURE POLICY

43%

of Australian licensees had policies referencing disclosure of AI use to consumers, in systems reviewed as of December 2023

SUPPLIER CONCENTRATION

73%

of reported cloud provision sits with the top three providers, against 44% of model provision and 33% of data

than a caveat Evident states.^{8,96} The single-platform procurement assumption that shaped budgets two years ago no longer describes the market.

Supervisors have named the risk, and in the UK the Treasury Committee is pushing them to act on it. The Financial Stability Board warned in October 2025 that reliance on a small number of specialized hardware, cloud, and pre-trained model suppliers creates exposure where alternatives are scarce.⁹⁷ The Bank of England has flagged that data and model providers may themselves become critical third parties.⁹⁸ In January 2026 the Treasury Committee called for major providers to be designated under the Critical Third Parties regime before the end of the year, and the Treasury's July adoption plan accepted that direction.^{47,67} MAS's proposed guidelines would require firms to manage third-party AI risk across the model lifecycle, though they were not yet in force at publication.⁴⁵ Poland's KNF explicitly requires analysis of over-dependence on any single ICT or AI supplier.⁸⁴

The available responses are a genuinely multi-model architecture, as BBVA operates by running OpenAI and Google stacks in parallel;^{3,70} documented exit plans tested rather than filed; and a board-level threshold above which single-provider dependence in customer-facing or decisioning workloads triggers escalation.

Technology ownership shapes the operating model inside the bank as well as its resilience and bargaining power. The technology and regional evidence, together with the supplier picture, supports implementation with discipline rather than delay. The final chapter converts those findings into a twelve-month sequence of actions that a bank can assign and document, then test.

11. An Operating Checklist for the Next Twelve Months

- 1 **Start internally, then expand to customers.** The business owner and AI risk lead should move customer-facing systems from review of every output to sampled review, and only then to direct interaction. Approval records, evaluation results, and the threshold for greater autonomy should be documented.³⁰
- 2 **Measure productivity after verification time.** Finance and transformation leaders should establish a baseline or control group and include the time employees spend checking outputs, then separate annual value into revenue uplift and cost avoidance.^{10,35}
- 3 **Measure service quality by complexity.** The customer-service owner should track resolution, satisfaction, complaints, repeat contacts, and escalation for each complexity tier every month. Automation should be narrowed when a quality floor is breached.²⁶
- 4 **Maintain a common control baseline.** The chief risk officer should keep a current AI inventory, assign a materiality classification and control owner to every system, document human oversight and explainability requirements, and map the additional rules in each jurisdiction.^{11,45}
- 5 **Treat AI output as bank communication.** Compliance and conduct-risk leaders should approve disclosure language, then test every customer channel and monitor complaints wherever an AI system communicates with customers.^{11,12,31}
- 6 **Track regulatory deadlines.** Regulatory affairs should maintain a dated register with a named owner and proof of submission. For significant euro-area institutions, the ECB action plan is due to the Joint Supervisory Team—the institution’s ECB supervisory team—by 31 October 2026; Poland’s inspection and penalty provisions start on 28 October 2026.^{14,80}

- 7 **Limit dependence on individual suppliers.** Technology-risk and procurement leaders should set a board-approved threshold for reliance on one cloud, model, or data provider, then maintain a concentration dashboard and complete practical exit tests.^{1,97}
- 8 **Test in the language customers actually use.** In markets where English is not the primary language, model owners and local-market leaders should benchmark frontier and local models against real native-language traffic, and should document both the privacy implications and the deployment rationale.^{75,78}
- 9 **Separate workforce planning from unverified automation claims.** Finance, human resources, and business owners should require seasonally adjusted, multi-quarter evidence on volume, quality, and workload before using AI to justify a headcount decision. The two published cases point in opposite directions and neither is a template: Commonwealth Bank reversed 45 redundancies after call volumes rose rather than fell, while Klarna disclosed a three-year reduction from 5,527 to 3,422 full-time employees and saw its service and operations costs rise anyway.^{25,27,99}
- 10 **Test customer assistants adversarially.** Security and AI-assurance teams should deliberately probe systems with hostile and misleading inputs, remediate failures, retest them, and retain logs sufficient to reconstruct an incident.
- 11 **Give security ownership of shadow AI.** The chief information security officer should monitor uploads to generative-AI services, maintain an exception register, and test the incident-response process.¹⁰⁰
- 12 **Map exposure to potentially critical providers.** In the UK, regulatory affairs and technology-risk teams should identify providers that may be designated as critical third parties, assess the services that depend on them, and maintain tested contingency plans.^{47,67}

Every source below was opened and checked against the claim it supports during the August 2026 verification pass. Sources whose

content could not be independently confirmed were removed and their claims re-sourced or rewritten; the list is numbered consecutively and contains no gaps. The list contains only sources the report relies on.

LEAVE SHADOW AI BEHIND

Hubi — a secure AI agent

Invite Hubi to join your team and watch it help your employees. It operates securely, respects privacy, and stays within the responsibilities you assign it.



Hubi is more than ChatGPT — it doesn't just reply; it takes action. It works within the platforms you already use, including Microsoft Teams and Slack, and is ready in minutes. All on your terms.

Invite Hubi to join your team

Bibliography

- 1 Bank of England / Financial Conduct Authority, «Artificial intelligence in UK financial services – 2024» (November 2024). <https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024>
- 2 Financial Conduct Authority, «Research Note: AI in UK financial services» (research-note page carrying the 2024 survey findings, updated 3 December 2025). <https://www.fca.org.uk/publications/research-notes/ai-uk-financial-services>
- 3 OpenAI, «BBVA puts AI at the core of banking with OpenAI». <https://openai.com/index/bbva/>
- 4 JPMorganChase, «LLM Suite named 2025 'Innovation of the Year' by American Banker». <https://www.jpmorganchase.com/about/technology/blog/llmsuite-ab-award>
- 5 American Banker, «Citi mandates AI prompt training for most employees». <https://www.americanbanker.com/news/citi-mandates-ai-prompt-training-for-most-employees>
- 6 FinTech Global, «BBVA rolls out ChatGPT Enterprise in global AI banking push» (12 December 2025; rollout to all 120,000 employees across 25 countries, up from 11,000). <https://fintech.global/2025/12/12/bbva-rolls-out-chatgpt-enterprise-in-global-ai-banking-push/>
- 7 Banking Dive, «Bank of America's chatbot Erica continues to amplify digital engagement» (15 July 2026; reporting Q2 2026 earnings-presentation figures for 24.6 million active users, 200 million quarterly interactions, and 70% digitally enabled sales). <https://www.bankingdive.com/news/bank-of-america-chatbot-erica-amplify-digital-engagement/825288/>
- 8 Evident Insights, «AI Use Case Trends in Banking, Q1 2026». <https://evidentinsights.com/insights/banking-use-case-trends-q1-2026>
- 9 Forrester, «DBS Bank's Billion-Dollar AI Dream, Realized» (on the control-group benchmarking methodology and the split between revenue uplift, cost savings and risk avoidance). <https://www.forrester.com/blogs/dbs-banks-billion-dollar-ai-dream-realized/>
- 10 DBS Group Holdings, «Annual Report 2025» (March 2026; FY2025 economic value of approximately S\$1 billion, 430+ use cases, 2,000+ models). <https://www.dbs.com/annualreports/2025/i/pdf/dbs-ar-2025.pdf>
- 11 Official Journal of the European Union / EUR-Lex, Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 amending Regulation (EU) 2024/1689 (published 24 July 2026, in force 27 July 2026), ELI: reg/2026/1744. <https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng>
- 12 Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin), «Überwachung von KI: BaFin erhält neue Kompetenzen» (press release, 29 July 2026; German AI Act implementing law, transparency duties from 2 August 2026, high-risk from 2 December 2027). https://www.bafin.de/SharedDocs/Veroeffentlichungen/DE/Pressemitteilung/2026/pm_2026_07_29_ki_verordnung.html
- 13 Board of Governors of the Federal Reserve System, «SR 26-2: Supervisory Guidance on Model Risk Management» (17 April 2026; issued in parallel as OCC Bulletin 2026-13 and FDIC FIL-15-2026). <https://www.federalreserve.gov/supervisionreg/srletters/SR2602.pdf>

-
- 14 European Central Bank, letter to the CEOs of significant institutions SSM-2026-0301, «Addressing AI-enabled cybersecurity threats» (7 July 2026). https://www.bankingsupervision.europa.eu/press/letterstobanks/shared/pdf/2026/ssm.2026_letter_on_AI_enabled_cybersecurity_threats.en.pdf
- 15 Bankier.pl, «W tym banku klientów obsługuje sztuczna inteligencja. AI odpowiada na pytania, składa wnioski i robi przelewy». <https://www.bankier.pl/smart/w-tym-banku-klientow-obsloguje-sztuczna-inteligencja-ai-odpowiada-na-pytania-sklada-wnioski-i-robi-przelewy>
- 16 PRNews.pl, «Bank Millennium wprowadza pierwszego w Polsce chatbota realizującego transakcje». <https://prnews.pl/bank-millennium-wprowadza-pierwszego-polsce-chatbota-realizujacego-transakcje-439533>
- 17 CyberDefence24, «Innowacyjne rozwiązanie Banku Millennium. Chatbot jak pracownik banku». <https://cyberdefence24.pl/biznes-i-finance/innowacyjne-rozwiazanie-banku-millennium-chatbot-jak-pracownik-banku>
- 18 OCBC, «Annual Report 2024: Creating Value Through AI, Transforming for the Future». <https://www.ocbc.com/group/investors/annual-reports/2024-annual-report/creating-value-through-ai.page>
- 19 Bank of America, «Fast Facts» (Erica capabilities; more than 24.6 million clients, 3.6 billion cumulative interactions since 2018, more than 200 million interactions per quarter, and 97% employee adoption of Erica for Employees). <https://newsroom.bankofamerica.com/content/newsroom/company-overview/bank-of-america-fast-facts.html>
- 20 PKO Bank Polski, «PKO Bank Polski ma pierwszego agenta AI i szybciej odpowie na maile klientów» (5 August 2026; sovereign agent architecture, planned 30-agent expansion, 16,000+ szukAI users, in-house models, and RoBERTa downloads). <https://bankomania.pkobp.pl/finanse/pko-bank-polski-ma-pierwszego-agenta-ai-i-szybciej-odpowie-na-maile-klientow/>
- 21 Barchart / Bank of America, «BofA AI and digital innovations fuel 30 billion client interactions» (March 2026). <https://www.barchart.com/story/news/665679/bofa-ai-and-digital-innovations-fuel-30-billion-client-interactions>
- 22 OpenAI, «Nubank elevates customer experiences with OpenAI». <https://openai.com/index/nubank/>
- 23 Klarna, «Klarna AI assistant handles two-thirds of customer service chats in its first month» (27 February 2024; the company's claims for chat share, 700 full-time-agent work-equivalents, satisfaction, resolution time, and projected savings). <https://www.klarna.com/international/press/klarna-ai-assistant-handles-two-thirds-of-customer-service-chats-in-its-first-month/>
- 24 Reuters, «Sweden's Klarna says AI chatbots help shrink headcount» (27 August 2024; reporting the chief executive's account of the hiring freeze, exceptions for engineering roles, and workforce reduction primarily through attrition). <https://www.reuters.com/technology/artificial-intelligence/swedens-klarna-says-ai-chatbots-help-shrink-headcount-2024-08-27/>

-
- 25 Klarna Group plc, Form F-1 Registration Statement, US Securities and Exchange Commission (filed March 2025; full-time employees of 5,527, 4,352 and 3,422 at end-2022, 2023 and 2024, the stated reason for the reduction, and 96% generative-AI use among employees as of 31 August 2024). <https://www.sec.gov/Archives/edgar/data/2003292/000162828025012824/klarnagroupplcf-1.htm>
- 26 Bloomberg Television, «Klarna CEO: We're Giving AI More Customer Service Work, Not Less» (8 May 2025; interview on service quality and the pilot recruitment of flexible human customer-service workers). <https://www.youtube.com/watch?v=pNd4KlccWAY>
- 27 CX Dive, «Klarna says its AI agent is doing the work of 853 employees» (20 November 2025; 853 agent-equivalents, \$60m saved, two-thirds of inquiries, NPS 73, and Q3 service and operations costs of \$50m against \$42m a year earlier). <https://www.customerexperiencedive.com/news/klarna-says-ai-agent-work-853-employees/805987/>
- 28 CMSWire, «Bank of America's AI Playbook: What CX Leaders Can Learn From Erica's Long Game». <https://www.cmswire.com/digital-experience/bank-of-americas-ai-playbook-what-cx-leaders-can-learn-from-ericas-long-game/>
- 29 MojeBankowanie.pl, «Bank Millennium» (Instytucja Roku interview). <https://mojebankowanie.pl/instytucjaroku-pl/bank-millennium-15>
- 30 Monzo, «Engineering the Future of Customer Operations: The Monzo Ops Agent». <https://monzo.com/blog/engineering-the-future-of-customer-operations-the-monzo-ops-agent>
- 31 Praxikon, «The Digital Omnibus and the postponement of high-risk obligations to December 2027» (Regulation (EU) 2026/1744). <https://www.praxikon.com/en/posts/digital-omnibus-high-risk-postponement-december-2027>
- 32 Hong Kong Monetary Authority, «Consumer Protection in respect of Use of Generative Artificial Intelligence» (circular, August 2024). <https://brdr.hkma.gov.hk/eng/doc-ldg/docld/20241107-1-EN>
- 33 Fortune, «How Citi's CTO is rolling out new gen AI productivity tools to more employees across the globe» (April 2025). <https://fortune.com/2025/04/09/how-citis-cto-is-rolling-out-new-gen-ai-productivity-tools-to-more-employees-across-the-globe/>
- 34 Fox Business, «Goldman Sachs rolls out firmwide AI assistant to boost employee productivity». <https://www.foxbusiness.com/technology/goldman-sachs-announces-firmwide-launch-ai-assistant>
- 35 OCBC, «OCBC is first Singapore bank to roll out generative AI chatbot to all employees globally» (press release 24 October 2023; rollout to all 30,000 employees in November 2023; 50% faster task completion inclusive of output-checking time). <https://www.ocbc.com/group/media/release/2023/ocbc-is-first-singapore-bank-to-roll-out-generative-ai-chatbot-to-all-employees-globally.page>
- 36 PKO Bank Polski, «Tysiące pracowników PKO Banku Polskiego korzysta ze sztucznej inteligencji. Technologię w banku rozwija Centrum Kompetencji AI» (November 2025). <https://www.pkobp.pl/media/aktualnosci/informacje-o-banku/tysiacze-pracownikow-pko-banku-polskiego-korzysta-ze-sztucznej-inteligencji-technologie-w-banku-rozwija-centrum-kompetencji-ai>
- 37 The Register, «Bank reverses decision to replace 45 staff with chatbot» (August 2025). https://www.theregister.com/2025/08/22/commonwealth_bank_chatbot_fail_rehiring/
- 38 OpenAI, «Morgan Stanley uses AI evals to shape the future of financial services». <https://openai.com/index/morgan-stanley/>

-
- 39 Morgan Stanley, «Launch of AI @ Morgan Stanley Debrief». <https://www.morganstanley.com/press-releases/ai-at-morgan-stanley-debrief-launch>
- 40 CIO Dive, «Banks accelerate AI deployments as agentic tools gain traction» (Evident data). <https://www.ciodive.com/news/banks-accelerate-ai-adoption-agentic-automation-evident-insights/757463/>
- 41 Google Cloud, «How Wells Fargo is using Google Cloud AI to empower its workforce with agentic tools» (August 2025). <https://cloud.google.com/blog/topics/financial-services/wells-fargo-agentic-ai-agentspace-empowering-workers/>
- 42 Banking Dive, «Wells Fargo taps Google Cloud for agentic AI push» (BNY Eliza and Goldman Sachs Devin disclosures). <https://www.bankingdive.com/news/wells-fargo-google-cloud-agentic-AI/757057/>
- 43 Agentic AI Foundation / PR Newswire, «Agentic AI Foundation Welcomes 57 New Members, Gaining Major Financial Services Players and APAC Leaders» (12 August 2026; Wells Fargo joins as a gold member, citing open collaboration and interoperability). <https://www.prnewswire.com/news-releases/agentic-ai-foundation-welcomes-57-new-members-gaining-major-financial-services-players-and-apac-leaders-302850143.html>
- 44 wnp.pl, «Największy polski bank wdrożył do pracy pierwszego agenta AI. Niebawem będzie miał ich trzydzieści» (2026; the thirty planned agents and their intended domains). <https://www.wnp.pl/tech/najwiekszy-polski-bank-wdrozyl-do-pracy-pierwszego-agenta-ai-niebawem-bedzie-mial-ich-trzydziesci,1087592.html>
- 45 Monetary Authority of Singapore, «Consultation Paper on Guidelines on Artificial Intelligence Risk Management» (13 November 2025; consultation closed 31 January 2026). https://www.mas.gov.sg/-/media/mas-media-library/publications/consultations/bd/2025/final_consultation_paper_on_guidelines_on_ai_risk_management_forrelease.pdf
- 46 International Monetary Fund, «How Agentic AI Will Reshape Payments», IMF Note 2026/004 (22 April 2026; delegated authority, Know Your Agent). <https://www.imf.org/en/publications/imf-notes/issues/2026/04/22/how-agentic-ai-will-reshape-payments-575560>
- 47 HM Treasury, «Financial Services AI Adoption Plan» (14 July 2026; ten recommendations, Critical Third-Party designation, agentic payments readiness, UK consumer adoption and trust figures). <https://www.gov.uk/government/publications/ai-adoption-plan-financial-services/financial-services-ai-adoption-plan>
- 48 Hong Kong Monetary Authority, «HKMA announces second cohort of GenA.I. Sandbox to advance responsible A.I. innovation» (October 2025). <https://www.hkma.gov.hk/eng/news-and-media/press-releases/2025/10/20251015-4/>
- 49 VentureBeat, «Wells Fargo's AI assistant just crossed 245 million interactions – no human handoffs, no sensitive data exposed» (April 2025). <https://venturebeat.com/ai/wells-fargos-ai-assistant-just-crossed-245-million-interactions-with-zero-humans-in-the-loop-and-zero-pii-to-the-llm>
- 50 Wells Fargo Newsroom, «Wells Fargo announces expansion of strategic relationship with Google Cloud». <https://newsroom.wf.com/news-releases/news-details/2025/Wells-Fargo-announces-expansion-of-strategic-relationship-with-Google-Cloud/default.aspx>

-
- 51 American Banker, «Wells Fargo plans to unleash Google's agentic AI bank-wide» (5 August 2025; rollout of generative and agentic AI to all 215,000 employees). <https://www.americanbanker.com/news/wells-fargo-plans-to-unleash-googles-agentic-ai-bank-wide>
- 52 Sullivan & Cromwell, «Federal Banking Agencies Issue Revised Guidance on Model Risk Management» (April 2026; on the announced but unissued interagency request for information covering generative and agentic AI). <https://www.sullcrom.com/insights/memo/2026/April/OCC-Fed-FDIC-Issue-Revised-Guidance-Model-Risk-Management>
- 53 GARP, «SR 11-7 in the Age of Agentic AI: Where the Framework Holds – and Where It Strains». <https://www.garp.org/risk-intelligence/operational/sr-11-7-age-agentic-ai-260227>
- 54 Skadden, «CFPB Applies Adverse Action Notification Requirement to Artificial Intelligence Models». <https://www.skadden.com/insights/publications/2024/01/cfpb-applies-adverse-action-notification-requirement>
- 55 Davis Polk, «FTC announces new enforcement initiative targeting deceptive AI practices». <https://www.davispolk.com/insights/client-update/ftc-announces-new-enforcement-initiative-targeting-deceptive-ai-practices>
- 56 Hunton, «Colorado AI Act Amended and Effective Date Delayed». <https://www.hunton.com/privacy-and-cybersecurity-law-blog/colorado-ai-act-amended-and-effective-date-delayed>
- 57 European Commission, «AI Act | Shaping Europe's digital future». <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- 58 Law and Technology, «AI literacy: the Digital Omnibus rewrites Article 4 of the AI Act» (on the shift from an obligation of result to an obligation of effort). <https://lawandtechnology.eu/en/ai-literacy-digital-omnibus-article-4-ai-act/>
- 59 White & Case, «EU AI Omnibus enters into force, amending the AI Act» (on the new Annex XIV nomenclature codes and code AIH 0401 for agentic AI). <https://www.whitecase.com/insight-alert/eu-ai-omnibus-enters-force-amending-ai-act>
- 60 European Banking Authority, «AI Act: implications for the EU banking and payments sector» (factsheet, November 2025). <https://www.eba.europa.eu/sites/default/files/2025-11/d8b999ce-a1d9-4964-9606-971bbc2aaf89/AI%20Act%20implications%20for%20the%20EU%20banking%20sector.pdf>
- 61 European Central Bank Banking Supervision, Supervisory Newsletter, «AI's impact on banking: use cases for credit scoring and fraud detection» (20 November 2025). https://www.bankingsupervision.europa.eu/press/supervisory-newsletters/newsletter/2025/html/ssm.nl251120_1.en.html
- 62 European Systemic Risk Board, «Warning ESRB/2026/3 on systemic cyber risks stemming from frontier artificial intelligence models» (adopted 25 June 2026; announced 7 July 2026; published in OJ C, C/2026/3795, 16 July 2026). https://www.esrb.europa.eu/pub/pdf/warnings/esrb.warning260625_on_systemic_cyber_risks_stemming_from_frontier_ai_models~ef424708cf.en.pdf
- 63 European Systemic Risk Board, «Frontier AI models could strain cyber resilience in the financial system, ESRB warns» (press release, 7 July 2026; systemic cyber risk assessed as "severe" in June 2026, up from "elevated" in March). <https://www.esrb.europa.eu/news/pr/date/2026/html/esrb.pr260707-4e1b68241a.en.html>

-
- 64 Joint Committee of the European Supervisory Authorities, «ESA Statement: Toward a consistent and risk-based approach for ICT risks from frontier AI models» (JC 2026 25, 31 July 2026). https://www.esma.europa.eu/sites/default/files/2026-07/JC_2026_25_ESA_statement_on_frontier_AI_models.pdf
- 65 BCLP, «AI Regulation in Financial Services: Turning Principles into Practice». <https://www.bclplaw.com/en-US/events-insights-news/ai-regulation-in-financial-services-turning-principles-into-practice.html>
- 66 Financial Conduct Authority, «FCA announces second cohort for AI Live Testing». <https://www.fca.org.uk/news/press-releases/fca-announces-second-cohort-ai-live-testing>
- 67 House of Commons Treasury Committee, «Artificial intelligence in financial services» (20 January 2026; 84 written submissions, findings on stress testing, Critical Third Parties and the Senior Managers Regime). <https://committees.parliament.uk/committee/158/treasury-committee/news/211401/current-approach-to-ai-in-financial-services-risks-serious-harm-to-consumers-and-wider-system/>
- 68 Stephenson Harwood, «The Bank of England and the FCA's survey on AI in UK financial services». <https://www.stephensonharwood.com/insights/the-bank-of-england-and-the-fcas-survey-on-ai-in-uk-financial-services/>
- 69 BBVA, «BBVA expands its agreement with OpenAI to 11,000 ChatGPT licences for the bank's employees». <https://www.bbva.com/en/innovation/bbva-expands-its-agreement-with-openai-to-11000-chatgpt-licences-for-the-banks-employees/>
- 70 BBVA, «BBVA drives AI adoption through talent» (30 April 2026; global agreements with both OpenAI and Google, and average monthly usage of ChatGPT and Gemini per employee). <https://www.bbva.com/en/innovation/bbva-drives-ai-adoption-through-talent/>
- 71 NatWest Group, «NatWest & OpenAI collaborate to accelerate cutting-edge AI transformation» (March 2025). <https://www.natwestgroup.com/news-and-insights/news-room/press-releases/ai-and-data/2025/mar/natwest-open-ai-collaborate-to-accelerate-cutting-edge-ai-transf.html>
- 72 OPI-PIB, «Model Polish RoBERTa 8k od OPI i PKO BP pomaga walczyć z hejtem» (RoBERTa 8k as the basis for the Sójka / Bielik Guard safety model). <https://opi.org.pl/model-polish-roberta-od-opi-pomaga-walczy-z-hejtem/>
- 73 bank.pl, «Bankowość przyszłości już teraz w Banku Millennium». <https://bank.pl/bankowosc-przyszlosci-juz-teraz-w-banku-millennium/>
- 74 Sztuczna Inteligencja (OPI), «Banki patrzą w przyszłość (i widzą SI)». <https://www.sztuczna inteligencja.org.pl/banki-patrze-w-przyszlosc-i-widza-si/>
- 75 Jan Kocoń, Maciej Piasecki, Arkadiusz Janz et al., «PLLuM: A Family of Polish Large Language Models» (arXiv 2511.03823, 5 November 2025; consortium technical report covering the 140-billion-token corpus and the 18 open-access models). <https://arxiv.org/abs/2511.03823>

-
- 76 PLLuM, «Projekty» (official consortium page; the eight member institutions and the intelligent citizen-assistant objective for public administration). <https://pllum.org.pl/projects>
- 77 bank.pl, «Bielik to nie jedyny polski model językowy». <https://bank.pl/bielik-to-nie-jedyny-polski-model-jezykowy/>
- 78 SpeakLeash / AGH, «Bielik: technical report» (arXiv, Open PL LLM Leaderboard results). <https://arxiv.org/pdf/2505.02410>
- 79 Ministerstwo Cyfryzacji, «Ustawa o systemach AI – bezpieczny rozwój sztucznej inteligencji w Polsce» (2026). <https://www.gov.pl/web/cyfryzacja/ustawa-o-systemach-ai--bezpieczny-rozwoj-sztucznej-inteligencji-w-polsce>
- 80 Ustawa z dnia 3 lipca 2026 r. o systemach sztucznej inteligencji (Dz.U. 2026 poz. 1003), ISAP – Internetowy System Aktów Prawnych. <https://isap.sejm.gov.pl/isap.nsf/DocDetails.xsp?id=WDU20260001003>
- 81 Kancelaria Żuraw & Partners, «Ustawa o systemach sztucznej inteligencji podpisana przez Prezydenta». <https://zglegal.pl/ustawa-o-systemach-sztucznej-inteligencji-podpisana-przez-prezydenta/>
- 82 CyberDefence24, «Ustawa o systemach AI. Organ nadzoru ruszy w tym roku». <https://cyberdefence24.pl/polityka-i-prawo/polska/sztuczna-inteligencja-pod-nadzorem-nowa-komisja-ruszy-jeszcze-w-tym-roku>
- 83 Digital HR, «Ustawa o systemach sztucznej inteligencji opublikowana» (staged commencement of KRiBSI's control and penalty provisions). <https://digitalhr.pl/nawosci/ustawa-o-systemach-sztucznej-inteligencji-opublikowana-opracowanie-digital-hr/>
- 84 Komisja Nadzoru Finansowego, «Rekomendacje dla podmiotów rynku finansowego dotyczące wpływu modeli Frontier AI na cyberbezpieczeństwo» (27 July 2026). https://www.knf.gov.pl/?articleId=99168&p_id=18
- 85 ITwiz, «UKNF ostrzega sektor finansowy przed AI. Firmy muszą na nowo zdefiniować cyberbezpieczeństwo». <https://itwiz.pl/uknf-ostrzega-sektor-finansowy-przed-ai-firmy-musza-na-nowo-zdefiniowac-cyberbezpieczenstwo/>
- 86 DBS, «DBS' AI-Powered Digital Transformation» (the 2022 baseline: SGD 180 million of economic value, 600+ models and 300 use cases). <https://www.dbs.com/artificial-intelligence-machine-learning/artificial-intelligence/dbs-ai-powered-digital-transformation.html>
- 87 The Fifth Person, «10 things I learned from the 2026 DBS Group AGM» (April 2026; two-thirds of employees using DBS-GPT). <https://fifthperson.com/2026-dbs-group-agm/>
- 88 Fortune, «OCBC rolls out AI-enabled "avatar banking", still pledges to hire 600 new advisors» (July 2026; 600 relationship managers, annual AI spend, CEO Tan Teck Long on workforce). <https://fortune.com/2026/07/01/ocbc-avatar-banking-ai-wendy-wayne-ceo-tan-teck-long/>

-
- 89 Fortune, «DBS Singapore to cut temp jobs as AI replaces roles» (25 February 2025; about 4,000 temporary and contract roles over three years, roughly 1,000 new AI roles, 13,000 employees to be reskilled, and Piyush Gupta on repurposing staff). <https://fortune.com/asia/2025/02/25/dbs-singapore-to-cut-temp-jobs-ai-piyush-gupta/>
- 90 Monetary Authority of Singapore, «Emerging Risks and Opportunities of Generative AI for Banks» (Project MindForge). <https://www.mas.gov.sg/-/media/mas-media-library/schemes-and-initiatives/ftig/project-mindforge/emerging-risks-and-opportunities-of-generative-ai-for-banks.pdf>
- 91 Allen & Gledhill, «MAS consults on proposed Guidelines for Artificial Intelligence Risk Management». <https://www.allenandgledhill.com/sg/publication/articles/31741/mas-consults-on-proposed-guidelines-for-artificial-intelligence-risk-management>
- 92 KPMG India, «RBI FREE-AI Committee report on Framework for Responsible and Ethical Enablement of Artificial Intelligence» (26 August 2025; 26 recommendations under six pillars and seven sutras). <https://kpmg.com/in/en/insights/2025/08/rbi-free-ai-committee-report-on-framework-for-responsible-and-ethical-enablement-of-artificial-intelligence.html>
- 93 Hong Kong Monetary Authority, «Regulators launch GenA.I. Sandbox++ to foster A.I. innovation across financial services» (March 2026). <https://www.hkma.gov.hk/eng/news-and-media/press-releases/2026/03/20260305-3/>
- 94 ASIC, «24-238MR ASIC warns governance gap could emerge in first report on AI adoption by licensees». <https://www.asic.gov.au/about-asic/news-centre/find-a-media-release/2024-releases/24-238mr-asic-warns-governance-gap-could-emerge-in-first-report-on-ai-adoption-by-licensees/>
- 95 MinterEllison, «ASIC urges stronger AI governance for AFS and credit licensees» (REP 798 detail; 43% of licensees had policies referencing disclosure of AI use to affected consumers). <https://www.minterellison.com/articles/asic-urges-stronger-ai-governance-for-afs-and-credit-licensees>
- 96 Evident Insights, «AI Outcomes Report» (methodology of the Use Case Tracker). <https://evidentinsights.com/insights/outcomes-report/>
- 97 Financial Stability Board, «Monitoring Adoption of Artificial Intelligence and Related Vulnerabilities in the Financial Sector» (October 2025). <https://www.fsb.org/uploads/P101025.pdf>
- 98 Bank of England, «Financial Stability in Focus: Artificial intelligence in the financial system» (April 2025). <https://www.bankofengland.co.uk/financial-stability-in-focus/2025/april-2025>
- 99 Information Age (ACS), «CBA reverses AI-driven job cuts, admits 'error'». <https://ia.acs.org.au/article/2025/cba-reverses-ai-driven-job-cuts--admits--error-.html>
- 100 American Banker, reporting on CB Financial Services' shadow-AI breach (discovered 5 May 2026; Form 8-K Item 1.05 filed 11 May 2026). <https://www.americanbanker.com/news/a-bank-breaks-its-silence-on-its-shadow-ai-breach>

ABOUT THE AUTHOR

Polina Fomenkova

Author

Polina Fomenkova is a content strategist with over ten years of experience in tech and B2B SaaS. She specializes in content on AI, startups and technology, turning complex ideas into plain language.

Disclosure, not capability, now separates the leaders.

BBVA has taken ChatGPT Enterprise past 100,000 employees, JPMorgan Chase reaches roughly 250,000, and Citi trained 175,000 employees across 80 markets. Bank of America's Erica reached 24.6 million active users. Almost none of them publish what any of it did for containment, cost, or service quality.

Which of the four AI layers the evidence actually supports, what the EU AI Act, SR 26-2 and Poland's new supervisor now oblige, and why only 15 of the 50 largest banks disclose an aggregate return figure at all.

100 sources: Bank of England, FCA, ECB, BaFin, BBVA, DBS, Evident Insights, OpenAI.

Ayro

Polina Fomenkova
First edition, 2026
Tidio for Tidio Agents
www.getayro.ai

ISBN 978-83-68606-32-4

