EXECUTIVE WHITEPAPER

AI IN BANKING

WHEN RISK CARRIES MATERIAL CONSEQUENCES

New technologies can materially increase productivity and reduce operating costs. Without adequate controls, however, financial, operational, regulatory and reputational consequences can exceed those benefits.

Dušica Živanović

AI Governance & Control Architecture | High-Stakes AI Systems | Execution Risk

October 2026

DECISION ARCHITECTURE

assessment.dusicazivanovic.com

Executive Summary

AI adoption in financial services is already producing measurable process-level productivity gains, while enterprise-level value remains harder to evidence. The 2026 Cambridge Centre for Alternative Finance report finds positive productivity impacts highest in technology, data and product functions (79%), followed by back-office and operations (75%) and front-office roles (69%). At the same time, 76% of large financial institutions report difficulty measuring the value of AI deployment, and only 34% of traditional financial institutions report increased profitability from AI.

The evidence therefore supports a distinction between **operational efficiency** and **enterprise financial return**. Material benefits are already visible in individual banking processes, but material failures can create financial loss, operational disruption, regulatory exposure and reputational damage at the same time.

Use of material. *In this whitepaper, material consequences are consequences capable of significantly affecting the institution's financial position, operations, regulatory obligations, customers or reputation. The term is not used as a universal legal, regulatory or accounting threshold.*

Evidence of value: selected 2026 banking disclosures

Area	AI application	Reported positive effect	Source
Fraud detection / payment monitoring	Behavioural analysis and transaction anomaly detection	Bank of Ireland: customer payment-fraud losses −25%; false positives −87.1%; alert volumes −85%; case-handling time halved.	Bank of Ireland, 19 Jun 2026 [2]
Transaction screening / financial crime	AI screening and case prioritisation	JPMorganChase: more than 2× screening volume while halving manual operator checks.	JPMorganChase Annual Report 2025, published 2026 [3]
KYC / Source of Wealth	Agentic AI embedded in KYC workflow	Bank of Singapore / OCBC: Source-of-Wealth preparation reduced from up to 10 days to about 1 hour.	OCBC, 29 Jul 2026; Annual Report [4]
Customer service	GenAI copilots supporting service agents	Bank of America: average call time nearly 1 minute shorter. BBVA Italy/Germany: average handling time reduced by >15%.	Bank of America, 21 Jul 2026; BBVA, 31 Jul 2026 [5][6]
Software development	Coding assistants, testing, review and refactoring	OCBC: coding/testing effort −20–30%. HSBC: 15% efficiency in coding time. Deutsche Bank: 1.5–2.5 hours/week saved; 10× gains on some specific tasks.	OCBC; HSBC; Deutsche Bank, 2026 [7][8][9]
AI development / deployment	MLOps and automated governance/development components	BBVA: development time reduced by 20–75% across four pilot initiatives, depending on the use case and initial MLOps maturity.	BBVA, 23 Jul 2026 [10]
Corporate credit assessment	Agentic AI prepares review-ready first-draft credit memo	DBS: >70 tasks within workflow; target ≥30% reduction in time spent on credit work. Target, not realised KPI.	DBS, 19 Aug 2026 [11]
Internal Audit	AI-supported data validation, testing, code and documentation	BBVA: expected ~10% productivity improvement in data-intensive audits. Projection, not realised KPI.	BBVA, 10 Mar 2026 [12]

When the Benefit Becomes Exposure

A material failure can impose costs well beyond the process in which efficiency gains were achieved. The examples below are deliberately different: an AI-enabled external fraud, an automated control failure, and unauthorised use of an AI application with customer data.

The examples distinguish risks created by external use of AI from risks arising when automated or AI-enabled capabilities operate inside the institution.

Observed event / practice	What occurred	Consequence / exposure	Source
Fideuram / Intesa Sanpaolo — Italy	AI-enabled impersonation was used to induce a series of international transfers.	Approximately €95m transferred; around €53m recovered; about €36m remained unaccounted for at the reporting date. Investigation ongoing.	Reuters, 25 Sep 2026 [13]
Community Federal Savings Bank — US	Automated suspicious-activity alert triage had deficiencies in logic, data and methodology; alerts that should have escalated were auto-closed.	OCC enforcement action and required BSA/AML remediation. OCC does not identify the automated system as AI.	OCC Consent Order, Apr 2026 [14]
Community Bank / CB Financial Services — US	Non-public customer information was handled using an unauthorised AI-based software application.	Material cybersecurity incident; disclosed data included customer names, Social Security numbers and dates of birth; containment, investigation, notifications and strengthened controls followed.	SEC Form 8-K, May 2026 [15]

The Consequence Is Not One-Dimensional

AI risk does not materialise within organisational risk silos. The same control failure can produce direct financial loss, operational remediation, regulatory or legal exposure and reputational damage. EBA's June 2026 Risk Assessment Report identifies operational failures, bias, data-quality issues, algorithmic-trading/procyclicality risk, conduct and legal risk, enhanced cyberattacks, third-party reliance and geopolitical risk among AI-related exposures, and notes that materialisation can result in financial loss, reputational damage and potentially systemic implications.

Initial failure	Financial	Operational	Regulatory / legal	Reputational
Fraudulent execution	Direct loss; recovery costs	Investigation; recovery; control redesign	Control / compliance consequences	Loss of confidence
Incorrect automated decision	Loss; compensation; remediation	Reprocessing; model/process redesign	Conduct, discrimination or decision requirements	Customer harm / trust
Unauthorised data access or use	Remediation; potential penalties	Data/system redesign; containment	Privacy, data protection, AI governance	Loss of customer trust
Inadequate AI deployment	Implementation and remediation cost	Workflow disruption; rollback; added controls	Depends on affected regulated process	Customer / market confidence
Third-party failure	Loss or revenue exposure	Service interruption; substitution	ICT / outsourcing / DORA exposure	Service-continuity confidence

Context. As AI becomes embedded in banking processes, technology-related failures can transmit across financial, operational, regulatory and reputational risk categories. EBA reports approximately 3.9 million operational-risk loss events in EU/EEA banks in 2025 and €15.2bn of materialised operational-risk losses. These

figures are not AI losses; they illustrate the material scale of operational failure in the banking environment into which AI is being deployed.

The Control Question

How can institutions capture the productivity benefits of AI without granting more information access, decision influence or execution authority than their control environment can support?

External threat

How can an institution prevent an external AI-enabled actor from producing a valid financial action with material consequences?

Internal authority

How can an institution prevent an AI system from accessing, processing, influencing or executing beyond its authorised environment?

Control Principles

This whitepaper proposes four control principles for assessing whether material AI authority can be supported.

01	**Define the permissible information space** What the system may access, process and use, and which relationships may legitimately influence an outcome.
02	**Define decision and execution authority** What the AI may recommend, decide or execute; authority thresholds; required human or independent approval.
03	**Establish execution-time controls** Conditions that must be satisfied before a material action can be executed; identity, authority and transaction-specific validation where relevant.
04	**Monitor boundaries and escalation** Independent runtime monitoring, boundary violations, stop conditions, escalation and evidence.

INFORMATION → INFLUENCE → DECISION → EXECUTION

What is permitted? | What is authorised? | What must be verified?

MONITORING · CONTROL · ESCALATION · EVIDENCE

What the Organisation Must Be Able to Answer

Before AI authority is expanded, the organisation must be able to demonstrate that the process, data, controls and operating architecture can support it.

The assessment begins with the questions below and becomes progressively more specific as authority and consequence increase.

READINESS BY DIMENSION

Dimension	Question the organisation must answer
Business justification	Is the business and economic case sufficiently defined?
Nature of work	What work is AI expected to perform: analysis, assistance, decision-making or execution?
Process stability	Is the underlying process sufficiently stable and defined to support AI involvement?
Process interdependence	Which upstream and downstream processes can affect, or be affected by, the AI-enabled process?
Data readiness	Are the required data available, reliable, legally usable and sufficiently controlled?
Initial process suitability	Is this process suitable for the requested level of AI involvement and authority?
Control readiness	Can the existing control environment support the consequences of the authority being requested?
Build / operate	Who builds or operates the AI component?
Interaction	Who interacts with the AI?

HOW THE ASSESSMENT IS DEFINED

Assessment stage	What is defined
Process Risk Classification	Inherent risk, control effectiveness, residual risk and assessment intensity.
Process & Authority	Purpose, role, authority and consequence classifications.
Data & Controls	Automation category, inherent risk, preliminary data-use result and required action.
Architecture Check	Architecture gap assessment.

The purpose is not to maximise AI autonomy. It is to determine what level of AI authority the organisation can demonstrate that its process, data and control environment can support.

Decision Architecture Assessment

The relevant governance decision is not whether AI is acceptable in the abstract. It is what level of information access, influence, decision authority and execution a specific process can support, under which conditions, and on what evidence.

My work applies Decision Architecture to planned and existing AI-enabled processes. The assessment connects business justification, process stability, data eligibility, control readiness, authority and third-party exposure to identify the supported authority boundary, the open control conditions and the evidence required before that boundary can be expanded.

Institutions considering a new use case, moving a pilot toward production or reviewing an existing deployment may submit the process for an initial assessment.

Request an initial assessment → assessment.dusicazivanovic.com

Sources

[1] Cambridge Centre for Alternative Finance, 2026 Global AI in Financial Services Report: Adoption, Impact and Risks, Apr 2026. Productivity: technology/data/product 79%; back office/operations 75%; front office 69%; 76% of large FIs report difficulty measuring AI value; 34% of traditional FIs report increased profitability.
https://www.jbs.cam.ac.uk/faculty-research/centres/alternative-finance/publications/2026-global-ai-in-financial-services-report/

[2] Bank of Ireland, "AI technology sees customer 'payment fraud' losses at Bank of Ireland drop 25%", 19 Jun 2026.
https://www.bankofireland.com/about-bank-of-ireland/press-releases/2026/ai-technology-sees-customer-payment-fraud-losses-at-bank-of-ireland-drop-25/

[3] JPMorganChase, Annual Report 2025 / CIB shareholder letter, published 2026: transaction screening volume more than doubled while manual operator checks were halved.
https://www.jpmorganchase.com/ir/annual-report/2025/ar-ceo-letter-petno-rohrbaugh

[4] OCBC / Bank of Singapore, "OCBC harnesses agentic AI to sharpen and quicken onboarding of wealthy customers", 29 Jul 2026; OCBC Annual Report 2025.
https://www.ocbc.com/group/media/release/2026/ocbc-harnesses-agentic-ai-to-quicken-onboarding-customers.page

[5] Bank of America, "Bank of America Enhances EricaAssist with Generative AI to Help Employees Resolve Client Needs Faster", 21 Jul 2026.
https://newsroom.bankofamerica.com/content/newsroom/press-releases/2026/07/bank-of-america-enhances-ericaassist-with-generative-ai-to-help-.html

[6] BBVA, "BBVA Cuts Customer Inquiry Response Times by More Than 15% in Italy and Germany Thanks to AI", updated 31 Jul 2026.
https://www.bbva.com/en/innovation/artificial-intelligence/

[7] OCBC Annual Report 2025, Operations and Technology: AI reduced coding and testing effort by 20–30%.
https://www.ocbc.com/group/investors/annual-reports/2025-annual-report/year-in-review/operations-and-technology.page

[8] HSBC, "Transforming HSBC with AI": more than 20,000 developers use coding assistants; 15% efficiency in time spent coding.
https://www.hsbc.com/who-we-are/hsbc-and-digital/hsbc-and-ai/transforming-hsbc-with-ai

[9] Deutsche Bank, "How AI is reshaping the software behind everyday banking", 14 Jul 2026: 1.5–2.5 hours/week saved; 10× gains on a number of specific tasks.
https://www.db.com/news/detail/20260714-how-ai-is-reshaping-the-software-behind-everyday-banking?language_id=1

[10] BBVA, "MLOps transformation at BBVA: from data platform to scalable ML", 23 Jul 2026: development time reduced by 20–75% across four pilot initiatives, depending on the use case and initial MLOps maturity.
https://www.bbva.com/en/innovation/mlops-transformation-at-bbva-from-data-platform-to-scalable-ml/

[11] DBS, "DBS scales agentic AI to transform way of working for corporate bankers", 19 Aug 2026: >70 tasks; ≥30% time-reduction goal.
https://www.dbs.com/newsroom/DBS_scales_agentic_AI_to_transform_way_of_working_for_corporate_bankers_freeing_up_time_for_more_strategic_client_engagements

[12] BBVA, "How BBVA Uses An AI Assistant to Analyze Data in Internal Audit", 10 Mar 2026: expected ~10% productivity improvement in data-intensive audits.
https://www.bbva.com/en/innovation/how-bbva-uses-an-ai-assistant-to-analyze-data-in-internal-audit/

[13] Reuters, "AI messaging scam costs Italy's top bank Intesa millions, sources say", 25 Sep 2026.
https://www.reuters.com/legal/government/ai-messaging-scam-costs-italys-top-bank-intesa-millions-sources-say-2026-09-25/

[14] Office of the Comptroller of the Currency, Consent Order against Community Federal Savings Bank, dated 24 Apr 2026 and released in May 2026. The order describes an automated alert triage system; it does not identify that system as AI.
https://www.occ.gov/static/enforcement-actions/eaAA-ENF-2025-21.pdf

[15] CB Financial Services, Inc., SEC Form 8-K, dated 7 May and filed 11 May 2026: unauthorised AI-based software application used in handling non-public customer information.
https://www.sec.gov/Archives/edgar/data/1605301/000160530126000021/cbfv-20260507.htm

[16] European Banking Authority, Risk Assessment Report, Jun 2026: operational risk, AI-related risk categories, 3.9m loss events and €15.2bn materialised operational-risk losses in 2025.
https://www.eba.europa.eu/publications-and-media/publications/risk-assessment-report-june-2026

Evidence note: Bank-reported operational results are identified as such and have not been independently validated for this whitepaper. Measurement periods, baselines and methodologies may differ across institutions. Targets and expected productivity gains are not presented as realised outcomes. EBA operational-risk figures provide sector context and are not attributed to AI.