MAVERIC
ACCELERATE NEXT

From AI Pilots to Profit

Engineering Outcomes in Banking

Why AI stalls: the outcome gap, not a tech gap

Most banks do not fail at AI for lack of models, platforms, or funding - they fail for lack of engineered outcomes. BCG survey states that only 5% of companies are achieving AI value at scale. Another 35% are scaling up their efforts and seeing some value. Without explicit value definitions, pilots multiply but do not convert, causing pilot-to-production leakage, reactive governance, inconsistent value realization, and eroding executive confidence.

A measurable, observable outcome chain is the difference between interesting experiments and sustained enterprise value. As per Evident, ROI from AI is still years away. Many of the banks are using indicators like time saved or adoption milestones to show benefits through AI adoption. This whitepaper aims to capture the need to engineer outcomes as critical to drive AI adoption.

Focus on outcomes:

Engineering outcomes means establishing an explicit chain from AI decisions and actions to operational KPIs and financial impact, orchestrated end to end so that value delivered is both provable and repeatable. It highlights the shift from traditional output-based service delivery models to outcome based, co-investment driven partnerships.

For banking organizations, it is imperative that they anchor AI adoption in measurable value outcomes. Evident AI index in banking shows that 32 banks now report AI use cases with tangible outcomes, up more than 20% from last year.

Outcome focused metrics span quantitative as well as qualitative across areas such as operational KPIs, business KPIs and experience KPIs, responsible KPIs etc.

Operational KPIs: Categorized into time savings (process or people), cost savings (existing vs AI led processing), or quality improvement (reducing rework).

e.g., Morgan Stanley built an AI model to enable legacy code rewriting (COBOL and PERL) to make them AI ready, freeing up developer hours.

Business KPIs: Including areas that are revenue focused – cross sell, conversions and new acquisitions, aimed at fraud reduction or connected with capital requirement.

e.g., RBC's wealth advisors use AI for meeting summarization and research investments freeing advisor time for revenue upliftment.

Furthermore, business KPIs also include experience related areas like customer satisfaction scores (CSAT/NPS), interaction quality, user adoption rates etc.

e.g., Truist bank built an AI tool that enabled them to understand customer sentiment about their products, leading to cost savings on account of reduced complaints.

Responsible KPIs: Including areas like explainability scores, bias/fairness tests, audit rates, incident resolution, regulatory compliance etc.

Outcome Chains: Making Value Traceable and Auditable

Outcome chains establish clear linkages between AI implementations and measurable business outcomes, leveraging continuous observability and transparency. At its core, outcome chains create clear pathways connecting AI model outputs to intermediate process improvements, translating further into measurable business KPIs and then finally leading to quantifiable financial impact. This approach ensures that every AI implementation can directly be attributed to quantifiable value creation, eliminating the ambiguity that often plagues AI ROI calculations.

As an example for AML, the first level involves arriving at specific decision points like suspicious activity scores or alert classification. The second level shows how AI decisions translate into operational improvements like automation – saving time, cost or quality, inference or decision aid – driving forecast, remediation or recommendation and transformation – influencing the culture of data led action. Finally, these translate into financial value like investigation cost reduction or regulatory fine avoidance.

Observability and Governance: Maintaining Chain Integrity

Observability and governance is foundational to making the outcomes a reality. We need to consistently monitor the data quality, model behaviour and alignment to KPIs. Live dashboards need to be created that capture the outcome chains along with alerts for any model bias, fairness and risks involved with model hallucination. Market conditions, regulatory changes, shifts in customer behaviour and model drift can impact AI models mandating regular validation and recalibration. Real time model output can be used as feedback for iterative model refinement. Monitoring systems can benefit by setting up lead and lag indicators for proactive management of AI systems.

Equally important is the need to ensure AI systems and frameworks are aligned to regulatory and governance frameworks like ISO 27K, EU AI act, NIST AI Risk Management and Basel III implications for AI in risk management. This will ensure that value creation does not come at the expense of regulatory compliance or operational risk.

Engineering outcomes with Agentic AI, Contextual Data Catalogs and Safe by design AI

Intent towards engineering measurable business outcomes can be made a reality with the help of agentic AI to drive autonomous goal-oriented agents, supported by contextual catalog that provides trustworthy, decision grade data, on a bedrock of safe by design AI architecture that ensures governance and auditability at every step.

Agentic Approach: It involves creating AI agents that plan, breakdown tasks and execute multi step workflows closely aligned to finalised outcome metrics. In addition to building relevant connectors to core banking systems and surround applications like CRM, payments, there is also a strong need to have orchestration agents to coordinate multi agent workflows with conflict resolution and rollbacks. Presence of appropriate feedback loops also ensure that agents are constantly updated on external landscape changes, policy changes as well as fine tuning required for the models based on decision outcomes driving continuous improvement.
According to BCG survey, agents account for 17% of total AI value in 2025 and expected to almost double to 29% by 2028. Top 5% of organizations allocate 15% of their AI budget to agents, further expanding the value chasm between them and their slower rivals.

Contextual Data Catalog: Moving beyond provision of datasets to providing decision ready, business context driven data. Unified metadata and business glossary supported by lineage ensures that AI agents can benefit from having access to and use of the right, high trust data. The access to data can also be made role based or purpose based duly taking into consideration the need for privacy norms through tokenization and masking of Personally Identifiable Information (PII). AI agents can also be deployed for ensuring data catalog is current, through initiatives like auto classification, tagging, lineage validation, so that data management is seamless and invisible.

Safe by Design AI: The ability to convert risk into a design constraint rather than an afterthought by embedding security, privacy, fairness, explainability and compliance at the architecture level itself. It covers responsible AI controls through fairness tests, bias mitigation, and quick responses to any incidents on account of model errors. Provision of human in the loop oversight brings together the necessary interventions for high-risk decisions, helps obviate any out of policy actions and address any escalation paths. Close regulatory alignment and consistent observability mechanisms through action logs, decision rationales further cements confidence in model outputs mapping to business outcomes.

According to BCG, 72% of companies already report unmanaged AI security risks. It's imperative that any firm adopting AI puts in place relevant guard rails and reconfigure workflows to combine autonomous agents with human oversight.

The way forward

Financial institutions can move forward towards value realization through AI adoption through the approach outlined below.

1. Identify focus streams or domains with the high levels of business impact and technical feasibility in driving AI led transformation. Across these domains, identify measurable business outcomes, establish the outcome chain and KPIs – operational, business, financial, experience and responsibility.
2. Conceptualise the AI stack comprising of relevant AI agents, orchestrator agents, contextual data fabric and safe by design guardrails. Establish target benchmarks, intermediate milestones for success metrics and measurement mechanisms to avoid any issues attributing success. Ensure adherence to governance standards like ISO 27001, EU AI act etc.
3. Create and deploy the agentic AI workflows with human in the loop. Establish the relevant integrations to core systems and set up processes for observability and audit. Create relevant real time dashboards to monitor outcome chain and KPI adherence.
4. Monitor actively primary indicators and secondary indicators leading to KPIs. Drive A/B comparisons wherever feasible to shift gears quickly onto different models.
5. Continuously optimise, evolve data catalogs and guard rails with change in external regulations and data structures. Templatise agents, features data catalogs for repeatability. Embed feedback loops for model updates and alerts on model drift, fairness, hallucinations.

Operationalise outcomes not experiments

Closer look at the banks that are front runners in deploying AI at scale and demonstrating value from Ai initiatives as outlined in the Evident AI index, indicate a blueprint comprising of value led prioritization of usecases, outcome linked KPIs, and continuous improvement loops with due importance to explainability, governance and transparency. BCG report indicates that future built organizations are 4.6 times more likely to have fit for purpose guardrails in place and 2.6 times more likely to have rigorous tracking of AI value by setting out clear AI driven KPIs. To convert AI investments into durable improvements in business parameters like cost to income, operational efficiency and customer lifetime value, financial institutions need to look at mastering engineered outcomes without sacrificing governance and transparency.

Authored by

P Venkatesh,
Co-founder and whole-time Director

As the Co-founder and whole-time Director at Maveric, P Venkatesh (PV) leads the global thought leadership function aimed at shaping and promoting Maveric's perspectives as well as expertise in the banking technology space. By building relationships with industry influencers, partners and BankTech ecosystem leaders, PV drives creation of impactful frameworks, methodologies and landscape reports that provide informed perspectives on new age technologies that shape the BankTech space.

Anoop Karumathil Melethil,
Vice President, Marketing

Anoop has two decades of experience in conceptualizing, creating, and driving technology-focused marketing initiatives across tier 1 as well as mid-size global IT services organizations. He has built his expertise in building marketing functions from the ground up, encompassing areas including but not limited to media, public relations, influencer marketing, digital marketing, employer branding, thought leadership, brand outreach, field marketing, ABM, and internal communication.

The Maveric **Edge**

Established in 2000, Maveric Systems is a niche banking and financial services technology specialist driving operations and technology transformation with AI, data and automation. Combining deep domain knowledge with AI-powered services and solutions we solve complex CXO challenges across retail banking, corporate banking, wealth management and capital markets.

Our 25 years of domain mastery ensures that every AI initiative is rooted in contextual relevance and precision. Our approach is anchored in dedicated domain, technology and data AI Centres of Excellence. From quick wins to enterprise scale - brought to life through our knowledge frameworks, co-created at speed, with multi-model flexibility, powered by agentic orchestration, secured by governance guardrails, we deliver Engineered Outcomes built on design principles that are Agentic, Safe by Design, and Contextual.