SFTI – working group 'Implementing AI'

'Bridging the AI PoC–Production Gap – Keys to Deployment Success in Swiss Financial Industry'

Study Report

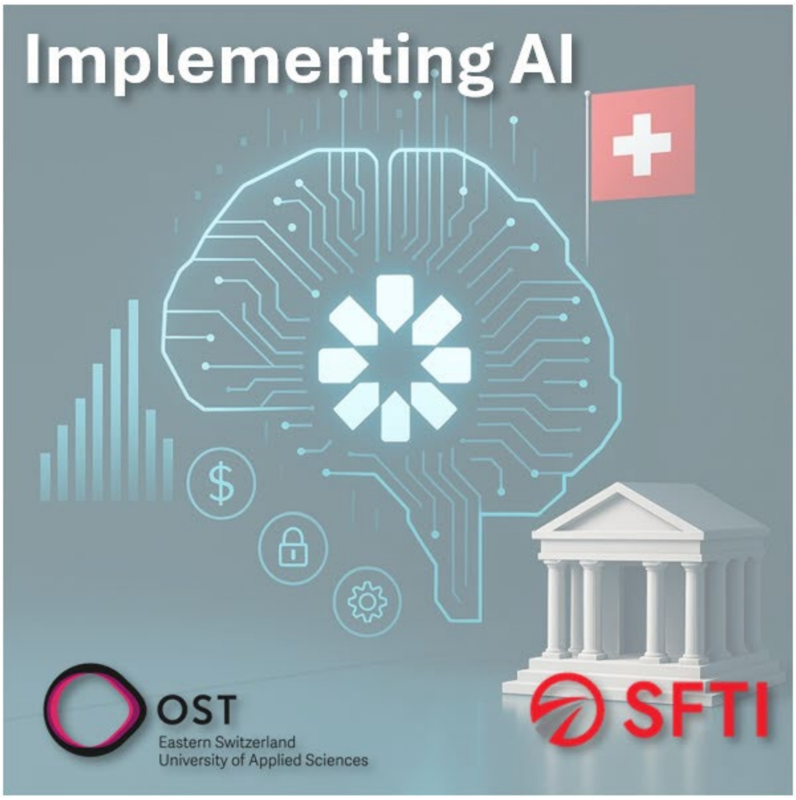

Authorship: Swiss FinTech Innovations and Eastern Switzerland University of Applied Sciences
Release: Version 1.0
Date: 29.10.2025

Authors

OST Eastern Switzerland University of Applied Sciences, Institute for Finance & Law

Stefan Neumann

SFTI (Swiss Fintech Innovations)

Stephanie Wickihalder

About OST

Eastern Switzerland University of Applied Sciences (OST) is a dynamic, innovative university that strengthens the Eastern Switzerland region with forward-looking initiatives and makes a significant contribution to its economic and social development. For more information about *OST*, please refer to http://www.ost.ch.

About SFTI

SFTI (Swiss Fintech Innovations) is an independent association of Swiss financial institutions committed to drive collaboration and digital innovations in the financial services industry. For more information about *SFTI*, please refer to http://www.sfti.ch.

Table of Contents

Management Summary

The adoption of Artificial Intelligence (AI) in the Swiss financial industry has accelerated in recent years, yet a significant share of initiatives remains trapped in pilot stages, unable to deliver sustained business impact. This follow-up study, titled *"Bridging the AI PoC–Production Gap – Keys to Deployment Success in Swiss Financial Industry"*, directly addresses this challenge. Led by OST (Eastern Switzerland University of Applied Sciences), Institute for Finance & Law, in cooperation with SFTI (Swiss FinTech Innovations) and supported by ELCA Advisory, the study provides a rigorous and practice-oriented analysis of how Swiss banks and insurers can transition AI projects from proof-of-concept (PoC) into fully productive and scalable implementations.

Objectives and Approach

The primary objective of this study is to equip Swiss financial institutions with concrete strategies, frameworks, and tools to overcome the "pilot purgatory" phenomenon. The focus lies on identifying actionable pathways to unlock AI's business value while remaining aligned with regulatory, governance, and operational realities when moving from pilot to deployment (?). Specific objectives include:

- **Validating AI business cases and expectation management** to ensure realistic ROI and stakeholder confidence.
- **Defining robust model selection, testing, and monitoring practices** to build trust and facilitate regulatory approval.
- **Addressing infrastructure and system dependencies** while ensuring data sovereignty and confidentiality.
- **Establishing governance and compliance-by-design principles** to embed trust and accelerate approval processes.
- **Developing a pragmatic success checklist and best-practice framework** for scaling AI initiatives.

The study employed a mixed-methods design: a comprehensive online survey among 24 Swiss financial institutions (including two-thirds of SFTI members) provided quantitative breadth, complemented by in-depth expert interviews offering qualitative insights. This approach ensured both strategic relevance and practical grounding.

Key Findings

1. **AI Use Case Life Cycle:** The portfolio of AI initiatives shifted upstream, with more institutions engaged in ideation and small pilots (40%), while fewer reached deployment (17%) or scaling (11%). Stricter governance and ROI discipline slowed deployments but increased success rates: 65% of production cases met or exceeded expectations.
2. **Business Case Realization:** Efficiency remains the dominant driver (100% of respondents), followed by internal user experience (71%). Cost reduction has declined as a decisive factor (42%), reflecting a shift toward productivity, quality gains, and employee enablement. Projects that reached production generally delivered on business case expectations.
3. **Decision Criteria for AI Applications:** Hosting, cybersecurity (17% each), and data sovereignty (15%) dominate vendor selection, confirming that compliance and risk management remain critical gates for deployment. Transparency and open-source flexibility are secondary but rising considerations.
4. **Roadblocks to Scaling AI Beyond PoC:** The most significant barriers are data privacy, security, and quality issues, as well as regulatory uncertainty and weak business-case proof. Technical limitations and legacy system dependencies further impede scaling, while cultural resistance and management support play a lesser role.

5. **Success Factors:** Winning deployments are driven by **data access**, **executive sponsorship**, and **robust governance frameworks**, reinforced by strong IT–business cooperation and MLOps capabilities. Testing, monitoring, and compliance-by-design approaches emerged as critical enablers.
6. **AI Model Card Expectations:** Respondents ranked **bias, risks, and limitations**, as well as **intended use and evaluation**, as the most important aspects for model documentation. Environmental impact and other non-core criteria received lower emphasis.

Scalability and Recommendations

The study demonstrates that scaling AI in Swiss finance requires a **production-first mindset**, anchored in disciplined data foundations, robust MLOps, and compliance-by-design. Institutions should introduce stage gates tied to deployability, cap concurrent PoCs, and require clear path-to-production artifacts at project inception. Business case discipline – tying outcomes to measurable KPIs and assigning accountable owners – is essential to sustain credibility.

Key recommendations include:

- **Mandate production-readiness from the start:** Require deployable architectures, governance structures, and observability at PoC sign-off.
- **Invest in data and MLOps foundations:** Shared feature stores, automated monitoring, and reproducible pipelines are decisive separators between pilots and scale.
- **Embed compliance early:** Standardized model cards and policy-by-design shorten approval cycles.
- **Focus business cases on user value and productivity:** Prioritize employee-facing AI that improves throughput, quality, and adoption.
- **Balance security with innovation:** Ensure sovereignty and compliance through hosting and licensing safeguards, while leveraging scalable vendor solutions.

Conclusion

This follow-up study confirms that while many Swiss financial institutions remain in early AI adoption stages, those that succeed in moving beyond PoCs achieve tangible business impact. By bridging the gap between pilots and production, institutions can unlock AI's transformative potential in efficiency, risk management, and user experience. The study provides a pragmatic roadmap – combining empirical survey evidence, case-based insights, and actionable tools – that positions financial industry players at the forefront of AI industrialization in Swiss financial services.

Acknowledgment

We extend our heartfelt gratitude to all individuals and organizations who contributed to this study. Your insights, expertise, and openness in sharing your experiences have been invaluable to the depth and quality of our findings.

In particular, we wish to thank **Chandra Shekaran, Christian Sebregondi, Fabian Keller, Laure Manueddu, Nicolas Zahn, Pascal Wild, Patrik Schmid, Peter Ivan, Rahul Karandikar, Richard Hess, Stefan Jeker, and Volkmar Beck** for their time, guidance, and thoughtful contributions. Your perspectives have greatly enriched the analysis and helped ensure that this report reflects the realities and ambitions of the Swiss financial industry.

We would like to express our **special appreciation to ELCA Advisory, namely Patrik Schmid and Nicolas Zahn,** for their outstanding support throughout all phases of the study. Their ongoing collaboration, expert guidance, and valuable content contributions were instrumental in shaping the quality, structure, and practical relevance of this report. This study would not have achieved its current depth and clarity without their dedicated partnership.

Thank you all for your engagement, collaboration, and commitment to advancing the responsible and effective use of AI in financial services.

-

1. Introduction and Study Approach

The increasing integration of Artificial Intelligence (AI) into the financial services sector continues to reshape operations, promising significant improvements in efficiency, compliance, and decision-making. However, many Swiss financial institutions remain challenged in moving from exploratory or pilot stages into fully productive deployments. Despite extensive experimentation with proofs-of-concept (PoCs), nearly half of initiatives remain stuck at the pilot stage, with only a minority reaching deployment or scaling. This "PoC-to-production gap" represents a critical barrier to realizing AI's full potential in the Swiss financial industry.

This follow-up study, titled *"Bridging the AI PoC–Production Gap – Keys to Deployment Success in Swiss Financial Industry"*, was initiated by SFTI (Swiss FinTech Innovations) and conducted in collaboration with OST (Eastern Switzerland University of Applied Sciences, Institute for Finance & Law), supported by ELCA Advisory. It builds directly on the findings of the previous study *"A Scalable Framework for Implementing Artificial Intelligence in Swiss Financial Institutions"* and addresses the persistent challenge of transitioning AI from pilots into scaled, value-generating deployments.

The study is designed to provide Swiss financial institutions with concrete, evidence-based insights into overcoming organizational, technical, and regulatory barriers. It aims to outline actionable pathways for banks and insurers to unlock AI's business value while ensuring compliance, governance, and operational readiness in a highly regulated and data-sensitive environment.

1.1 Study Objectives

The primary objective of this follow-up study is to identify the key enablers, success factors, and practical strategies necessary for Swiss financial institutions to scale AI initiatives effectively beyond pilot stages. The study complements the previous framework by focusing specifically on the transition from PoC to productive deployment.

Specific objectives include:

1. **Business Case Realization and Expectation Management**
   Provide guidance to ensure that AI projects are grounded in realistic assumptions, deliver measurable ROI, and sustain stakeholder confidence throughout their lifecycle.
2. **Model Governance and Testing Practices**
   Provide guidance on robust model selection, testing, validation, and monitoring approaches that comply with Swiss regulatory requirements and build institutional trust in AI.
3. **Infrastructure and Data Governance**
   Address system dependencies, integration with legacy infrastructures, and confidentiality requirements, while supporting effective data utilization and transparency.
4. **Regulatory Compliance and Risk Management**
   Highlight how to embed compliance-by-design practices, ensuring that AI solutions meet stringent Swiss regulatory standards, including FINMA expectations and data protection requirements.
5. **Scalability and Industrialization**
   Develop strategies, frameworks, and tools that enable institutions to replicate successful deployments across business units efficiently and sustainably.
6. **Practical Application and Knowledge Transfer**
   Deliver actionable insights, playbooks, and best-practice frameworks that institutions can directly apply to current AI initiatives to accelerate their path to production.

1.2 Study Methodology

To ensure both breadth and depth, the study employs a mixed-methods approach, combining quantitative survey data with qualitative expert insights:

1. **Survey Insights**
   A structured survey was distributed among 24 institutions across the Swiss financial services landscape, including two-thirds of SFTI member institutions as well as select additional banks and insurers. The survey collected data on:
   - Current stages of AI use cases (ideation, PoC, deployment, scaling).
   - Levels of business case realization and expectation alignment.
   - Barriers and roadblocks encountered in moving from PoC to production.
   - Success factors and enablers supporting industrialization.
   - Institution-specific practices for model governance, testing, monitoring, infrastructure, and data management.
2. **Analysis of Findings**
   Survey results were analyzed to identify industry-wide trends, such as common bottlenecks, shifts in business case drivers, and the influence of compliance and governance. The data was contextualized through benchmarking against global research on enterprise AI adoption, ensuring that Swiss findings were compared to international patterns.
3. **Expert Interviews**
   In-depth, semi-structured interviews were conducted with selected institutions, capturing both success stories and stalled initiatives. These interviews provided qualitative depth, illustrating how specific approaches, decisions, and governance structures influence the transition from PoC to production.
4. **Collaborative Expertise**
   The study draws on academic rigor from OST's Competence Center for Banking & Finance and practical expertise from industry practitioners and ELCA Advisory. This combination ensures that results are both theoretically sound and directly applicable to real-world contexts.

1.3 Report Structure

This report is organized into the following chapters:

1. **Management Summary**
   A concise overview of the study's objectives, key findings, and recommendations.

1. **Introduction and Study Approach**
   An introduction to the study's purpose, objectives, methodology, and structure.

2. **AI in Swiss Financial Services – Context and Challenges**
   An exploration of the AI adoption landscape, highlighting the barriers and drivers specific to Switzerland, including regulatory requirements, organizational readiness, and global benchmarking.
3. **Survey and Interview Findings – From PoC to Production**
   A detailed presentation of survey results and interview insights, covering use case life cycles, business case expectations, success factors, and roadblocks to deployment.
4. **Keys to Deployment Success – Best Practices and Success Factors**
   Identification and analysis of the critical enablers for scaling AI beyond pilots, including governance, compliance-by-design, data and MLOps infrastructure, and IT–business collaboration.
5. **Framework and Practical Tools for Scaling AI**
   A proposed framework and actionable playbook, including a success checklist and

maturity model to guide institutions in assessing and enhancing their readiness for production deployment.

6. **Conclusion and Recommendations**
   Final reflections on bridging the PoC–production gap and practical recommendations for Swiss financial institutions to achieve sustainable, scalable AI deployment.

This follow-up study offers a robust roadmap for overcoming the barriers to AI deployment in the Swiss financial industry. By combining survey evidence, case-based insights, and academic rigor, it delivers practical strategies that enable institutions to transition from pilots to production and unlock AI's full potential as a transformative force in finance.

2. AI in Swiss Financial Services – Context and Challenges

2.1 Adoption Snapshot in Switzerland

Swiss banks and insurers have accelerated AI exploration, but many initiatives still stall before production. Building on the previous SFTI–OST study, which already showed a concentration in pilot phases (≈43% pilots; 29% deployed; 19% scaling), our new survey and interviews confirm an upstream shift: more ideation and PoCs, fewer deployments and scaled solutions. In 2025, ideation increased markedly while the shares of deployment and scaling tightened, signaling stricter governance gates, scarcer platform capacity, and higher ROI scrutiny before go-live.

What this means: The "pilot purgatory" is less about executive buy-in and more about foundational readiness—data access/quality, privacy and security controls, explainability, and MLOps discipline required for regulated, durable operations.

2.2 Business Drivers in 2025

Efficiency remains the dominant motive, but the framing has evolved from pure cost-out to measurable productivity and quality (time-to-answer, error reduction). Internal user experience has surged—especially employee-facing copilots and retrieval/assist tools—becoming a decisive gate in use-case selection. Cost reduction is still relevant but less decisive; selective growth angles (sales enablement/new business) are emerging, yet remain secondary to operational impact.

2.3 Barriers to Industrialization

The main obstacles to crossing from PoC to production are structural rather than cultural:

- **Data & privacy/security:** fragmented and often uncleansed data, stringent confidentiality, and high assurance requirements slow industrialization.
- **Regulatory clarity & model risk:** explainability, documentation, and ongoing monitoring raise the bar for approval.
- **Technical/MLOps readiness & legacy integration:** reproducible pipelines, monitoring, rollback, lineage, and integration into complex core landscapes are often underbuilt.
- **Expectation management:** PoCs that lack production-ready architecture or realistic benefit baselines falter at gates.

2.4 Regulatory & Supervisory Context (Switzerland and EU)

Switzerland follows a principle-based, technology-neutral stance. The revised Federal Act on Data Protection (FADP) emphasizes transparency, data minimization, and individual rights. FINMA's Guidance 08/2024 articulates supervisory expectations around governance, inventories of AI use, risk classification, explainability, testing/monitoring, human oversight, and outsourcing controls—without creating AI-specific "hard law," but setting binding expectations for supervised institutions through a risk-based lens. In parallel, the EU AI Act's risk-based regime (and its extraterritorial effects) heightens obligations for high-risk use (e.g., AML/KYC, credit), including documentation, conformity and monitoring; Swiss firms serving EU clients should align Swiss governance with EU-grade transparency and oversight.

Implication: Early "compliance-by-design"—model cards, traceability, fairness monitoring, human-in-the-loop, and centralized AI inventories—shortens approval cycles and reduces re-work.

2.5 What This Means for the Study

The Swiss context is one of **high ambition, high constraints**. Portfolios expanded upstream, but production capacity became the bottleneck. This study therefore focuses on **bridging mechanisms**: production-first stage-gates, data/MLOps foundations, governance and model-risk controls, and benefit ownership—so that fewer, better-architected use cases cross the line and deliver durable value in regulated environments.

3. Survey and Interview Findings – From PoC to Production

This chapter presents the empirical findings of the study, drawing on a structured survey of 24 Swiss financial institutions and a series of follow-up interviews with selected practitioners. The objective was to capture both the **breadth of industry-wide patterns** and the **depth of practical experiences** regarding the transition of AI initiatives from proof-of-concept (PoC) stages into productive deployments.

The results highlight not only the **current distribution of AI use cases across the life cycle** but also the **business case outcomes, barriers, and enablers** shaping success. By combining quantitative evidence with qualitative insights, this chapter provides a comprehensive view of the state of AI industrialization in Switzerland and identifies the critical issues that need to be addressed in order to overcome the PoC–production gap.

3.1 Respondent Demographics and Survey Design

The survey conducted in summer 2025 provides a comprehensive snapshot of the Swiss financial services landscape, with participation from a wide range of institutions and roles. In total, **24 institutions** took part in the study, representing a cross-section of the industry and ensuring both breadth and depth in the collected insights.

3.1.1 Participating Institutions

The largest group of respondents were **regional and cantonal banks (10 institutions)**, reflecting the importance of mid-sized players in the Swiss financial ecosystem. These banks are often key adopters of innovative technologies while operating under tight resource constraints, making their perspectives particularly valuable (see Figure 1).

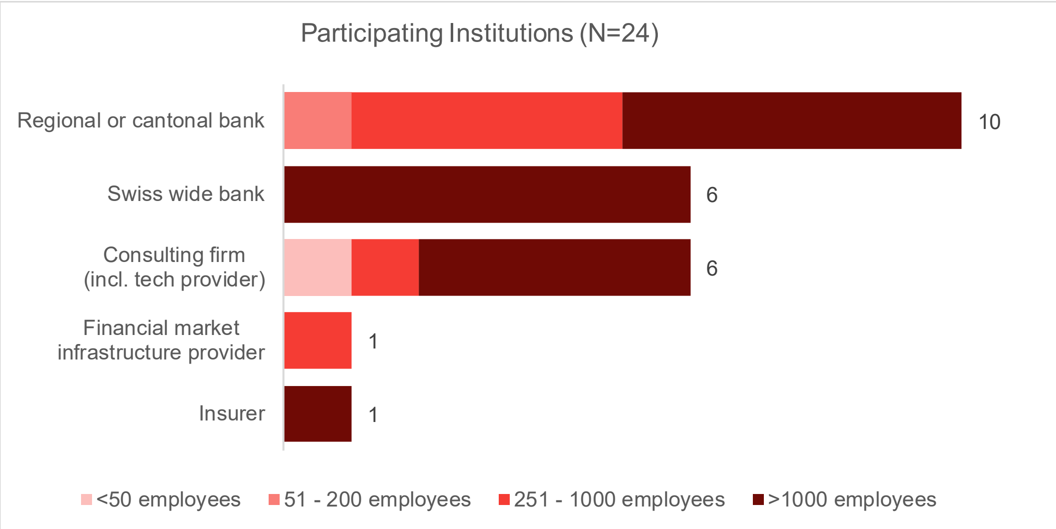

Figure 1: Participating Institutions (N=24)

Swiss-wide banks (6) and **consulting/technology providers (6)** add balance and technological expertise to the sample. Smaller but still relevant contributions came from a **financial market infrastructure provider (1)** and an **insurance company (1)**. While insurance was underrepresented in absolute numbers, the sector's inclusion still brings sectoral diversity to the survey base.

The sample also spans institutions of different sizes, from **small firms with fewer than 50 employees** to **large incumbents with over 1,000 employees**. This cross-size representation provides insights into how maturity, scale, and organizational resources shape AI adoption strategies.

3.1.2 *Responding Roles*

The roles of respondents illustrate a strong representation of both decision-making and implementation functions (see Figure 2). **Executive management (8)** and **business project leads (6)** together account for more than half of the responses, ensuring that strategic and project-level perspectives are well anchored in the data.

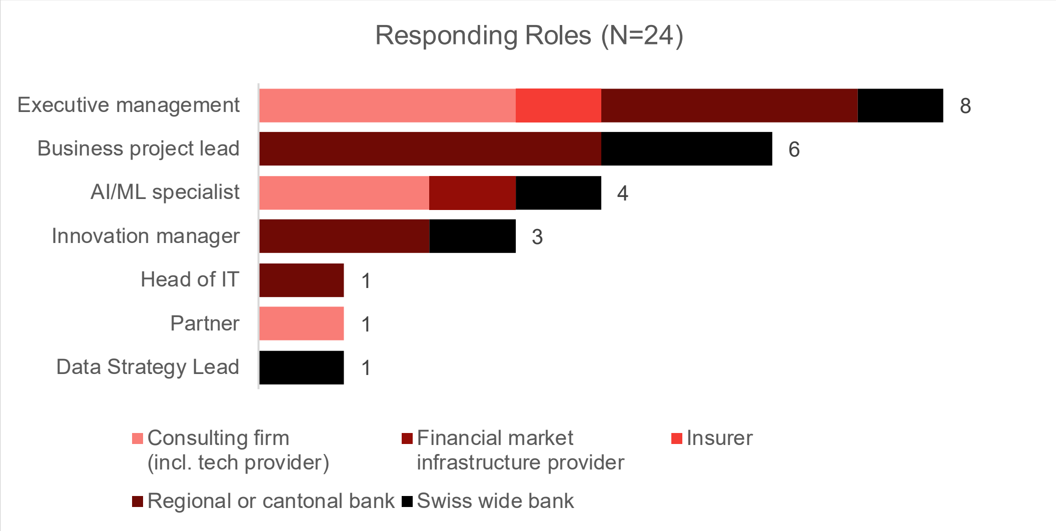

Figure 2: Responding Roles (N=24)

On the technical side, **AI/ML specialists (4)** and **innovation managers (3)** contributed depth in terms of model development, experimentation, and adoption practices. Additional voices included a **Head of IT, a partner, and a data strategy lead**, strengthening the diversity of perspectives.

3.1.3 *Validity and Balance*

The mix of roles demonstrates a **good balance of strategy and execution**: insights capture both the priorities of senior decision-makers and the practical challenges faced by implementers.

The sample is **banking-heavy**, consistent with SFTI's current constituency, but includes representation from insurers and infrastructure providers. This ensures that findings are broadly relevant across financial services, while acknowledging sector-specific nuances.

Finally, the **cross-size coverage**—from small regional banks to large incumbents—enables the study to highlight challenges that arise at different scales, such as resource constraints in smaller institutions versus governance complexity in larger ones.

3.2 *AI Use Case Life Cycle*

The survey results reveal a marked **shift in the distribution of AI initiatives across the use case life cycle**, highlighting both the continued dynamism of experimentation and the persistent bottlenecks in industrialization. Compared with earlier findings from the 2024 study, portfolios have regressed upstream: a larger share of projects remains in **ideation and pilot testing**, while fewer reach deployment or scaling (see Figure 3).

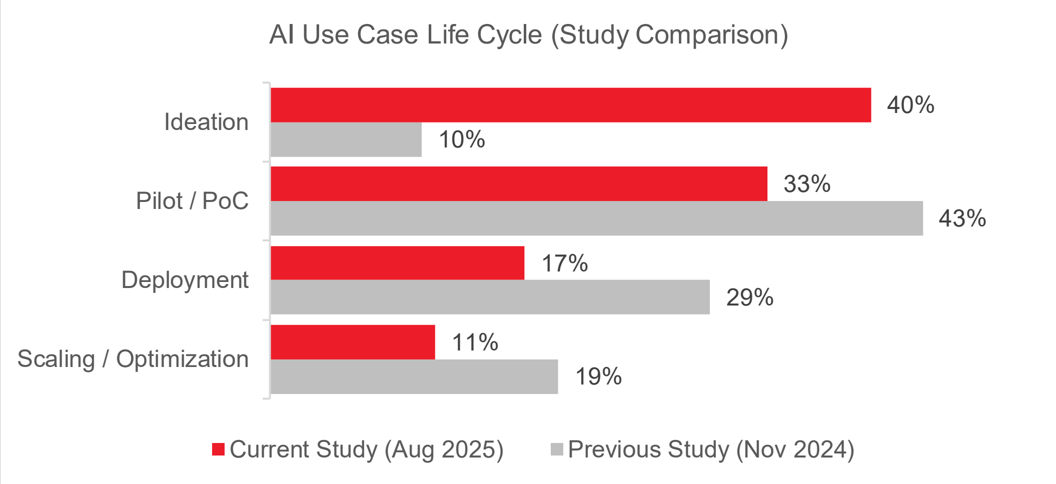

Figure 3: AI Use Case Life Cycle – Distribution of initiatives across ideation, pilot, deployment, and scaling (N=24)

3.2.1 *Current Distribution*

The survey results reveal a clear redistribution of AI initiatives across the lifecycle stages. Compared to the previous year, the balance has shifted markedly toward early-stage activities, with many institutions expanding ideation and feasibility efforts while tightening progression criteria for deployment. This reflects both the wave of new generative AI explorations and a more disciplined approach to production readiness. The following breakdown illustrates how AI initiatives are currently distributed across the key stages of their lifecycle:

- **Ideation and feasibility analysis:** Increased significantly to around **40% of all initiatives** (from ~10% in 2024), reflecting the surge in new explorations, particularly around generative AI.
- **Development and prototyping / pilot testing:** Remains substantial, confirming that institutions continue to invest in proving feasibility and testing limited-scale implementations.
- **Full deployment and integration:** Declined to **17%** of use cases (from 29% in 2024).
- **Scaling and optimization:** Now only **11%** of cases (compared to 19% previously), highlighting how few projects make it beyond initial production.

3.2.2 *Causes of the Upstream Shift*

The pronounced movement of AI initiatives toward earlier stages in the lifecycle is driven by a combination of technological, organizational, and regulatory factors. Institutions are simultaneously broadening their exploratory portfolios and tightening the conditions required for production approval. This shift reflects both a renewed wave of innovation—particularly around generative AI—and an increased focus on risk control, data governance, and return-on-investment validation. The following points summarize the main drivers behind this upstream trend:

1. **Generative AI momentum:** The release of new tools and platforms has triggered a wave of exploratory pilots and ideation exercises.
2. **Stricter governance gates:** Increasing regulatory expectations (e.g., explainability, documentation, model inventories) have raised the bar for production approval.
3. **ROI discipline:** Business case requirements are now tighter, with projects expected to show a clear path-to-value before advancing.
4. **Resource constraints:** Limited capacity of platform, IT, and data teams creates bottlenecks that delay progression to production.
5. **Portfolio hygiene:** Some initiatives were reclassified back to ideation or terminated if requirements hardened or feasibility was disproven.

3.2.3 *Interpretation*

The findings indicate that **Swiss institutions are in a "reset and explore" phase**. Many new projects are initiated, but fewer progress beyond pilots. This does not necessarily imply failure—rather, it reflects a **quality-over-quantity strategy**, where stricter gates ensure that only projects with robust governance, data readiness, and business value advance.

Notably, institutions that succeed in moving projects to production tend to achieve **higher success rates**. In fact, **65% of production cases met or exceeded business case expectations** (see Chapter 3.3), underscoring the survivorship effect of tougher entry barriers.

3.2.4 *Key Implication*

The life cycle distribution confirms the persistence of the **PoC-to-production gap**. While ideation and experimentation are vibrant, scaling remains elusive. Bridging this gap requires not just enthusiasm for AI innovation, but also systematic readiness in data governance, compliance-by-design, and MLOps infrastructure. These elements are decisive in enabling AI use cases to move downstream in the funnel toward productive and scalable impact.

3.3 *Business Case Expectations and Realization*

A central finding of the survey is that, although **fewer projects make it to production**, those that do are significantly more likely to deliver on their business case. **65% of production cases met or exceeded expectations** (see Figure 4). This demonstrates that while the funnel has narrowed, the quality and reliability of outcomes has improved.

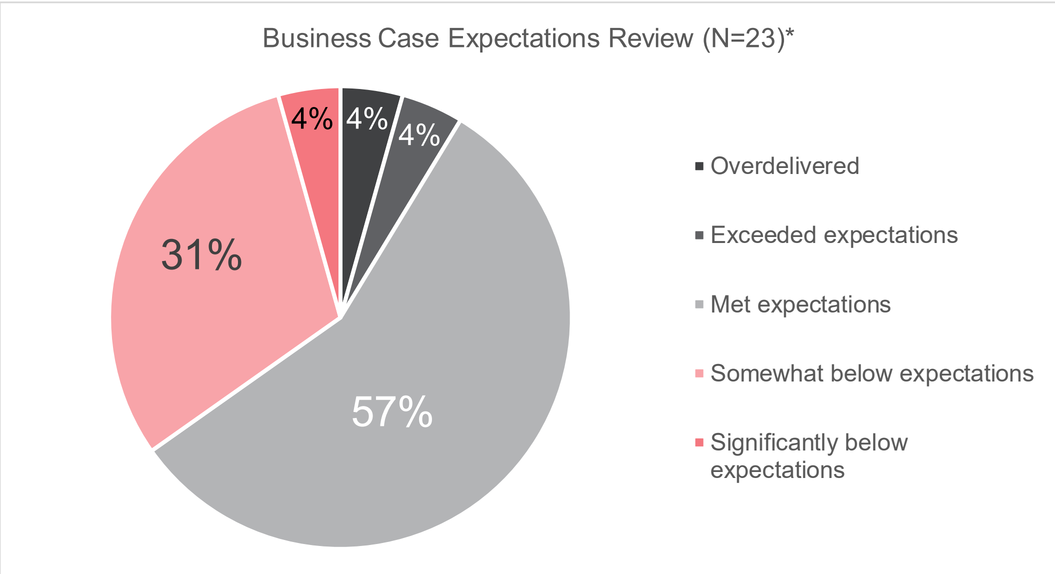

Figure 4: Business Case Realization – Share of production cases meeting/exceeding expectations (N=24)

3.3.1 *Meeting Expectations in Production*

The survey results show that, while fewer AI projects reach production, those that do deliver tangible and measurable value. This confirms a trend toward greater selectivity and maturity in AI implementation, as institutions apply stricter criteria for advancing PoCs to deployment. The following key observations illustrate how production-stage initiatives are performing against their original expectations:

- **Survivorship effect:** Stricter governance, compliance requirements, and ROI gates mean that only the most robust projects progress. As a result, the "hit rate" among deployed cases is higher.

- **Reset & explore dynamic:** The surge in new ideation, particularly in generative AI, creates an apparent regression upstream. Yet earlier cohorts that cleared tougher approval processes now deliver dependable results.
- **Quality over quantity:** The narrowing funnel reflects a deliberate strategy to focus on fewer, but better-architected, initiatives.

3.3.2 Shifts in Business Case Drivers

Beyond delivery performance, the study also highlights a notable evolution in how institutions define and prioritize the business value of AI. Decision criteria have matured from narrow cost-saving motives toward broader efficiency, user experience, and quality objectives in terms of reaching set targets. The points below summarize the main shifts in business case drivers observed across participating institutions (see Figure 5):

- **Operational efficiency** is now universal (100%). The focus has shifted from cost-cutting to measurable productivity gains such as process throughput, error reduction, and time-to-answer.
- **Internal user experience** has risen dramatically (71%), reflecting the adoption of employee-facing copilots, workflow assistants, and retrieval tools. Making work easier for staff has become a decisive factor for deployment.
- **Cost reduction** has declined in importance (42%), suggesting that pure savings cases are less persuasive compared to efficiency and quality-driven metrics.
- **Risk and compliance** remain important (29%), but are increasingly embedded as a baseline requirement rather than a primary driver.
- **Growth use cases** (sales enablement, new business opportunities) are cautiously emerging (≈17%), but remain secondary to efficiency and user enablement.

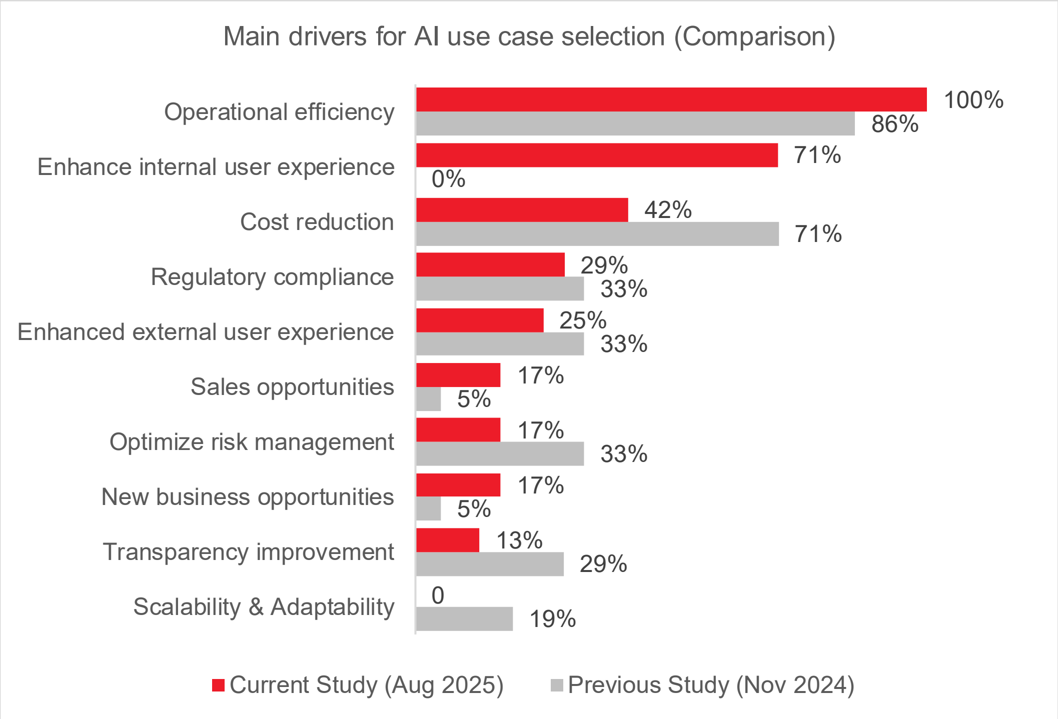

Figure 5: Business Case Drivers – Shift from cost reduction toward productivity and user experience

3.3.3 *Interview Insights*

Qualitative interviews confirmed that **business owners demand tangible, measurable outcomes** rather than abstract promises of efficiency. Successful production cases are those tied to **clear KPIs**, such as minutes saved per process, error rates reduced, or throughput gains. Importantly, projects with an **accountable business owner** overseeing adoption and benefit tracking were consistently reported as more successful.

3.3.4 *Interpretation*

The findings show that **expectation realism and ROI discipline have become central to AI adoption**. Institutions that enforce production-first thinking and demand clear "path-to-prod" artifacts at the PoC stage—such as reference architectures, observability, and benefit baselines—are better able to secure durable results.

3.3.5 *Key Implication*

The Swiss financial industry is transitioning from a "cost-saving" to a **productivity and user-value paradigm** in AI. This evolution reflects maturity: efficiency gains are still paramount, but now framed in terms of usability, adoption, and measurable ROI. To sustain credibility and investment, institutions must embed **business case discipline** throughout the lifecycle, ensuring that expectations are realistic and outcomes are tracked.

3.4 *Roadblocks to Scaling AI Beyond PoC*

While enthusiasm for AI remains high, survey results underscore that the **biggest obstacles to scaling AI are structural, not cultural**. Respondents consistently highlighted challenges linked to data readiness and governance, compliance requirements, and technical readiness, while organizational buy-in and user acceptance were rated as less critical (see Figure 6).

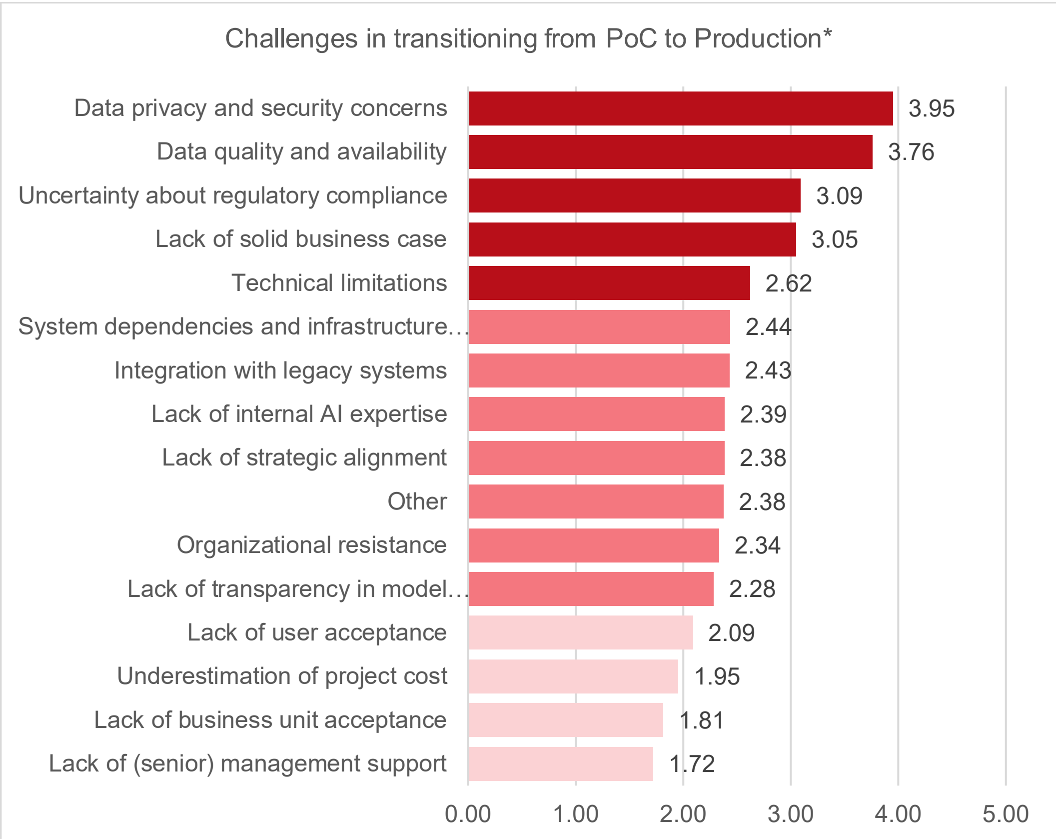

Figure 6: Roadblocks to Scaling AI – Ranked significance of barriers (N=24)
* Rated on a scale from 1 = Low to 5 = High, and 0 = not relevant

3.4.1 *Top-Tier Barriers*

The most significant obstacles to scaling AI beyond the proof-of-concept phase are structural rather than cultural, which also indicates an increased overall “acceptance” of AI compared to the previous study. Respondents consistently identified challenges rooted in data, compliance, and technical readiness as the primary factors preventing industrialization. The following points outline the top-tier barriers that most strongly limit progress toward production:

1. **Data Privacy & Security**
   - Strict controls and confidentiality requirements slow adoption.
   - Legal and reputational risks compel institutions to enforce conservative policies around data use.
   - Interviewees noted that gaining approval for sensitive data usage often extends timelines significantly.
2. **Data Quality & Availability**
   - Fragmented, inconsistent, or siloed datasets undermine reliability.
   - Lack of harmonized data infrastructure means that many PoCs cannot be replicated or scaled across functions.
   - As one executive commented, “We don’t fail because of algorithms—we fail because of data.”
3. **Regulatory Uncertainty**
   - Unclear expectations around explainability, bias mitigation, and model risk approval create hesitation.
   - Respondents reported that uncertainty about how regulators will interpret new AI applications often leads to projects being stalled.
4. **Weak Business Cases**
   - ROI often remains insufficiently defined at the pilot stage.
   - Users and sponsors question the added value when benefits are not quantified or tied to KPIs.
5. **Technical Limitations**
   - Missing MLOps infrastructure, reproducibility, and monitoring capabilities prevent progression to production.
   - Legacy systems and integration challenges add further complexity, particularly in mid-sized banks.

3.4.2 *Mid-Tier Barriers*

Beyond the major structural hurdles, several mid-level barriers also constrain AI deployment. These issues typically do not block projects entirely but create friction that slows execution or increases costs. The factors below highlight additional challenges that, while secondary, still require attention to ensure smooth scaling:

- **System Dependencies & Legacy Integration**: Dependence on core banking platforms complicates rollout and increases costs.
- **AI Expertise Gaps**: Lack of in-house specialists slows progress; outsourcing is not always aligned with governance expectations.
- **Cost Transparency**: Institutions cited difficulties in fully understanding the long-term operational costs of AI solutions.

3.4.3 *Lower-Tier Barriers*

At the lower end of the spectrum, respondents rated organizational and cultural issues as relatively minor impediments. Senior management support for AI is generally strong, and user acceptance is improving as familiarity increases. Nevertheless, these softer aspects continue to influence adoption speed and should not be overlooked. The following points summarize the barriers considered least critical by participating institutions:

- **Management Support & Sponsorship**: Rated relatively low as a blocker. Senior leaders are generally supportive of AI initiatives.
- **User Acceptance**: While relevant for adoption, it was not seen as the decisive bottleneck compared to compliance or data challenges.

3.4.4 *Interpretation*

The findings demonstrate that **AI in Swiss financial institutions is not blocked by cultural resistance, but by foundational readiness** – AI seems to be broadly "accepted". Scaling fails less from a lack of vision than from difficulties in proving compliance, integrating into legacy infrastructures, and ensuring data quality.

3.4.5 *Key Implication*

To bridge the PoC–production gap, institutions must focus on **fixing the fundamentals first**:

- Establishing robust data governance and harmonization and ensure golden source of cleansed data.
- Clarifying regulatory expectations and embedding compliance early ("compliance-by-design").
- Building or acquiring scalable MLOps capabilities.
- Enforcing realistic and measurable business cases before advancing pilots.

Without addressing these foundations, organizational readiness and change management—though important—cannot by themselves ensure successful industrialization.

3.5 *Success Factors for Deployment*

The survey findings and interview insights reveal that **successful AI deployments in Swiss financial institutions are consistently enabled by a small set of decisive factors**. These success factors are rooted in governance, collaboration, and infrastructure and data readiness, and they highlight that scaling AI requires both technical capabilities and organizational alignment (see Figure 7).

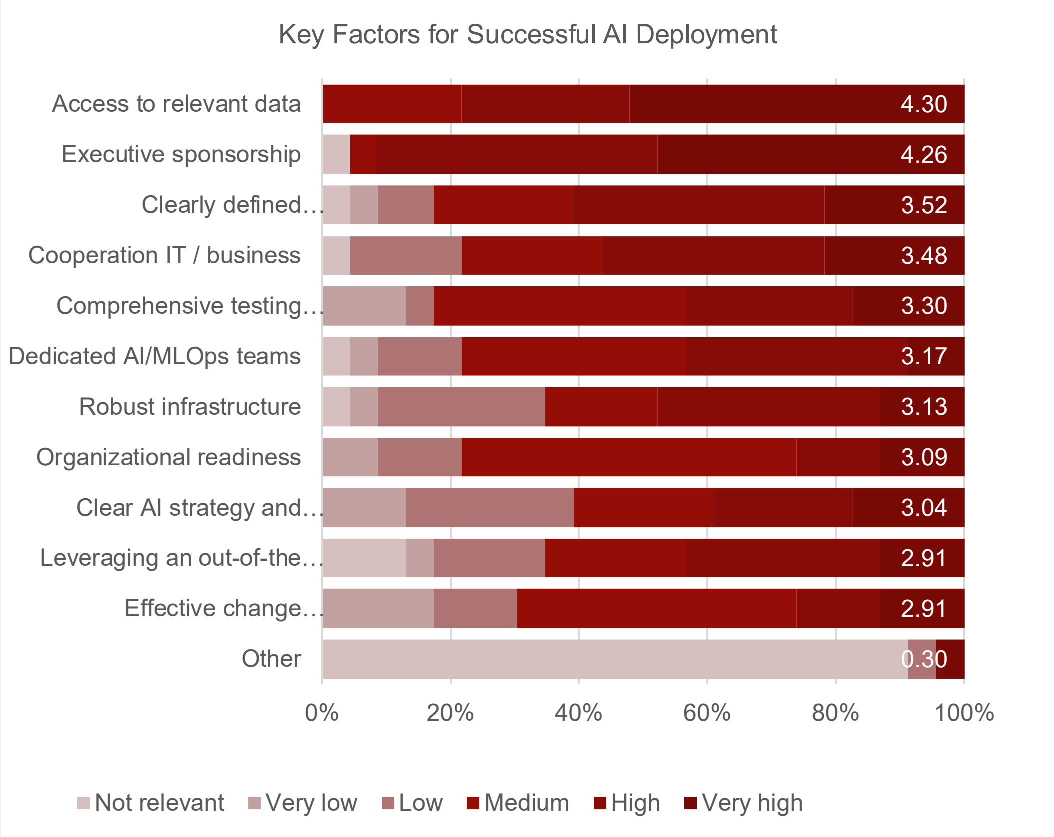

Figure 7: Success Factors for Deployment – Ranked enablers and accelerators (N=24)

3.5.1 *Top Drivers of Success*

Institutions that successfully bridge the PoC–production gap demonstrate a clear pattern of organizational and operational strengths. Their achievements are not based on experimental innovation alone but on well-established foundations in data quality, leadership commitment, and governance maturity. The following key drivers were identified as the most decisive factors contributing to sustainable AI deployment success:

1. **Access to Reliable Data**
   - Institutions that achieve seamless and secure access to relevant datasets report faster transitions from PoC to production.
   - Data readiness (quality, harmonization, accessibility) is the single most cited determinant of deployment success.
2. **Executive Sponsorship**
   - Senior leadership support ensures prioritization, funding, and speed in decision-making.
   - Sponsorship is most effective when paired with expert teams that can translate strategic vision into technical execution.
3. **Governance Frameworks**
   - Clear governance rules, including AI inventories, risk classification, and explainability standards, create confidence among compliance teams and regulators.
   - Institutions with robust governance report shorter approval cycles and fewer post-deployment issues.
4. **IT–Business Cooperation**
   - Shared ownership between technical and business units helps bridge the gap between prototype and product.

  - Interviewees stressed that alignment of IT and business goals is critical to move projects through deployment gates.

5. **Testing & Monitoring Frameworks**
  - Continuous validation (bias detection, drift monitoring, stress testing) increases trust in model robustness.
  - Respondents emphasized that structured testing and monitoring not only reduce risks but also accelerate acceptance and confidence.

3.5.2 *Strong Enablers*

In addition to the top-tier success drivers, several organizational and technical enablers further accelerate AI deployment. While not always decisive on their own, these factors strengthen institutional readiness and create the operational stability needed for scaling. The points below summarize the most frequently mentioned enablers supporting successful transitions into production:

- **Dedicated AI/MLOps Teams:** Institutions with specialized teams for deployment and monitoring show higher success rates.
- **Robust Infrastructure:** Institutions investing in scalable architectures, within Swiss legal frameworks, progress more quickly.
- **Clear AI Strategy & Roadmap:** Alignment with business priorities and long-term planning creates clarity and focus.

3.5.3 *Pragmatic Accelerators*

Finally, respondents identified a set of pragmatic accelerators that help shorten time-to-value and improve adoption in practice. These measures typically complement structural success factors by addressing implementation efficiency and user engagement. The following examples illustrate practical levers that institutions use to translate AI potential into tangible business outcomes more quickly:

- **Out-of-the-Box Solutions:** Pre-built applications with lower configuration effort shorten time-to-value.
- **Change Management:** Proactive engagement with users builds trust and adoption, mitigating the risk of resistance at rollout.

3.5.4 *Interview Insights*

Several interviewees highlighted that **early compliance integration** (“compliance-by-design”) was pivotal. By involving compliance teams at the start, institutions avoided delays at the production gate. Another recurring theme was **benefit ownership**: assigning accountable business owners to monitor adoption and track KPIs was seen as a major differentiator for success.

3.5.5 *Interpretation*

These findings underscore that **winning deployments are not driven by experimentation alone but by strong execution capabilities**. Institutions that combine **data readiness, leadership sponsorship, governance clarity, and IT–business collaboration** are best positioned to move beyond pilots. Testing and monitoring frameworks, together with specialized teams, ensure that AI initiatives remain robust and scalable after go-live.

3.5.6 *Key Implication*

To scale AI successfully, Swiss financial institutions must **treat deployment as a structured, governed process**, not a natural extension of experimentation. Effective deployment depends on:

- Reliable data foundations.
- Senior leadership sponsorship with accountability.

- Early and robust governance frameworks.
- Continuous monitoring and benefit ownership.

Together, these success factors form the backbone of sustainable AI industrialization.

3.6 *Decision Criteria for AI Application Selection*

When evaluating AI applications and vendor solutions, Swiss financial institutions place a clear emphasis on **security, compliance, and sovereignty**. These criteria consistently outweigh factors such as openness, customizability, or even cost considerations. The findings reflect the regulated and risk-sensitive environment in which banks and insurers operate (see Figure 8).

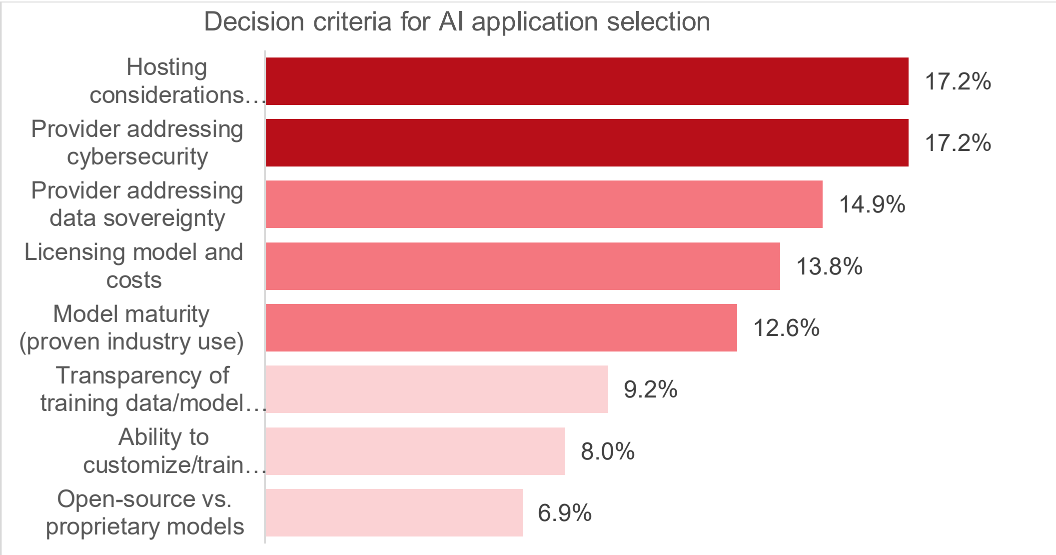

Figure 8: Decision Criteria for AI Application Selection – Ranked factors (N=24)

3.6.1 *Most Decisive Criteria*

When selecting AI applications and vendors, Swiss financial institutions apply rigorous evaluation criteria focused on security, compliance, and long-term scalability. The decision process reflects the sector's high sensitivity to data protection and regulatory alignment, with hosting models, cybersecurity, and data sovereignty emerging as the most influential factors. The following key criteria represent the primary considerations driving AI application and vendor selection across institutions:

1. **Hosting & Cybersecurity (17% each)**
   - Institutions prioritize where AI runs (on-premise vs. cloud) and how data security is enforced.
   - Concerns include encryption standards, access controls, and compliance with ISO and FINMA requirements.
   - Cybersecurity is treated as a prerequisite rather than a negotiable feature.
2. **Data Sovereignty (15%)**
   - Control over data location and jurisdiction remains a decisive factor.
   - Respondents stressed the importance of ensuring that sensitive client data stays within Switzerland or under Swiss/EU legal protection.
3. **Licensing & Costs (14%)**
   - While cost efficiency matters, the emphasis is on contractual clarity and scalability.
   - Institutions favor licensing models that align with ROI expectations and allow for future scaling.

4. **Model Maturity (13%)**
    - Banks look for assurance that selected models have been tested in real-world financial contexts.
    - Proven adoption in peer institutions or reference cases is an important gate for vendor selection.

3.6.2 *Secondary Criteria*

Beyond the core decision drivers, institutions also evaluate a range of secondary factors that influence vendor selection and deployment strategies. While these aspects rank lower in quantitative importance, they often become decisive in later stages of due diligence—particularly for regulatory audits and operational resilience. The following points summarize the secondary criteria that complement the main selection framework:

- Transparency & Customizability (9% and 8%)
    - While lower on the ranking, institutions acknowledged that explainability and configuration flexibility are essential for regulatory audits and long-term resilience.
    - Interviews highlighted that these aspects often become deal-breakers at later stages of evaluation.
- Open-Source vs. Proprietary (7%)
    - Most institutions accept vendor lock-in for stability and compliance assurances.
    - Open-source was rarely prioritized, but some interviewees saw it as a potential hedge against concentration risk.

3.6.3 *Interview Insights*

Participants noted that **risk and compliance requirements dominate selection decisions**. One project lead explained: *“Even if a tool promises high efficiency, if it cannot demonstrate regulatory alignment on hosting, security, and data sovereignty, it simply won’t pass procurement.”* Another respondent emphasized that cost is rarely decisive alone: *“Cost discussions only happen after compliance has been assured.”*

3.6.4 *Interpretation*

The results underline a dual reality: while **business drivers (efficiency, user experience)** define which use cases are prioritized (see Chapter 3.3), **technical and compliance safeguards** determine which vendors and tools are actually chosen. The dominance of hosting, security, and sovereignty demonstrates that in Switzerland, **AI industrialization is gated by trust and risk management**.

3.6.5 *Key Implication*

For Swiss financial institutions, application selection is a **compliance-first process**. To accelerate adoption, vendors must:

- Provide transparent evidence of compliance with hosting, encryption, and localization requirements.
- Demonstrate proven maturity and successful use cases in the financial sector.
- Offer licensing models that align with the long-term ROI of the AI business case.

Institutions should not overlook transparency and customizability: while ranked lower, these factors will become increasingly important for explainability and resilience as regulatory expectations continue to evolve.

3.7 *AI Model Card Expectations*

The survey explored which aspects of **AI model cards** are considered most important for Swiss financial institutions. Model cards serve as structured documentation of AI systems, detailing their purpose, design, limitations, and risks. In line with Switzerland's risk-sensitive and compliance-driven environment, institutions placed the highest value on elements directly linked to **trust, explainability, and regulatory alignment** (see Figure 9).

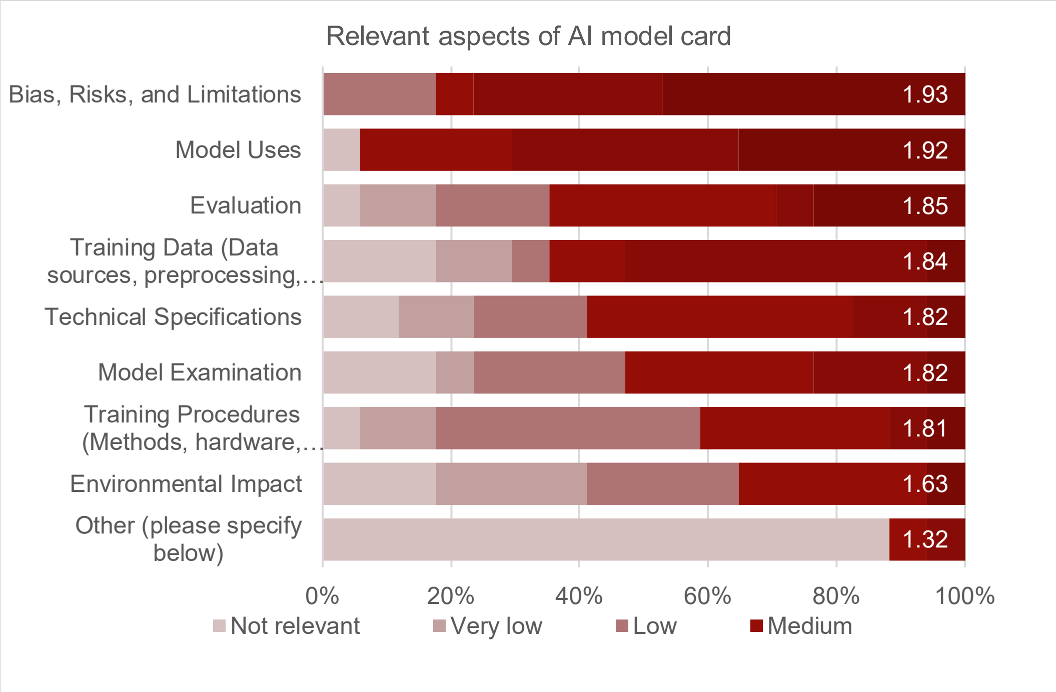

Figure 9: AI Model Card Aspects – Ranking of importance (N=24)

3.7.1 *Top-Rated Aspects*

Survey responses show clear alignment among institutions regarding what constitutes high-quality model documentation. The most valued elements of an AI model card are those that ensure transparency, reliability, and regulatory compliance. Respondents emphasized that model cards must go beyond technical summaries to include clear statements of risk, validation, and intended use. The following aspects were rated as the most important for trustworthy and regulator-ready AI documentation:

1. **Bias, Risks & Limitations (1.93)**
   - The most critical factor is transparency on where models may fail or behave unpredictably.
   - Institutions expect explicit disclosure of risks, known limitations, and bias concerns.
   - This aligns with FINMA's emphasis on fairness and explainability in supervisory expectations.
2. **Intended Use (1.92)**
   - Clear articulation of the scope and boundaries of model application is essential.
   - Respondents stressed the importance of avoiding "function creep," where models are applied in unintended contexts without validation.
3. **Evaluation and Validation (1.85)**
   - Institutions value rigorous documentation of how models were tested, validated, and benchmarked.
   - Regular re-validation is expected, particularly for high-risk use cases such as AML, fraud detection, or credit scoring.

4. **Training Data Transparency (1.84)**
    - Financial institutions require clarity on the origin, quality, and representativeness of training data.
    - Transparency in data lineage is seen as a prerequisite for regulatory acceptance.

3.7.2 Mid-Tier Priorities

In addition to the top-rated factors, participants identified a group of mid-tier priorities that, while less visible to non-technical stakeholders, are essential for robust model governance. These elements support technical validation, auditability, and operational control, forming the foundation of responsible model management. The points below summarize these secondary priorities:

- **Technical Specifications (1.82), Model Examination (1.82), and Training Procedures (1.81):** These aspects matter particularly for IT and validation teams, but were slightly less emphasized by business and compliance stakeholders.
- They nonetheless provide critical evidence during audits and reviews.

3.7.3 Lower Emphasis

Certain aspects of AI model documentation received comparatively limited attention among respondents. While factors such as environmental impact and additional metadata are recognized as relevant in broader ESG and transparency discussions, they are not yet central to current regulatory or operational priorities. The following findings summarize the areas of lowest emphasis within the Swiss financial sector:

- **Environmental Impact (1.63):**
  While relevant in a broader ESG context, sustainability aspects of AI were not prioritized in the Swiss financial sector at this stage.
- **Other (1.32):**
  Respondents did not identify additional categories beyond those provided, suggesting that the listed dimensions adequately reflect institutional needs.

3.7.4 Interview Insights

Interviews confirmed that **risk and compliance considerations dominate model documentation**. One compliance officer remarked: *"If we cannot explain where the model may fail and how it was tested, it will never pass regulatory scrutiny."* Another participant stressed that **training data transparency** is increasingly critical, particularly for generative AI, where concerns about provenance and copyright have emerged.

3.7.5 Interpretation

The strong emphasis on **risks, limitations, intended use, and validation** illustrates that Swiss institutions view model cards as a regulatory safeguard as much as an operational tool. This focus aligns closely with FINMA's expectations (Guidance 08/2024) and with international developments such as the EU AI Act.

3.7.6 Key Implication

For model cards to enable scaling, institutions should:

- Standardize their use across all AI applications, not just high-risk ones.
- Ensure that documentation highlights risks, intended use, and validation evidence in accessible, non-technical language.
- Treat training data transparency as a baseline requirement for compliance and trust.

By embedding **model cards as a compliance-by-design tool**, Swiss financial institutions can accelerate the approval and deployment of AI systems, while mitigating risks and strengthening stakeholder confidence.

3.8 Synthesis: From Findings to Deployment Challenges

The combined survey and interview findings present a clear picture of the current state of AI adoption in Swiss financial institutions: **AI portfolios are expanding in ideation and pilot testing, but fewer projects reach production and even fewer scale**. Yet once in production, projects deliver strong outcomes, with **65% of cases meeting or exceeding business case expectations**. This creates a dual reality of both **promise and constraint**.

3.8.1 Key Themes Across Findings

The synthesis of survey and interview insights reveals a coherent picture of how Swiss financial institutions are navigating the transition from AI experimentation to production. Despite a widening gap between ideation and scaling, the overall direction reflects growing maturity, stronger governance, and a clearer understanding of value realization. The following key themes summarize the most important patterns observed across all categories of findings:

1. **The Upstream Shift**
   - AI portfolios have regressed upstream, with ideation and pilots dominating while deployment and scaling shares have declined.
   - This reflects both renewed enthusiasm for experimentation (especially with generative AI) and stricter governance gates.
2. **Business Case Discipline**
   - Institutions now emphasize measurable productivity and user experience gains over pure cost savings.
   - Business case credibility is reinforced by accountable ownership, clear KPIs, and conservative ROI baselines.
3. **Foundational Barriers**
   - Data privacy, security, and quality remain the top bottlenecks, followed by regulatory uncertainty and lack of MLOps maturity.
   - Technical readiness and integration with legacy infrastructures limit progression from PoC to production more than cultural resistance or management support.
4. **Success Factors**
   - Winning deployments are consistently enabled by strong data foundations, executive sponsorship, robust governance, IT–business cooperation, and testing/monitoring frameworks.
   - Early compliance integration and benefit ownership emerge as decisive accelerators.
5. **Selection and Governance Criteria**
   - Hosting, cybersecurity, and data sovereignty dominate vendor choices, underscoring the sector's risk-sensitive nature.
   - Model cards are valued most for their ability to disclose risks, limitations, intended use, and validation methods, reinforcing compliance-by-design.

3.8.2 Interpretation

The synthesis of findings highlights a paradox: **Swiss financial institutions are both cautious and successful.** Cautious, because many projects stall before production due to heightened requirements around compliance, data, and technical readiness. Successful, because those projects that clear these hurdles tend to deliver reliably on expectations.

This duality demonstrates that the **PoC-to-production gap is not a matter of willingness but of readiness**. Institutions are willing to invest, but without mature data governance, scalable MLOps, and embedded compliance, projects cannot advance. Conversely, when these fundamentals are in place, AI projects are not only deployed but also deliver durable business value.

3.8.3 *Key Implication*

The challenge for Swiss financial institutions is therefore to **build the foundations that allow more projects to cross the gap**:

- Treat compliance and risk as design principles, not as afterthoughts.
- Prioritize investments in data, infrastructure, and monitoring capabilities.
- Introduce production-first stage gates, ensuring that PoCs are architected for scale from the outset.
- Tie business cases to measurable KPIs, with accountable ownership and transparent benefit tracking.

Only by addressing these systemic barriers can the industry move from an abundance of ideas to a sustainable pipeline of **trusted, value-generating AI solutions at scale**.

3.9 *Other Notable Observations*

While the preceding chapters have focused on the main analytical pillars of this study — including the AI use-case lifecycle, business-case realization, selection criteria, success factors, and scaling challenges — the survey and interview data also revealed a number of additional, noteworthy patterns.

These findings do **not alter the frameworks, interpretations, or recommendations presented in previous and following sections**, nor do they suggest revisions to the success equation that will be introduced later in the study. Instead, they provide **complementary insights** that deepen the understanding of the Swiss financial sector's evolving AI landscape. Each observation captures a distinctive tendency or emerging development which, while not central to the study's analytical model, is nevertheless valuable to document and may inspire future inquiry or targeted initiatives.

1. **Uneven Institutional Readiness Across Bank Sizes**

   Survey results reveal clear asymmetries in AI maturity between large, nationally operating banks and smaller regional or cantonal institutions.

   Whereas larger players report fewer new pilots but higher integration depth and more robust governance, smaller banks show abundant ideation activity yet struggle to move beyond pilot stages. This suggests the emergence of a **two-speed AI economy**: smaller institutions continue to explore broadly, while larger ones consolidate around a few industrialized, high-value use cases. The divergence points to differing resource levels, compliance capacities, and data readiness — not differing ambitions.

2. **The Rise of Generative AI Exploration**

   A recurring theme in qualitative responses is the proliferation of generative AI pilots — including LLM-based assistants, summarization tools, and document-processing models.

   These initiatives have rapidly expanded ideation pipelines but often sit **outside formal MLOps or data-governance frameworks**, creating friction once regulatory or security review begins. The result is a "reset and explore" dynamic: enthusiasm for generative applications drives experimentation, yet institutionalization lags behind. The trend signals a **shift from analytical to generative AI** as the primary innovation domain, accompanied by the need to extend governance and control mechanisms to new model types.

3. **Compliance and Risk Functions as Emerging Enablers**

   Contrary to earlier perceptions of compliance as a blocker, several respondents describe a more **collaborative and enabling role** of risk and compliance teams.

   Institutions that involve compliance functions early in the AI lifecycle — embedding "policy-by-design" approaches — report **shorter approval times** and smoother transitions to production. This cultural shift indicates that regulatory engagement is maturing and becoming a competitive advantage rather than a constraint.

4. **Model Documentation and the Limited Standardization of AI Model Cards**

   Although most respondents maintain some form of model documentation, only a minority use a structured "model card" approach.

   Where formalized documentation exists, it focuses primarily on model risks, limitations, and validation results, with less emphasis on explainability, reusability, or sustainability criteria. The absence of **industry-wide templates or minimum documentation standards** reduces comparability and complicates cross-institutional audits. Harmonization in this area could significantly increase transparency and regulatory confidence in the future.

5. **Low Consideration of Environmental and Sustainability Aspects**

   Among the various model-card dimensions evaluated, **environmental impact scored lowest** (1.63 on average).

   This suggests that sustainability has not yet become an established element of AI governance within the Swiss financial industry. While understandable given current regulatory priorities, this represents a potential blind spot as **ESG-aligned AI policies** and energy-efficiency metrics begin to enter European and international supervisory discourse.

6. **Vendor Dependency and Limited Openness**

   Decisions on AI platforms and vendors continue to be dominated by **hosting location, cybersecurity, and data sovereignty** considerations.

   While these priorities are legitimate given the sector's sensitivity to client data and regulatory scrutiny, they result in a bias toward **closed, proprietary ecosystems**. Transparency, customizability, and open-source options rank markedly lower. The data thus point to a growing **vendor-lock-in risk** that may limit flexibility and long-term resilience of the Swiss AI ecosystem.

7. **The "Shadow AI" Phenomenon**

   A subtle but telling theme is the spread of unofficial or "shadow" AI usage within organizations — employees experimenting with public LLM tools or copilots for daily tasks.

   While these activities can enhance productivity, they also **circumvent governance, security, and data-protection standards**. This phenomenon reflects the tension between innovation agility and regulatory discipline and underlines the importance of establishing internal policies that enable responsible experimentation without exposure to compliance risk.

8. **Lack of Quantitative Impact Tracking**

   Even among institutions with AI in production, **systematic post-implementation tracking of business benefits** remains rare.

   Most respondents rely on anecdotal evidence or subjective satisfaction feedback rather than defined KPI frameworks. This limits the ability to measure true value realization and weakens the feedback loop into portfolio steering. Embedding benefit measurement into

MLOps and governance routines could significantly strengthen accountability and investment confidence.

9. **The Emergence of Cross-Functional AI Product Squads**

   Several successful institutions have transitioned from project-based to **product-oriented AI delivery structures**, combining IT, data science, and business expertise in permanent “AI squads.”

   These cross-functional teams treat AI solutions as evolving products with defined owners, budgets, and lifecycle management. The model accelerates deployment and aligns AI initiatives more closely with business value streams — a notable **organizational innovation** within the Swiss financial sector.

10. **The Hidden Cost of Data Quality Debt**

    Respondents consistently cite **data quality and accessibility as primary obstacles**, but open-ended feedback reveals a more structural root cause: data ownership fragmentation, legacy architectures, and insufficient metadata management. These factors create persistent “data debt” that hinders reliable model training and reuse.

    The challenge is therefore not only technical but **organizational and cultural**, requiring clear data stewardship roles and long-term investment in foundational data governance — the very bedrock upon which scalable AI depends.

Collectively, these observations depict a **dynamic but unevenly maturing AI landscape** in Swiss financial services.

They underscore that beyond the core challenges of scaling and governance, a second layer of organizational, cultural, and structural issues continues to shape AI’s trajectory.

While none of these elements necessitate revisions to the success frameworks or recommendations presented throughout the study, they **enrich the overall picture** and offer valuable signals for where institutions — and the industry at large — may focus their next steps in advancing AI maturity.

4. Keys to Deployment Success – Best Practices and Success Factors

Building on the survey and interview findings presented in Chapter 3, this chapter distills the critical factors that determine whether AI initiatives in Swiss financial institutions successfully progress from proof-of-concept (PoC) to productive deployment. While barriers such as data quality, compliance uncertainty, and technical readiness often slow progress, the evidence also shows a consistent set of **success factors and best practices** that enable institutions to bridge the gap.

The analysis demonstrates that successful deployments are not the result of isolated pilots but of **systematic readiness** across data, governance, infrastructure, and organizational alignment. Institutions that combine strong leadership sponsorship with robust data foundations, clear compliance-by-design principles, and close IT–business collaboration are far more likely to achieve durable and scalable AI impact.

This chapter organizes these enablers into **practical success factors, best practices, and guiding principles**, providing a roadmap for institutions seeking to industrialize their AI initiatives. The following sections synthesize insights from survey evidence, interview narratives, and international benchmarks, highlighting actionable strategies that can be applied across the Swiss financial sector.

4.1 Strategic Leadership and Business Alignment

The survey and interview findings consistently demonstrate that **strategic leadership and business alignment are decisive prerequisites** for moving AI projects beyond the PoC stage. Institutions that approach AI adoption as a strategic initiative—rather than as isolated technology experiments—report higher success rates in reaching production and delivering measurable business value. Three dimensions stand out as particularly critical: executive sponsorship, disciplined business case management, and realistic expectation setting.

4.1.1 Executive Sponsorship

Senior leadership support was ranked among the most important success factors across survey respondents. Effective sponsorship provides:

- **Budget and resources:** Without senior-level backing, AI initiatives often remain underfunded and unable to progress beyond pilots.
- **Prioritization:** Sponsorship ensures that projects are recognized as strategically relevant and not treated as side experiments.
- **Pace and decisiveness:** With leadership support, approval cycles are shorter, and resistance from compliance, IT, or operations teams can be addressed swiftly.

Interviewees emphasized that sponsorship is most effective when it goes beyond symbolic endorsement. Projects where **executives actively monitored progress, engaged with compliance teams, and pushed for organizational alignment** were consistently more likely to succeed in deployment.

4.1.2 Business Case Discipline

Clear and credible business cases emerged as a central determinant of production success. Respondents stressed that AI projects must:

- Be tied to **concrete KPIs** such as process throughput, error reduction, time-to-complete, or cost avoidance.
- Have **accountable business owners** responsible for benefit realization, rather than leaving accountability solely with IT or innovation teams.
- Be grounded in **conservative ROI assumptions**, avoiding exaggerated promises that cannot be sustained during production.

Institutions that enforced such discipline reported higher confidence from stakeholders and were better able to secure resources for scaling. The shift observed in the survey—from cost reduction to **productivity and user experience gains** as primary drivers (see Chapter 3.3)—further underlines the importance of aligning AI initiatives with business priorities that resonate with users and management alike.

4.1.3 *Expectation Management*

The findings also reveal the importance of **managing expectations across the lifecycle**. Overly ambitious PoCs—designed without a clear path to production—tended to disappoint and stall. By contrast, institutions that framed PoCs as **experiments within a structured funnel**, including the possibility of early termination, were able to maintain credibility even when specific initiatives did not advance.

Interviewees highlighted that **accepting PoC failures as part of portfolio learning** was a mark of maturity. One respondent noted: *"We don't measure success by every PoC going live. Success is when we know early which ideas are worth scaling."* This mindset reduces the risk of sunk costs and supports more selective, quality-focused progression to deployment.

4.1.4 *Interpretation*

Together, these insights underline that **strategic leadership and business alignment provide the foundation for AI industrialization**. Without executive sponsorship, robust business cases, and realistic expectations, even technically sound projects are unlikely to cross the PoC–production gap. Conversely, when leadership prioritizes AI, ties it to measurable business outcomes, and manages expectations, institutions create the organizational momentum needed for scaling.

4.1.5 *Key Implication*

Swiss financial institutions should treat AI deployment as a **strategic, business-driven initiative**. This requires:

- Visible and active sponsorship from senior leadership.
- Business cases tied to KPIs and benefit ownership.
- Realistic ROI expectations and a willingness to terminate weak PoCs early.

In doing so, institutions can align AI initiatives with their broader strategic objectives, ensuring not only deployment success but also sustained value creation.

4.2 *Data Foundations and MLOps Infrastructure*

Across survey responses and interviews, **data readiness and MLOps capabilities** emerged as the strongest technical determinants of whether AI initiatives successfully transition from PoC to production. While strategy and leadership define direction, it is the quality of data and the robustness of model operations that determine executional success. Institutions that lack these foundations struggle to move beyond pilots, regardless of their ambition.

4.2.1 *Data Readiness*

Data readiness was cited as the **single most decisive success factor** in the survey (see Figure 7, Chapter 3.5). Institutions repeatedly emphasized that AI projects "fail on data, not on algorithms." Key requirements include:

- **Data quality and harmonization:** Reliable, consistent, and representative datasets are essential. Fragmented or siloed data remains one of the main barriers.
- **Accessibility and governance:** Controlled but efficient access to relevant datasets allows faster development and validation cycles.

- **Confidentiality and compliance:** Adherence to strict privacy and security standards, including encryption, localization, and auditability, ensures trust and regulatory approval.

Interviewees noted that institutions with **centralized data platforms or feature stores** were able to accelerate AI deployment significantly by reducing duplication and ensuring reproducibility across use cases.

4.2.2 MLOps Capabilities

MLOps—the set of practices for managing machine learning across its lifecycle—was highlighted as a **critical differentiator between pilots and production**. Respondents reported that without standardized MLOps practices, even promising pilots remain "one-offs." Institutions with mature MLOps practices demonstrated:

- **Reproducible pipelines:** Automated, versioned workflows from data ingestion to model deployment.
- **Monitoring and observability:** Continuous drift detection, performance tracking, and rollback mechanisms.
- **Validation frameworks:** Rigorous testing for bias, fairness, robustness, and explainability before go-live.
- **Traceability:** Full lineage of datasets, model versions, and decisions, ensuring compliance with FINMA and EU AI Act expectations.

4.2.3 Infrastructure Scalability

Institutions also pointed to the need for **scalable and modular infrastructures** that allow reusability across business units. Success was most evident in organizations that invested in:

- **Shared feature stores:** enabling consistency across models and reducing redundant engineering work.
- **API-driven architectures:** allowing AI components to be integrated flexibly into core systems.
- **Hybrid or compliant cloud solutions:** balancing efficiency and innovation with strict sovereignty requirements.

Smaller institutions often rely on vendor-provided platforms to fill these gaps, while larger banks increasingly invest in internal infrastructure. Both approaches, when aligned with governance and compliance frameworks, enable sustainable scaling.

4.2.4 Interview Insights

Interviewees stressed that **data readiness and MLOps maturity are non-negotiable foundations** for industrialization. One respondent summarized: *"We no longer ask if an algorithm works — we ask if it can be monitored, scaled, and audited."* This shift in focus illustrates the industry's transition from experimentation to operationalization.

4.2.5 Interpretation

The results confirm that **data and MLOps are the backbone of successful AI deployment**. Without high-quality, accessible, and secure data, models cannot be trusted. Without mature MLOps, pilots remain siloed experiments. Together, these foundations create the operational reliability and compliance confidence required to move AI from proof-of-concept to production at scale.

4.2.6 Key Implication

Swiss financial institutions must prioritize investment in **data readiness and governance and MLOps infrastructure** as a strategic enabler. This requires:

- Building or adopting shared, high-quality data platforms.
- Embedding monitoring, validation, and traceability into every model lifecycle.

- Treating infrastructure not as a support function but as a **product in itself**, with dedicated teams ensuring its scalability and resilience.

By doing so, institutions can transform AI initiatives from isolated pilots into **repeatable, industrialized solutions** that deliver consistent business value under Swiss regulatory conditions.

4.3 *Governance, Compliance, and Risk Management*

The findings of this study underscore that **governance and compliance are not peripheral issues but central enablers of AI deployment success**. Swiss financial institutions operate in a highly regulated environment, where regulatory expectations from FINMA and the forthcoming EU AI Act shape every stage of the AI lifecycle. Institutions that embed compliance and governance **from the start of the project** are far more likely to progress beyond PoCs, while those that treat compliance as an afterthought often stall.

4.3.1 *Compliance-by-Design*

The survey and interviews highlighted **compliance-by-design** as a best practice for shortening deployment cycles and building trust. Institutions that involve compliance officers and risk managers at the earliest stages of PoC design reported:

- Faster approval processes when moving to production.
- Clearer documentation of risks and mitigations.
- Greater confidence from regulators and internal audit teams.

Instead of being seen as a brake on innovation, compliance-by-design becomes a **strategic accelerator**, reducing rework and avoiding last-minute blockers.

4.3.2 *Governance Frameworks*

Robust governance structures are a defining feature of successful AI adopters. Respondents emphasized the importance of:

- **AI inventories:** centralized registries of all AI systems in use, including purpose, risk classification, and performance metrics.
- **Risk classification frameworks:** consistent methods for assessing AI risks across both traditional and generative AI use cases.
- **Accountability lines:** clear assignment of roles and responsibilities across the AI lifecycle, ensuring human oversight and decision authority remain intact.
- **Explainability standards:** requirements that models be interpretable to regulators, customers, and internal stakeholders alike.

These practices align closely with **FINMA Guidance 08/2024**, which calls for transparency, accountability, and independent oversight of AI systems in supervised institutions.

4.3.3 *Model Documentation and Transparency*

The survey results on **AI model cards** (see Chapter 3.7) demonstrate the strong institutional demand for standardized documentation. Institutions expect model cards to disclose:

- Risks, limitations, and bias considerations.
- Intended use and application boundaries.
- Validation and monitoring methods.
- Training data origin and quality.

By embedding such documentation requirements into governance frameworks, institutions not only meet regulatory expectations but also strengthen internal trust in AI systems.

4.3.4 *Risk Management Practices*

Effective governance extends into robust risk management. Respondents and interviewees stressed the following aspects:

- **Continuous monitoring:** tracking for data drift, model degradation, and unexpected outputs.
- **Independent reviews:** regular third-party or cross-team audits of models to ensure compliance and reliability.
- **Outsourcing controls:** rigorous due diligence and contractual safeguards when relying on external vendors.
- **Fairness and bias monitoring:** systematic checks to prevent discriminatory outcomes in areas such as credit scoring or AML.

These practices echo both **FINMA's supervisory priorities** and the EU AI Act's high-risk classification regime, which places strict obligations on financial AI applications.

4.3.5 *Interview Insights*

Interviews confirmed that **regulatory alignment is a critical determinant of deployment success**. One compliance manager remarked: *"The fastest way to kill a PoC is to ignore compliance until the end. The fastest way to scale one is to embed compliance from the beginning."* Another executive noted that robust governance not only helps with regulators but also strengthens internal decision-making: *"When we can clearly explain risks and controls, our board is much more willing to greenlight production."*

4.3.6 *Interpretation*

The evidence shows that governance and compliance are **gatekeepers, not add-ons**. Institutions that lack structured governance frameworks face uncertainty, delays, and regulatory pushback. By contrast, those that integrate compliance early and adopt clear governance models create an environment of trust and predictability that enables scaling.

4.3.7 *Key Implication*

For Swiss financial institutions, **governance and compliance are the bridge between experimentation and industrialization**. Institutions should:

- Mandate compliance-by-design from PoC inception.
- Maintain centralized AI inventories and consistent risk classification frameworks.
- Adopt standardized model documentation (model cards) across all use cases.
- Implement continuous monitoring, independent reviews, and fairness audits.

These practices not only satisfy regulatory expectations but also build the organizational confidence needed to deploy AI responsibly and at scale.

4.4 *Adoption, Change Management, and Scaling*

While data foundations, infrastructure, and governance provide the technical and regulatory backbone of AI deployment, **adoption and change management determine whether AI solutions generate real value in daily operations**. Survey and interview findings confirm that organizational readiness is not the biggest barrier compared to compliance and data issues, yet institutions that actively manage adoption achieve far smoother scaling and higher impact.

4.4.1 *IT–Business Cooperation*

A recurring theme across responses is the need for **shared ownership of AI projects** between IT and business functions. Successful institutions avoid siloed pilots by ensuring:

- **Joint project teams** with representation from both business units and technical experts.
- **Clear accountability** for deployment and ongoing operations that extends beyond the innovation team.
- **Co-creation of requirements and KPIs**, ensuring that business users define success metrics while IT ensures feasibility.

Interviewees stressed that projects framed as "business-led but tech-enabled" were the most likely to secure long-term adoption.

4.4.2 *Change Management*

Adoption does not happen automatically. Institutions reported that user acceptance is strengthened by:

- **Early involvement of end users** in design and testing phases, increasing trust and usability.
- **Transparent communication** of AI's role and limitations, reducing concerns about "black-box" systems.
- **Training programs** that enable employees to use AI effectively and understand its outputs.
- **Cultural readiness**: framing AI as an augmentation tool that supports employees rather than a replacement.

Survey evidence confirmed that while **user resistance was not a top-rated barrier** (see Chapter 3.4), proactive change management was nonetheless a key enabler of successful adoption.

4.4.3 *Pragmatic Accelerators*

Institutions also highlighted pragmatic measures that helped them accelerate deployment and adoption:

- **Out-of-the-box solutions:** Pre-built applications with proven compliance and configuration reduced complexity and time-to-value.
- **Vendor maturity:** Partnering with providers that could demonstrate regulatory alignment and reference cases in financial services.
- **Ecosystem partnerships:** Leveraging networks such as SFTI and academic institutions (e.g., OST) for knowledge sharing and benchmarking.

These accelerators helped smaller and mid-sized institutions in particular, which often lack the in-house resources for bespoke development.

4.4.4 *Scaling Practices*

Moving from one successful deployment to multiple scaled solutions requires deliberate practices:

- **Stage-gates with production-first criteria:** Ensuring every PoC includes a defined path-to-production and required documentation.
- **Playbooks and templates:** Codifying lessons learned into reusable guides for future projects.
- **Benefit tracking:** Continuous measurement of KPIs post-deployment to validate business value and secure investment for scaling.
- **Portfolio governance:** Maintaining a live heatmap of AI use cases to balance experimentation with industrialization.

Institutions that adopted these practices reported higher consistency in advancing from pilots to repeatable, enterprise-wide deployments.

4.4.5 *Interview Insights*

One project lead noted: *"Scaling isn't just about technology; it's about trust. When business users feel ownership and see benefits tracked in their KPIs, adoption follows naturally."* Another emphasized the importance of **playbooks**: *"After our first deployment, we created a template process. That was the turning point for scaling — we didn't have to reinvent the wheel every time."*

4.4.6 *Interpretation*

The findings demonstrate that **successful AI adoption requires both structured governance and proactive cultural engagement**. Technical readiness may open the door to production, but without user trust, training, and shared ownership, AI solutions risk underutilization. Conversely, when business units are engaged, change is managed openly, and benefits are measured transparently, scaling becomes repeatable and sustainable.

4.4.7 *Key Implication*

For Swiss financial institutions, adoption and scaling depend on **treating AI as both a technical and organizational transformation**. Institutions should:

- Ensure joint IT–business ownership of AI initiatives.
- Invest in change management and training to build user trust.
- Leverage out-of-the-box solutions and ecosystem partnerships to accelerate deployment.
- Institutionalize scaling practices through stage-gates, playbooks, and continuous benefit tracking.

By embedding these practices, institutions can ensure that AI moves from being a set of promising pilots to a **core operational capability** that delivers sustained business impact.

4.5 *Synthesis: The Success Equation*

The preceding sections demonstrate that **AI deployment success is not achieved by any single factor**, but by the interaction of multiple enablers across leadership, data, governance, and adoption. Institutions that combine these elements consistently outperform those that treat AI as isolated pilots or technology experiments. The study therefore synthesizes the findings into a simple but powerful formula — the **Success Equation** (see Figure 10):

Figure 10: The Success Equation

4.5.1 *Elements of the Success Equation*

The analysis of the survey data and institutional experiences reveals that successful AI deployment depends on the interplay of four essential dimensions. Together, these elements form a practical "success equation" that determines whether projects can transition from PoC to scalable production. Each component—value case, data readiness and MLOps, governance and controls, and adoption—contributes a critical layer to sustainable implementation. The following points summarize the key elements of this success formula:

1. **Value Case**
   - Clear, measurable business outcomes tied to KPIs.
   - Accountable business ownership and realistic ROI expectations.

- Framed around productivity, user experience, and quality rather than exaggerated cost savings.

2. **Data Readiness & MLOps**
   - High-quality, harmonized, and accessible datasets.
   - Robust MLOps pipelines with monitoring, drift detection, and traceability.
   - Infrastructure designed for reusability and scalability across multiple projects.
3. **Governance & Controls**
   - Compliance-by-design embedded from inception.
   - Centralized AI inventories, risk classification frameworks, and explainability standards.
   - Model documentation (model cards) covering risks, validation, and training data.
4. **Adoption**
   - Joint ownership between IT and business functions.
   - Proactive change management, training, and cultural readiness.
   - Scaling practices including stage-gates, playbooks, and benefit tracking.

4.5.2 Interaction and Multiplication

The equation is deliberately framed as a **multiplication rather than addition**. This reflects the reality that weaknesses in any one domain undermine the entire initiative:

- A strong business case cannot succeed without reliable data.
- Excellent MLOps cannot compensate for lack of compliance alignment.
- Robust governance fails if users do not adopt the system in practice.

Only when all four elements are present and reinforcing one another can institutions achieve **sustainable production impact**.

4.5.3 Interview Insights

Several interviewees highlighted this interdependency. One AI lead observed: *"We had data and algorithms, but no governance — it stalled. Later, we had governance and strategy, but no MLOps — it stalled again. It only worked once all pieces came together."* This confirms that **holistic readiness** is essential.

4.5.4 Key Implication

The **Success Equation** provides a practical lens for assessing institutional readiness. Swiss financial institutions should use it as a checklist to identify where weaknesses may be undermining their AI ambitions. By strengthening all four dimensions in parallel, institutions can systematically bridge the PoC–production gap and position themselves for sustainable AI-driven value creation.

4.6 Overall Conclusion

The analysis of success factors has shown that bridging the PoC–production gap requires a holistic interplay of strategy, data, governance, and adoption. Institutions that master these elements achieve sustainable deployment impact, while those lacking in one or more areas remain trapped in pilot purgatory. To translate these insights into practical action, the next chapter introduces a **framework and set of tools** derived from the study's findings. This framework provides financial institutions with a structured approach to assess their readiness, implement best practices, and scale AI initiatives effectively across their organizations.

5. Framework and Practical Tools for Scaling AI

The previous chapters identified the barriers that keep many AI initiatives in the pilot stage and the success factors that enable a transition to production. Building on these insights, this chapter introduces a **structured framework** designed to support Swiss financial institutions in scaling AI responsibly and effectively. The framework translates the empirical findings of the study into a **practical toolbox** that institutions can apply to their own contexts, regardless of size, maturity, or functional focus.

At its core, the framework provides a **prioritization hierarchy**—illustrated in a seven-step pyramid—that shows how institutions can build the essential foundations first, and then layer on advanced practices to achieve sustainable impact. Alongside this conceptual structure, the chapter introduces **practical tools** such as checklists, stage-gate criteria, model cards, and maturity models. Together, these elements form a **comprehensive roadmap** to bridge the PoC–production gap and unlock lasting business value from AI initiatives in the Swiss financial sector.

5.1 Purpose and Design of the Framework

The survey and interview findings presented in Chapters 3 and 4 revealed a consistent set of barriers and success factors that determine whether AI initiatives in Swiss financial institutions progress from proof-of-concept (PoC) to scalable production. While these findings provide valuable insights on their own, institutions need a **structured framework** that translates the lessons into **practical guidance** for implementation.

The purpose of this framework is therefore twofold:

1. To provide a **comprehensive yet pragmatic structure** that captures the essential dimensions of successful AI deployment.
2. To equip Swiss financial institutions with **practical tools and checklists** that can be applied directly to assess readiness, design PoCs for production, and scale AI initiatives responsibly.

5.1.1 Principles of Design

The framework builds on three key design principles:

- **Practicality and Usability**
  The framework is not an abstract academic model but a **hands-on toolbox** for practitioners. Each element is accompanied by clear descriptions, guiding questions, and success conditions that allow institutions to apply the framework to their own projects.
- **Scalability and Flexibility**
  The framework is designed to be **scalable across institutions of different sizes and maturities**. While larger incumbents may apply it to enterprise-wide governance structures, regional or cantonal banks can adapt it to smaller teams and incremental deployments.
- **Regulatory Alignment and Risk Awareness**
  Given Switzerland's stringent regulatory environment, the framework incorporates **compliance-by-design principles** and aligns with expectations from FINMA Guidance 08/2024 as well as international standards such as the EU AI Act. This ensures that institutions adopting the framework not only enable operational efficiency but also build resilience and regulatory trust.

5.1.2 Scope of Application

The framework is **functionally agnostic**, meaning it can be applied across different use cases—from compliance and risk management to operations and support functions. While many survey respondents emphasized compliance-related AI use cases, the framework has been structured to serve as a **general guide for all AI applications in financial services**.

5.1.3 From Insight to Action

Most importantly, the framework translates the **success equation** identified in Chapter 4.5 into actionable components. It operationalizes the four critical dimensions of successful deployment—**Strategic Leadership & Business Alignment, Data Foundations & MLOps Infrastructure, Governance & Compliance, and Adoption & Scaling**—into concrete steps. Each dimension is broken down into **logical components and practical tools**, ensuring that institutions can move from insight to implementation.

In the following sections, we present an overview of the framework, outline its components, and introduce practical tools such as stage-gate templates, success checklists, and a maturity model that together form a **comprehensive roadmap for bridging the PoC–production gap in Swiss financial institutions**.

5.2 Framework Overview: The Four Dimensions

The framework builds directly on the **success equation** outlined in Chapter 4.5 and translates it into a **structured roadmap** for Swiss financial institutions. Its design reflects two complementary logics:

1. **The Four Core Dimensions** – the fundamental building blocks that must be present for AI projects to succeed:
    - Strategic Leadership & Business Alignment
    - Data Foundations & MLOps Infrastructure
    - Governance, Compliance & Risk Management
    - Adoption, Change Management & Scaling
2. **The Seven Action Steps** – a practical prioritization hierarchy, represented as a pyramid (see Figure 11), which shows how institutions can layer specific practices to build maturity step by step.

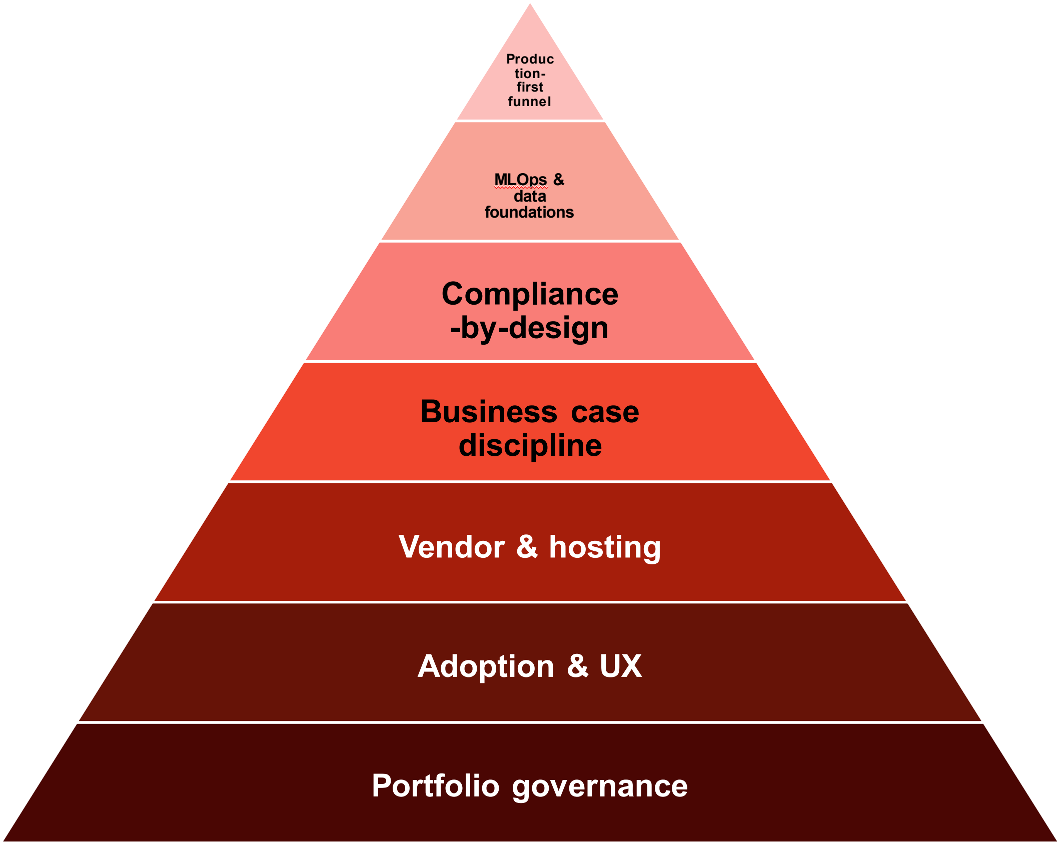

Figure 11: The Seven-Step Pyramid – Action hierarchy for bridging the PoC–production gap

5.2.1 *The Four Core Dimensions*

Each dimension captures a distinct set of success factors identified in the survey and interviews:

- **Strategic Leadership & Business Alignment**
  Ensures that AI initiatives are tied to measurable business outcomes, supported by executive sponsorship, and embedded in realistic ROI expectations.
- **Data Foundations & MLOps Infrastructure**
  Provides the technical reliability required for scaling, including high-quality data, reproducible pipelines, monitoring, and infrastructure designed for reusability.
- **Governance, Compliance & Risk Management**
  Embeds compliance-by-design and risk management into every stage, ensuring regulatory approval, accountability, and stakeholder trust.
- **Adoption, Change Management & Scaling**
  Focuses on organizational integration, user trust, cultural readiness, and systematic practices for replicating success across portfolios.

These four dimensions form the **conceptual backbone** of the framework.

5.2.2 *The Seven Action Steps*

To operationalize the four dimensions, the framework is further structured into **seven practical action steps**, prioritized in a pyramid sequence (see Figure 10). Each layer builds on the one below, emphasizing that institutions must **establish foundations first** before progressing to advanced scaling practices:

1. **Production-first funnel** – Mandate path-to-production artifacts for all PoCs; cap concurrent pilots; fund only the most promising 10–15%.
2. **MLOps & data foundations** – Build shared feature stores, continuous integration/continuous deployment (CI/CD) pipelines, monitoring, and infrastructure-as-product.

3. **Compliance-by-design** – Integrate risk and legal requirements from day one; implement standardized model cards.
4. **Business case discipline** – Tie initiatives to KPI baselines; assign accountable owners; enforce "kill/scale fast" decisions.
5. **Vendor & hosting** – Apply sovereignty and security as stage-gates; negotiate scalable licensing models.
6. **Adoption & UX** – Prioritize internal, employee-facing use cases; measure time saved and quality improvements.
7. **Portfolio governance** – Maintain use-case heatmaps; publish playbooks and templates for replication across units.

5.2.3 Linking the Four Core Dimensions and the Seven Action Steps

While the **four core dimensions** describe *what must be in place* for AI deployment to succeed, the **seven action steps** explain *how institutions can build these capabilities in practice*. The framework is therefore structured as a **two-layer system**:

- The four **dimensions** provide the **strategic foundation**—the thematic areas in which readiness must be achieved.
- The seven **action steps** represent the **implementation sequence**—the practical levers through which institutions strengthen each dimension over time.

Each action step contributes to one or more of the four dimensions, creating a matrix rather than a linear hierarchy. The seven steps thus serve as the *operational pathways* for developing maturity within the dimensions (see Table 1):

Core Dimension	Supported by Action Steps	Explanation
1. Strategic Leadership & Business Alignment	▪ *Step 1 (Production-first funnel)* ▪ *Step 4 (Business case discipline)*	▪ *These steps ensure that AI initiatives are strategically directed, measurable, and backed by accountable ownership.*
2. Data Foundations & MLOps Infrastructure	▪ *Step 2 (MLOps & data foundations)* ▪ *Step 5 (Vendor & hosting)*	▪ *They address the technical backbone—data quality, reproducibility, infrastructure scalability, and sourcing strategy.*
3. Governance, Compliance & Risk Management	▪ *Step 3 (Compliance-by-design)* ▪ *Step 7 (Portfolio governance)*	▪ *Together they establish consistent oversight, inventories, risk classification, and transparency across the AI portfolio.*
4. Adoption, Change Management & Scaling	▪ *Step 6 (Adoption & UX)* ▪ *Step 7 (Portfolio governance)*	▪ *These steps drive organizational integration, user acceptance, and replication of successful deployments.*

Table 1: Linking the Four Core Dimensions and the Seven Action Steps

This structure allows the framework to be used **both vertically and horizontally**:

- **Vertically**, institutions can climb the seven steps as a maturity sequence—building progressively from foundation to scale.
- **Horizontally**, they can assess their readiness across the four dimensions—ensuring balance between strategy, technology, governance, and adoption.

In short, the four dimensions define *what success looks like* in each area of AI deployment, while the seven action steps describe *how to get there*. Their interaction provides both a diagnostic map and a practical roadmap for moving AI from PoC to sustainable production.

5.3 Logical Components of the Framework

To move from conceptual dimensions and action steps toward **practical application**, the framework is broken down into **logical components**. These components describe *what must be in place* for each dimension of the framework to function effectively. In total, the framework consists of approximately **30–33 components**, grouped under the four core dimensions (see Chapter 5.2).

Each component is defined with its rationale and success conditions, ensuring institutions can assess their readiness systematically (see Table 1 - Table 4).

1. Strategic Leadership & Business Alignment	
Executive Sponsorship	▪ *Definition:* Visible and active support from senior leadership ▪ *Rationale:* Provides prioritization, resources, and speed in decision-making ▪ *Success Conditions*: Executives actively monitor projects, remove blockers, and engage compliance/IT teams
Business Case Discipline	▪ *Definition:* Use cases tied to clear KPIs and accountable business ownership ▪ *Rationale*: Ensures credibility and measurable value ▪ *Success Conditions:* ROI framed in terms of productivity, user experience, and quality gains, not only cost savings
Expectation Management	▪ *Definition:* Realistic ROI assumptions and acceptance of PoC termination ▪ *Rationale:* Prevents overpromising and protects institutional credibility ▪ *Success Conditions:* Structured funnel with stage-gates; lessons from failed PoCs integrated into portfolio learning
Portfolio Prioritization	▪ *Definition:* Stage-gates that cap PoCs and fund only the most promising 10–15% ▪ *Rationale:* Prevents resource dilution and increases hit rate ▪ *Success Conditions:* Clear entry/exit criteria and transparent decision-making

Table 2: Framework Component 1 - Strategic Leadership & Business Alignment

2. Data Foundations & MLOps Infrastructure	
Data Quality Standards	▪ *Definition:* Policies and controls to ensure accuracy, completeness, and representativeness of data ▪ *Rationale:* Poor data is the most cited cause of AI failure ▪ *Success Conditions:* Regular data audits and monitoring
Data Accessibility & Governance	▪ *Definition:* Secure, governed access to relevant datasets ▪ *Rationale:* Speeds up development and ensures compliance ▪ *Success Conditions:* Access rights embedded in data governance policies
Feature Stores	▪ *Definition:* Centralized, reusable repositories of engineered features ▪ *Rationale:* Improves reproducibility and scalability across use cases ▪ *Success Conditions:* Shared across projects, monitored for drift and accuracy
MLOps Pipelines	▪ *Definition:* Automated, version-controlled workflows for model lifecycle ▪ *Rationale:* Differentiator between pilots and production ▪ *Success Conditions:* CI/CD pipelines with full traceability
Monitoring & Observability	▪ *Definition:* Continuous monitoring of models in production ▪ *Rationale:* Detects drift, bias, and degradation early ▪ *Success Conditions:* Alerts, rollback mechanisms, and dashboards in place
Scalability & Infrastructure-as-Product	▪ *Definition:* Infrastructure designed for modularity and reusability ▪ *Rationale:* Prevents one-off solutions ▪ *Success Conditions:* Shared services (API-driven), hybrid hosting compliant with sovereignty requirements

Table 3: Framework Component 2 - Data Foundations & MLOps Infrastructure

3. Governance, Compliance & Risk Management	
Compliance-by-Design	▪ *Definition:* Involving compliance and risk functions from day one ▪ *Rationale:* Avoids late-stage blockers and accelerates approvals ▪ *Success Conditions:* Regulatory criteria embedded into PoC design
AI Inventories	▪ *Definition:* Centralized registries of all AI systems ▪ *Rationale:* Transparency for regulators, auditors, and management ▪ *Success Conditions:* Updated inventory with purpose, risk level, and metrics
Risk Classification Frameworks	▪ *Definition:* Consistent methods for categorizing AI risks ▪ *Rationale:* Prevents inconsistent definitions across institutions ▪ *Success Conditions:* Applied to all AI systems, including generative AI
Explainability Standards	▪ *Definition:* Clear requirements for interpretable models ▪ *Rationale:* Regulatory necessity and trust builder ▪ *Success Conditions:* Non-technical explanations available for key stakeholders
Model Documentation (Model Cards)	▪ *Definition:* Standardized documentation covering risks, limitations, validation, and training data ▪ *Rationale:* Supports compliance-by-design and internal trust ▪ *Success Conditions:* Mandatory for all production-ready AI systems
Independent Reviews & Oversight	▪ *Definition:* Regular third-party or cross-team audits of models ▪ *Rationale:* Strengthens credibility and regulatory acceptance ▪ *Success Conditions:* Findings integrated into lifecycle governance
Outsourcing Controls	▪ *Definition:* Governance mechanisms for vendor-provided AI ▪ *Rationale:* Mitigates operational and legal risks ▪ *Success Conditions:* Contracts defining accountability and compliance standards

Table 4: Framework Component 3 - Governance, Compliance & Risk Management

4. Adoption, Change Management & Scaling	
IT–Business Cooperation	▪ *Definition:* Joint ownership of AI projects ▪ *Rationale:* Ensures alignment of goals and resources ▪ *Success Conditions:* Co-created KPIs and co-led deployment teams
User Involvement	▪ *Definition:* Early involvement of end users in testing and design ▪ *Rationale:* Builds trust and usability ▪ *Success Conditions:* Pilot groups, feedback loops, and transparent communication
Change Management Programs	▪ *Definition:* Structured initiatives to prepare staff for AI adoption ▪ *Rationale:* Enhances cultural readiness and reduces resistance ▪ *Success Conditions:* Training sessions, transparent communication, focus on augmentation rather than replacement
Adoption Metrics	▪ *Definition:* Measurement of user acceptance and usage patterns ▪ *Rationale:* Ensures AI solutions are embedded in workflows ▪ *Success Conditions:* KPIs such as adoption rate, satisfaction, and time-to-value
Playbooks & Templates	▪ *Definition:* Standardized guides capturing lessons learned ▪ *Rationale:* Avoids reinventing the wheel ▪ *Success Conditions:* Updated after each deployment, shared across teams
Portfolio Heatmaps	▪ *Definition:* Visual overview of all AI use cases across lifecycle stages ▪ *Rationale:* Enables portfolio governance and resource prioritization ▪ *Success Conditions:* Maintained and updated regularly by AI governance teams
Benefit Tracking	▪ *Definition:* Continuous measurement of business impact post-deployment ▪ *Rationale:* Secures credibility and future investment ▪ *Success Conditions:* Metrics tied to business KPIs, tracked over time

Table 5: Framework Component 4 - Adoption, Change Management & Scaling

5.3.1 *Interpretation*

The breakdown into logical components allows institutions to **self-assess systematically**: each component can be checked for presence, maturity, and effectiveness. Taken together, they form a **comprehensive operationalization** of the four dimensions and seven action steps introduced in Chapter 5.2.

5.3.2 *Key Implication*

Swiss financial institutions should use these components as a **readiness checklist**, identifying strengths and gaps across their AI lifecycle. Addressing missing components systematically increases the likelihood of progressing from PoC to scalable production.

5.4 *Practical Tools for Implementation*

To ensure that the framework moves beyond conceptual guidance, it has been translated into a set of **practical tools**. These instruments provide Swiss financial institutions with ready-to-use templates and methods to **assess readiness, structure PoCs for production, and scale AI responsibly** – based on and adapted with their individual situation and context.

5.4.1 *Success Checklist (Playbook)*

The **success checklist** is designed as a quick-assessment tool covering the four dimensions of the framework. Each section contains guiding questions that can be answered with *yes/no/partially*. Gaps identified here indicate areas requiring improvement before scaling (see Table 5).

Example Success Checklist Questions	Yes	No	Partially
Strategic Leadership & Business Alignment			
▪ Has executive sponsorship been formally assigned?	☐	☐	☐
▪ Are business KPIs and ROI baselines defined for this initiative?	☐	☐	☐
Data Foundations & MLOps Infrastructure			
▪ Is the dataset clean, complete, representative, and free of major quality issues?	☐	☐	☐
▪ Does the project have a reproducible MLOps pipeline?	☐	☐	☐
Governance & Compliance			
▪ Has the use case been registered in the AI inventory?	☐	☐	☐
▪ Is a model card drafted, covering risks, limitations, and validation?	☐	☐	☐
Adoption & Scaling			
▪ Have end users been involved in the PoC testing phase?	☐	☐	☐
▪ Are adoption and benefit-tracking KPIs defined	☐	☐	☐

Table 6: The Success Checklist (Playbook)

Institutions can apply the checklist both **before approving a PoC** and **before production rollout** to ensure readiness.

5.4.2 *Stage-Gate Template*

The **stage-gate template** defines the mandatory criteria that a PoC must satisfy to progress to production. It helps institutions avoid resource drain on weak projects and ensures **production-first design** (see Table 6).

Stage-Gate Template / Criteria Examples	
Architecture & Infrastructure	▪ Deployment architecture documented ▪ Integration requirements validated
Compliance-by-Design	▪ Compliance team sign-off on risk assessment ▪ Compliance team sign-off on model documentation
MLOps Readiness	▪ Monitoring plan in place ▪ Rollback plan in place ▪ Reproducibility plan in place
Business Case	▪ KPIs defined ▪ Accountable business owner assigned ▪ ROI baseline documented
Adoption	▪ End-user pilot feedback collected ▪ Training plan drafted

Table 7: The Stage-Gate Template

The stage-gate should be managed by a **multidisciplinary review board** (business, IT, compliance, risk).

5.4.3 Model Card Template

The **model card template** ensures that all production-ready AI systems are documented in a standardized way (see Table 7). It builds on survey results showing the importance of transparency about risks, validation, and data (see Chapter 3.7). Institutions can embed this template into their **compliance approval workflows** to ensure consistency.

Core Elements of Model Card Template
▪ **Purpose & Intended Use:** Description of scope and limitations
▪ **Risks & Limitations:** Known failure modes, bias considerations, edge cases
▪ **Validation & Testing:** Performance benchmarks, fairness checks, stress tests
▪ **Training Data Transparency:** Data sources, representativeness, preprocessing methods
▪ **Governance & Oversight:** Responsible owner, human oversight mechanisms, review intervals

Table 8: Model Card Template (Core Elements)

5.4.4 Use Case Portfolio Heatmap

The **portfolio heatmap** provides a visual overview of all AI initiatives across their lifecycle stages—ideation, PoC, deployment, and scaling. This tool helps institutions **govern their AI portfolio strategically** by identifying bottlenecks and prioritizing resources (see Figure 12).

Maturity / Phase	Portfolio of Use Cases / Projects
Ideation 65%	4821 GenAI Client Advisor - Wealth Management - HPB 0376 ESG Insights Engine - Asset Management - CIO 9154 AML Anomaly Radar - Compliance - CCO 2640 KYC Document Copilot - Onboarding - CLM 7083 Smart Credit Underwriter - Retail Lending - CRO 5912 FX Trade Co-Pilot - Markets - HST 8307 Relationship 360 Scoring - Corporate Banking - HCC 1045 GenAI Policy Guard - Risk & Controls - CRO 3769 Advisor Note Summarizer - Wealth - HOA 5284 Synthetic Data Lab - Data & Analytics - CDO 6931 Instant Dispute Triage - Operations - COO 2178 Personalized Next Best Action – Dig. Channels - CDO 8456 Treasury Liquidity Forecaster - Treasury – GTR
Pilot / PoC 15%	1302 Client Voice Miner - CX & NPS – HCX 4795 Sanctions Screening Booster - Compliance - HFC 9028 Model Risk Explainability - Model Risk – MRM
Deployment 10%	3517 Cross-Border Tax Classifier - Tax & Legal - GCO 6642 Payments Fraud Shield - Payments – HOP
Sacling / Optimization 10%	2889 Sustainable Portfolio Optimizer - Wealth/AM - HSI 7403 GenAI Knowledge Search - Enterprise IT - CIO

Figure 12: The Use Case Portfolio Heatmap (Funnel) – Example Projects

Further **practical applications** of the portfolio heatmap (or funnel) are:

- Monitoring the overall balance between early-stage pilots and scaled deployments.
- Identifying clusters of stalled projects (e.g., PoCs blocked by data issues).
- Highlighting successes that can be replicated in other functions.
- Supporting strategic decisions on where to invest or terminate projects.

The heatmap should be updated regularly (e.g., quarterly) and reviewed by executive committees responsible for AI strategy.

5.4.5 Playbooks and Templates for Scaling

In addition to individual tools, institutions are encouraged to develop **playbooks**—internal manuals that codify best practices for AI deployment. These playbooks can include:

- **Case examples:** What worked, what failed, and why.
- **Reusable templates:** Standardized business case forms, compliance documentation, monitoring dashboards.
- **Guidelines for adoption:** Communication strategies, training modules, user acceptance testing protocols.

Playbooks reduce friction by avoiding repeated reinvention and enable **institutional learning across projects and teams**.

5.4.6 Interpretation

Together, these tools make the framework **immediately actionable**. They help institutions ensure that PoCs are designed with production in mind, that compliance and governance requirements are systematically addressed, and that scaling is supported by standardized documentation and portfolio oversight.

5.4.7 Key Implication

Swiss financial institutions should **embed these tools into their AI governance processes**, making them part of the formal lifecycle of every AI project. By doing so, they can institutionalize best practices, accelerate deployment, and reduce risks, thereby increasing their chances of bridging the PoC–production gap.

5.5 Maturity Model for Scaling AI

While the framework provides the dimensions and tools for successful AI deployment, institutions differ widely in their current stage of readiness. To support **self-assessment and structured progress**, the study proposes a **maturity model** for scaling AI. This model defines **five levels of maturity**, each characterized by distinct practices, capabilities, and organizational mindsets across the four dimensions of the framework (see Table 8 through to Table 12).

The maturity model serves two purposes:

1. To help institutions **diagnose their current level of AI readiness**.
2. To provide a **roadmap for progression**, highlighting what must be in place to move to the next level.

Level 1 – Exploratory	
▪ **Strategic Leadership & Alignment:** AI treated as experimentation, with limited executive involvement.	▪ *Typical Outcome:* Many PoCs, few credible paths to production.
▪ **Data & MLOps:** Datasets are fragmented, not cleansed; no reproducible pipelines.	
▪ **Governance & Compliance:** Ad-hoc handling of compliance, often deferred until late stages.	
▪ **Adoption & Scaling:** Isolated pilots with little user involvement; limited awareness of adoption needs.	

Table 9. The Maturity Model for Scaling AI – Level 1

Level 2 – Emerging	
▪ **Leadership:** Growing recognition of AI's strategic importance; sporadic executive sponsorship.	▪ *Typical Outcome:* Some production deployments, but inconsistent and fragile.
▪ **Data & MLOps:** Basic data governance structures; some manual monitoring of models.	
▪ **Governance:** Early compliance checks in selected pilots; no centralized AI inventory.	
▪ **Adoption:** Users consulted late; limited training or communication.	

Table 10: The Maturity Model for Scaling AI – Level 2

Level 3 – Managed	
▪ **Leadership:** Clear executive sponsorship; AI projects tied to defined business cases.	▪ *Typical Outcome:* First wave of stable production deployments; confidence building among stakeholders.
▪ **Data & MLOps:** Initial MLOps pipelines established; centralized feature stores piloted.	
▪ **Governance:** AI inventory introduced; risk classification frameworks tested.	
▪ **Adoption:** Structured change management introduced; training programs in place.	

Table 11: The Maturity Model for Scaling AI – Level 3

Level 4 – Industrialized	
▪ **Leadership:** AI embedded in strategic planning; KPIs tracked at portfolio level.	▪ *Typical Outcome:* AI projects scale across functions; regulators and management see AI as controlled and trustworthy.
▪ **Data & MLOps:** Mature data governance; reproducible CI/CD pipelines across multiple projects.	
▪ **Governance:** Compliance-by-design standard; model cards mandatory for all production models.	
▪ **Adoption:** Playbooks and templates used systematically; adoption metrics monitored.	

Table 12: The Maturity Model for Scaling AI – Level 4

Level 5 – Scaled	
▪ **Leadership:** AI viewed as an enterprise-wide capability; innovation portfolios linked directly to business strategy.	▪ *Typical Outcome:* Sustainable competitive advantage from AI, with measurable business impact across the institution.
▪ **Data & MLOps:** Fully integrated data platforms; automated monitoring, rollback, and auditability.	
▪ **Governance:** Governance frameworks harmonized with international standards (FINMA, EU AI Act); independent reviews embedded.	
▪ **Adoption:** AI solutions widely adopted by business users; continuous learning cycles ensure institutional resilience.	

Table 13: The Maturity Model for Scaling AI – Level 5

5.5.1 *Application of the Maturity Model*

Practical applications of the Maturity Model for Scaling AI include:

- Institutions can **self-assess** their current maturity by reviewing practices across the four dimensions.
- The model provides **guidance for progression**: each level defines the necessary preconditions for advancing to the next.
- By linking maturity levels to the framework tools (checklists, stage-gates, model cards, heatmaps), institutions can prioritize investments systematically.

5.5.2 *Interpretation*

The maturity model highlights that **scaling AI is a staged journey**. Institutions cannot skip foundational steps: for example, robust governance (Level 3) is required before true industrialization (Level 4) is possible. Similarly, cultural adoption at Level 4 is critical before an institution can achieve fully scaled AI impact at Level 5.

5.5.3 *Key Implication*

By situating themselves within the maturity model, Swiss financial institutions gain a **clear roadmap for institutional progress**. This allows them to benchmark against peers, identify missing capabilities, and prioritize actions that strengthen readiness for sustainable AI deployment.

5.6 *Application Scenarios*

To demonstrate the practical relevance of the framework, this section presents two short **application scenarios** based on insights from survey responses and interviews (see Table 13 and Table 14).

These anonymized vignettes illustrate how institutions encounter challenges in moving AI projects from PoC to production, and how the framework could be applied to improve outcomes.

Scenario A – A Successful Deployment	
Context	▪ A mid-sized regional bank developed an AI-powered **transaction monitoring tool** to improve Anti-Money Laundering (AML) compliance. The project began as a small pilot in 2023, focused on detecting suspicious transaction patterns more effectively than rule-based systems.
Challenges	▪ **Integrating** the tool with legacy core banking systems.
	▪ **Gaining approval** from compliance and risk departments to use sensitive client data.
Approach	▪ From the outset, the project had **executive sponsorship**, with the Chief Risk Officer acting as sponsor.
	▪ A **production-first funnel** was applied: the PoC was designed with clear integration requirements, monitoring, and rollback plans.
	▪ **Compliance-by-design** was enforced: risk and compliance teams were part of the PoC design, and a model card was drafted early.
	▪ Data access was supported through a **centralized feature store** that ensured traceability and auditability.
	▪ **Adoption was prioritized**: investigators were involved in testing, training programs were rolled out, and benefit tracking (false positive reduction, case throughput) was monitored.
Outcome	▪ The solution went into production in 2024 and **reduced false positives by 30%,** while improving case handling speed. The project has since been scaled to cover other monitoring tasks
Lesson	▪ The case demonstrates how **executive sponsorship, compliance-by-design, and robust data/MLOps foundations** form the decisive combination for successful deployment.

Table 14: Scenario A – A Successful Deployment

Scenario B – A Stalled Initiative	
Context	▪ A large Swiss-wide bank launched a pilot to use **AI-driven document analysis** for automating credit onboarding. The project aimed to improve efficiency by extracting and verifying customer data from unstructured documents.
Challenges	▪ The PoC was designed without a clear path to production.
	▪ Compliance was engaged only late, raising concerns about data confidentiality and explainability.
	▪ The business case promised large cost savings but lacked clear KPIs tied to measurable outcomes.
	▪ No reproducible MLOps pipeline was in place; the pilot ran in an isolated environment.
Outcome	▪ Despite technical success at the PoC stage, the project stalled before deployment. Compliance concerns and lack of benefit ownership prevented further investment, and the initiative was eventually discontinued.
Lesson	▪ This case illustrates the **risks of weak business case discipline and late-stage compliance engagement**. Applying the framework's **stage-gate template** and **success checklist** could have highlighted the project's deficiencies early, saving time and resources.

Table 15: Scenario B – A Stalled Initiative

5.6.1 Interpretation

These two scenarios illustrate the **contrast between structured and unstructured approaches** to AI deployment:

- In the successful case, strategic sponsorship, compliance integration, and data/MLOps readiness created the conditions for production success.
- In the stalled case, the absence of these fundamentals led to wasted effort despite technical feasibility.

5.6.2 Key Implication

The scenarios confirm that **bridging the PoC–production gap is not about proving technical capability, but about ensuring organizational readiness across all framework dimensions**. Applying the framework systematically enables institutions to identify weaknesses early and maximize their chances of scaling AI initiatives effectively.

5.7 Synthesis and Outlook

The framework presented in this chapter translates the empirical findings of the study into a **structured roadmap for action**. It combines the four core dimensions of success—**leadership and alignment, data and MLOps, governance and compliance, adoption and scaling**—with a **seven-step action hierarchy** that shows how institutions can build maturity layer by layer.

By breaking these elements down into **logical components** (Chapter 5.3) and equipping them with **practical tools** such as checklists, stage-gates, model cards, and portfolio heatmaps (Chapter 5.4), the framework provides Swiss financial institutions with **immediately usable instruments**. The maturity model (Chapter 5.5) adds a structured lens for self-assessment and progression, while the application scenarios (Chapter 5.6) illustrate how the framework works in practice—highlighting both success stories and common pitfalls.

5.7.1 Key Takeaways

The findings of this study converge around several overarching lessons that define what it takes to scale AI successfully within Swiss financial institutions. They emphasize that sustainable impact depends less on isolated technical excellence and more on the alignment of strategy, governance, data, and adoption capabilities. The following key takeaways summarize the core insights derived from the research and framework analysis:

- **Holistic readiness is essential:** Weaknesses in any one dimension—data, governance, leadership, or adoption—can stall projects, regardless of technical feasibility.
- **Foundations come first:** Institutions must establish robust data governance, MLOps, and compliance frameworks before scaling adoption.
- **Practical tools accelerate impact:** Standardized templates and playbooks reduce friction and ensure repeatability.
- **Maturity is a journey:** Institutions progress through distinct stages, each requiring specific capabilities and practices before advancing.

5.7.2 Outlook

As Swiss financial institutions continue to experiment with AI, the challenge is no longer *whether* AI can add value, but *how to industrialize it responsibly and at scale*. The framework developed in this study provides the **"how"**: a structured pathway for translating pilots into production-ready systems that deliver measurable business impact.

The next chapter builds on this foundation by presenting **conclusions and strategic recommendations** for the Swiss financial industry. These recommendations highlight what institutions, regulators, and industry associations such as SFTI can do collectively to enable sustainable AI deployment and ensure Switzerland remains at the forefront of innovation in financial services.

6. Conclusions and Recommendations

This final chapter summarizes the key conclusions of the study and provides strategic recommendations for institutions, regulators, and industry associations. Building on the empirical findings (Chapter 3) and the framework (Chapter 5), it highlights the central lessons learned about the **PoC–production gap** and outlines practical pathways for bridging it. The recommendations are designed to guide Swiss financial institutions in achieving sustainable AI impact while ensuring regulatory alignment and trustworthiness.

6.1 Key Conclusions of the Study

The study confirms that Swiss financial institutions are **highly engaged in AI experimentation** but continue to face significant challenges in moving from proof-of-concept (PoC) to full deployment and scaling. The survey and interviews revealed a consistent pattern of enthusiasm tempered by structural barriers.

6.1.1 The PoC–Production Gap

Most AI initiatives remain concentrated in the **ideation and pilot stages**, with fewer advancing into production and only a small share reaching scaling. Compared to 2024, portfolios have shifted upstream: ideation now accounts for around 40% of initiatives, while deployment and scaling have declined to 17% and 11% respectively. This demonstrates both a surge in new explorations—particularly around generative AI—and the persistence of structural barriers that hinder progression.

6.1.2 Survivorship Effect and Value Realization

Despite the bottlenecks, projects that do reach production **deliver solid value**. Survey results showed that **65% of production cases met or exceeded business case expectations**, underscoring a clear survivorship effect: stricter governance, ROI discipline, and compliance requirements mean that only robust projects advance, but those that do are more reliable and impactful.

6.1.3 Structural Barriers as Bottlenecks

The primary obstacles to scaling are **structural rather than cultural**. Respondents consistently identified:

- Data privacy and security constraints.
- Poor data quality and availability.
- Regulatory uncertainty around explainability, bias, and model risk.
- Lack of mature MLOps capabilities and integration with legacy systems.
  Weak or vague business cases also contribute to delays, while cultural resistance and lack of management support were rated as relatively minor barriers.

6.1.4 Success Factors and Enablers

Institutions that succeed in bridging the PoC–production gap consistently leverage a **common set of enablers**:

- Reliable access to high-quality data.
- Executive sponsorship combined with accountable business ownership.
- Compliance-by-design and robust governance frameworks, including AI inventories and model cards.
- Strong IT–business cooperation and clear benefit tracking.
- Dedicated MLOps teams and monitoring frameworks.

These success factors confirm that AI deployment is not primarily a question of algorithms, but of **organizational readiness, governance, and execution discipline**.

6.1.5 *Swiss Specificity*

The study highlights several **Swiss-specific characteristics** of AI adoption:

- **Security and data sovereignty** dominate application and vendor selection, reflecting strict data protection requirements and reputational sensitivity.
- **Regulatory alignment with FINMA and the EU AI Act** shapes how institutions design governance frameworks, particularly in compliance-heavy areas such as AML and KYC.
- Institutions adopt a **quality-over-quantity approach**: fewer projects cross into production, but those that do are more likely to succeed.

6.1.6 *Key Conclusion*

The Swiss financial industry is at a **critical inflection point**. Institutions are heavily engaged in AI exploration, but without systematic investment in data foundations, compliance integration, and production-first design, most initiatives will remain stuck in pilots. Those that succeed demonstrate that AI can deliver measurable impact—but scaling requires holistic readiness across leadership, data, governance, and adoption.

6.2 Strategic Recommendations for Swiss Financial Institutions

The conclusions of this study highlight that bridging the PoC–production gap requires **systematic readiness across multiple dimensions**. For Swiss financial institutions, the following strategic recommendations provide a clear roadmap to move from experimentation to sustainable AI impact (see Table 15):

6.2.1 Interpretation

These recommendations reflect a shift in mindset: **AI is no longer about experimentation, but about disciplined execution**. Institutions that succeed are those that embed compliance and governance, enforce business case realism, and invest in the data and MLOps capabilities that separate pilots from scalable deployments.

6.2.2 Key Implication

By applying these recommendations systematically, Swiss financial institutions can move beyond isolated pilots and establish **AI as a trusted, value-generating capability**, aligned with both business objectives and regulatory expectations.

1. Mandate Production-First Design
▪ Recommendation: Require every PoC to include a defined path to production, with reference architecture, monitoring, and compliance sign-off built in from the outset. ▪ Rationale: Projects without a production-first mindset risk remaining isolated pilots, wasting resources. ▪ Practical Step: Introduce stage-gate reviews (see Chapter 5.4) that explicitly test for production readiness before allocating further investment.
2. Invest in Data and MLOps Foundations
▪ Recommendation: Prioritize investment in centralized data platforms, feature stores, and MLOps pipelines that enable reproducibility, monitoring, and scalability. ▪ Rationale: Data quality and missing MLOps are the most decisive barriers to scaling. ▪ Practical Step: Treat data infrastructure as a strategic asset, not an IT support function; assign dedicated teams with product responsibility.
3. Embed Compliance-by-Design
▪ Recommendation: Involve compliance and risk functions from day one; make model cards, AI inventories, and explainability standards mandatory for all production-ready AI. ▪ Rationale: Regulatory uncertainty is a major blocker, and late-stage compliance reviews stall projects. ▪ Practical Step: Institutionalize compliance-by-design checklists in every PoC to shorten approval cycles and increase trust.
4. Enforce Business Case Discipline
▪ Recommendation: Tie all AI initiatives to measurable KPIs and assign accountable business owners for benefit realization. ▪ Rationale: Weak or vague ROI undermines confidence and slows deployment decisions. ▪ Practical Step: Require business case documentation with conservative ROI baselines at PoC sign-off; track benefits post-deployment.
5. Drive Adoption and Scaling Through Change Management
▪ Recommendation: Engage end users early, communicate transparently, and invest in training to build trust in AI-enabled workflows. ▪ Rationale: Adoption challenges are not the primary barrier, but proactive change management accelerates scaling and secures long-term value. ▪ Practical Step: Develop playbooks and training modules that can be reused across projects, lowering adoption hurdles.
6. Use Portfolio Governance to Balance Exploration and Industrialization
▪ Recommendation: Maintain a live portfolio heatmap of all AI initiatives across the lifecycle (ideation, PoC, deployment, scaling). ▪ Rationale: Without active portfolio management, institutions risk spreading resources too thin across too many pilots. ▪ Practical Step: Cap concurrent PoCs, and selectively fund the top 10–15% of initiatives with credible paths to production.

Table 16: Strategic Recommendations for Swiss Financial Institutions

6.3 Recommendations for Regulators and Supervisory Authorities

While the primary responsibility for bridging the PoC–production gap lies with financial institutions themselves, regulators and supervisory authorities play a decisive role in shaping the environment in which AI can be industrialized. The study findings show that **regulatory uncertainty, explainability requirements, and data protection rules** remain among the most cited barriers. Addressing these challenges through **clarity, proportionality, and dialogue** will accelerate responsible scaling of AI in Switzerland (see Table 16):

6.3.1 Interpretation

The study confirms that regulatory expectations are not merely constraints, but also **potential accelerators** of AI adoption. Institutions that integrate compliance early progress faster; regulators can reinforce this by providing clarity, proportionality, and harmonization.

6.3.2 Key Implication

By clarifying expectations, aligning with international frameworks, and fostering dialogue, regulators can help create an environment where AI adoption in Switzerland is **responsible, scalable, and globally competitive**. This approach strengthens trust in AI while ensuring that Switzerland maintains its reputation as a secure and innovative financial hub.

1. Clarify Supervisory Expectations
▪ Recommendation: Provide detailed and consistent guidance on how supervised institutions should manage AI-related risks. ▪ Rationale: Many projects stall due to uncertainty over what regulators will consider "sufficient" in areas such as explainability, fairness, and monitoring. ▪ Practical Step: Expand on FINMA Guidance 08/2024 with illustrative examples, compliance checklists, and expectations for both traditional and generative AI use cases.
2. Promote Proportionality
▪ Recommendation: Ensure supervisory expectations are risk-based and proportionate to the size and complexity of institutions. ▪ Rationale: Smaller banks and insurers often lack the resources of large incumbents and risk being overburdened by requirements designed for systemically important institutions. ▪ Practical Step: Publish proportional application guidelines, clarifying how expectations differ for large vs. small supervised entities.
3. Support Harmonization with International Standards
▪ Recommendation: Align Swiss supervisory expectations with emerging global frameworks, particularly the EU AI Act. ▪ Rationale: Many Swiss institutions serve EU clients and face extraterritorial obligations; alignment reduces compliance complexity and competitive disadvantages. ▪ Practical Step: Engage with EU regulators to harmonize expectations around model risk management, bias monitoring, and conformity assessments.
4. Foster Dialogue and Knowledge Exchange
▪ Recommendation: Establish regular dialogue formats (e.g., roundtables, workshops) between regulators, financial institutions, and technology providers. ▪ Rationale: Dialogue reduces uncertainty, builds trust, and allows regulators to stay informed about practical challenges in AI deployment. ▪ Practical Step: Create a joint industry-regulator AI working group, hosted by FINMA or in partnership with SFTI, to discuss evolving practices.
5. Encourage Innovation Within Boundaries
▪ Recommendation: Promote controlled environments (e.g., regulatory sandboxes) that allow institutions to experiment with AI under supervision. ▪ Rationale: This enables safe testing of new applications while ensuring that compliance risks are addressed early. ▪ Practical Step: Expand Switzerland's existing innovation programs to explicitly include AI-specific sandboxes with clear entry/exit criteria.

Table 17: Recommendations for Regulators and Supervisory Authorities

6.4 Recommendations for Industry Associations and Ecosystem Players

Industry associations such as **Swiss FinTech Innovations (SFTI)**, the **Swiss Bankers Association (SBA)**, and other ecosystem actors play a vital role in accelerating responsible AI adoption. While individual institutions must build their own capabilities, collective efforts at the ecosystem level can lower costs, reduce duplication, and strengthen Switzerland's positioning as a global leader in financial innovation (see Table 17).

6.4.1 Interpretation

Industry associations can play a **multiplier role**: by providing shared standards, benchmarking, training, and vendor dialogues, they enable institutions to overcome barriers that are too costly or complex to solve alone. At the same time, they help Switzerland establish a distinctive international profile in responsible AI.

6.4.2 Key Implication

Through proactive collaboration, associations such as SFTI and SBA can ensure that Swiss financial institutions are not only individually prepared, but also collectively positioned as leaders in responsible and scalable AI adoption. This **ecosystem-level readiness** is essential for maintaining Switzerland's competitiveness in a rapidly evolving global financial landscape.

1. Provide Shared Tools and Standards

- Recommendation: Develop and distribute standardized instruments such as model card templates, success checklists, and stage-gate criteria.
- Rationale: Shared standards reduce fragmentation, increase comparability, and lower compliance costs across the sector.
- Practical Step: Publish a "Swiss AI Governance Toolkit" for members, aligning with FINMA expectations and international best practices.

2. Facilitate Benchmarking and Peer Learning

- Recommendation: Organize regular surveys, workshops, and peer exchanges to benchmark AI maturity across institutions.
- Rationale: Institutions benefit from understanding where they stand relative to peers and learning from successful deployments.
- Practical Step: Extend the current SFTI survey into an annual benchmarking study, creating a maturity index for Swiss AI adoption.

3. Strengthen Talent Development and Training

- Recommendation: Support the creation of sector-wide training programs to close knowledge gaps in AI, compliance, and data science.
- Rationale: The shortage of AI and MLOps expertise is a recurring barrier, particularly for small and mid-sized banks.
- Practical Step: Develop joint training modules with universities (e.g., OST, ETH, HSG) focusing on compliance-by-design, AI governance, and applied MLOps.

4. Promote Collaboration Between Institutions and Technology Providers

- Recommendation: Act as a bridge between financial institutions and AI solution providers, fostering dialogue and pilot projects.
- Rationale: Many institutions lack the resources to independently evaluate emerging technologies.
- Practical Step: Establish ecosystem partnerships and curated "AI solution showcases" that help members identify vendors aligned with Swiss regulatory standards.

5. Position Switzerland as a Global Leader in Responsible AI

- Recommendation: Engage in international forums and thought leadership to promote Switzerland's strengths in trust, governance, and innovation.
- Rationale: Switzerland has the opportunity to distinguish itself globally as a hub for responsible financial AI, balancing innovation with security and compliance.
- Practical Step: Publish white papers and policy contributions highlighting the Swiss approach to responsible AI adoption, and coordinate with European and global industry associations.

Table 18: Recommendations for Industry Associations and Ecosystem Players

6.5 Outlook

The findings of this study make clear that Switzerland's financial industry is at a **turning point in AI adoption**. Institutions are deeply engaged in experimentation, with portfolios rich in ideation and pilots, yet the majority remain unable to progress beyond PoC into scaled, value-generating deployments. This **PoC–production gap** is not unique to Switzerland, but the country's strict regulatory environment and high standards for trust, security, and sovereignty make bridging it particularly challenging.

At the same time, the study has shown that institutions which successfully move AI into production achieve **measurable benefits**. With 65% of production cases meeting or exceeding expectations, the potential of AI is clear: productivity gains, enhanced compliance, improved user experience, and in some cases, new growth opportunities. The challenge lies not in proving technical feasibility, but in ensuring **institutional readiness**.

Looking forward, three developments will shape the next phase of AI adoption in Swiss financial services:

1. **Institutional Industrialization of AI**
   - Financial institutions must move from pilot projects to **systematic scaling**, treating AI as an enterprise-wide capability.
   - Investment in **data foundations, MLOps, and compliance-by-design frameworks** will determine which institutions can cross the gap.
2. **Regulatory Evolution**
   - FINMA's Guidance 08/2024 and the EU AI Act will increasingly set the standards for governance, explainability, and risk management.
   - Institutions that align early with these expectations will progress faster, while laggards risk being left behind.
3. **Ecosystem Collaboration**
   - Industry associations, regulators, and technology providers must work together to provide **shared tools, standards, and training**.
   - Switzerland has the opportunity to position itself as a **global leader in responsible financial AI**, combining innovation with security and trust.

6.6 Final Reflection

AI is no longer a question of potential, but of execution. The Swiss financial sector must now focus on **bridging the PoC–production gap** by combining strategic leadership, robust foundations, regulatory alignment, and proactive adoption practices. The framework and tools presented in this study provide a **practical roadmap** for doing so.

If institutions, regulators, and associations act in concert, Switzerland can not only unlock the full potential of AI within its financial industry but also reinforce its reputation as a **secure, innovative, and globally competitive financial hub** in the digital era.

7. Executive Conclusion

Bridging the PoC–Production Gap in Swiss Financial Services

The study confirms that Swiss financial institutions are highly active in AI experimentation, but most initiatives remain **stuck in ideation and pilot phases**. Only a minority reach production, and even fewer scale. Yet where projects do advance, outcomes are strong: **65% of production cases meet or exceed business case expectations**.

The challenge is therefore not technical feasibility, but **institutional readiness**. Data quality, compliance, and MLOps remain the decisive bottlenecks — while cultural resistance and leadership buy-in are comparatively minor.

Key Conclusions

Key Conclusions
▪ **PoC–Production Gap:** Portfolios shift upstream (40% ideation; only 17% deployment, 11% scaling).
▪ **Survivorship Effect:** Stricter gates mean fewer, but more reliable, production cases.
▪ **Structural Barriers:** Data, compliance, and legacy integration dominate roadblocks.
▪ **Success Factors:** Data readiness, MLOps maturity, compliance-by-design, executive sponsorship, and IT–business cooperation.
▪ **Swiss Specificity:** Security, sovereignty, and regulatory trust drive vendor and application selection.

Strategic Recommendations

Strategic Recommendations for Institutions
▪ Mandate production-first design: enforce stage-gates, architecture, and monitoring at PoC stage.
▪ Invest in data and MLOps: treat infrastructure as a product, build feature stores and pipelines.
▪ Embed compliance-by-design: model cards, AI inventories, and early risk team involvement.
▪ Enforce business case discipline: tie projects to KPIs, assign accountable owners, kill/scale fast.
▪ Drive adoption and scaling: engage users early, codify playbooks, track benefits transparently.
▪ Portfolio governance: cap concurrent pilots, selectively fund top 10–15% with credible paths to production.

Strategic Recommendations for Regulators
▪ Clarify supervisory expectations with practical examples.
▪ Apply proportionality to avoid overburdening small institutions.
▪ Align Swiss guidance with the EU AI Act to reduce cross-border complexity.
▪ Foster structured dialogue (roundtables, working groups).
▪ Support innovation through AI sandboxes.

Strategic Recommendations for Industry Associations
▪ Provide shared tools and standards (model cards, checklists).
▪ Facilitate benchmarking and annual maturity assessments.
▪ Develop joint training programs with academia.
▪ Broker dialogue with technology providers.
▪ Position Switzerland internationally as a hub for responsible financial AI.

Outlook

Switzerland stands at a **critical inflection point**: AI's potential is clear, but scaling requires systematic readiness. Institutions that invest in foundations, embed compliance, and manage adoption holistically will unlock durable business value. With coordinated effort between banks, regulators, and associations, Switzerland can establish itself as a **global leader in responsible, trustworthy, and innovative financial AI**.

8. Appendix

8.1 List of Tables

8.2 List of Figures

8.3 References

Ankenbrand, T., Bieri, D., Braunwalder, F., Brunner, U., Dietrich, A., Kaelin, G., Kronenberger, T., Ott, D., Reda Kouba, R., Sigrist, F., Viejra, A., & Walser, T. (2023). Data-driven banking in Switzerland: Data drives the next generation of banking. Hochschule Luzern – Institut für Finanzdienstleistungen Zug (IFZ). https://www.hslu.ch/en/lucerne-university-of-applied-sciences-and-arts/research/projects/detail/?pid=6271

Ankenbrand, T., Bieri, D., Gattlen, A. (2025). IFZ FinTech Study 2025 - An Overview of Swiss and Liechtenstein FinTech. Hochschule Luzern – Institut für Finanzdienstleistungen Zug (IFZ). https://hub.hslu.ch/retailbanking/download/ifz-fintech-study/ https://www.hslu.ch/en/lucerne-university-of-applied-sciences-and-arts/research/projects/detail/?pid=3702

Ankenbrand, T., Bieri, D., Reichmuth, L., Stengel, C., Wickihalder, S., & Yilmaz, E. A. (2023). GPT for Financial Advice. . Hochschule Luzern – Institut für Finanzdienstleistungen Zug (IFZ). https://hub.hslu.ch/retailbanking/download/gpt-for-financial-advice/

Artificial Intelligence Act (2024). In Wikipedia. https://en.wikipedia.org/wiki/Artificial_Intelligence_Act

Blattmann, U., Fischer, T., Ettlin, J. (2025). IFZ Studie - Bank-IT und Sourcing 2025. Hochschule Luzern – Institut für Finanzdienstleistungen Zug (IFZ). https://drive.switch.ch/index.php/s/hWr90YJy89we8fF and https://blog.hslu.ch/bankingservices/

Byong-Je Kim et. al. (2025). Byong-Je Kim, Seunghoo Jeong, Bong-Kyung Cho, & Ji-Bum Chung. AI Governance in the Context of the EU AI Act: A Bibliometric and Literature Review Approach. arXiv. https://arxiv.org/abs/2502.03468

Chambers Switzerland. (2025). Artificial Intelligence 2025 – Switzerland: Trends and developments. Chambers Practice Guides. https://practiceguides.chambers.com/practice-guides/artificial-intelligence-2025/switzerland/trends-and-developments

Cheung et. al. (2025). Cheung, D., Lewis, D., et al. Mapping the Regulatory Learning Space for the EU AI Act. arXiv. https://arxiv.org/abs/2503.05787

Consultancy.eu (2025). The EU AI Act: The impact on financial services institutions.”. Consultancy.eu. https://www.consultancy.eu/news/11237/the-eu-ai-act-the-impact-on-financial-services-institutions

Delaram Golpayegani, Isabelle Hupont, Cecilia Panigutti, Harshvardhan J. Pandit, Sven Schade, Declan O’Sullivan, & Dave Lewis. (2024). AI Cards: Towards an Applied Framework for Machine-Readable AI and Risk Documentation Inspired by the EU AI Act. arXiv. https://arxiv.org/abs/2406.18211

Eastwood, B. (2024). Financial services’ deliberate approach to AI. MIT Sloan Ideas Made to Matter. https://mitsloan.mit.edu/ideas-made-to-matter/financial-services-deliberate-approach-to-ai

European Union. (2024). Artificial Intelligence Act (Regulation of the European Parliament and of the Council).

EY (2025). Regulating AI in Financial Services: What It Means for Your Business. EY. https://www.ey.com/en_ch/services/ai/eu-ai-act-readiness

EY. (n.d.). The EU AI Act: What it means for your business. Retrieved 2025, from https://www.ey.com/en_ch/insights/forensic-integrity-services/the-eu-ai-act-what-it-means-for-your-business

FINMA (2023). Artificial intelligence in the Swiss financial market. FINMA Dossier. https://www.finma.ch/en/documentation/dossier/dossier-fintech/kuenstliche-intelligenz-im-schweizer-finanzmarkt-2023/

FINMA (2024). Guidance 08/2024: Governance and risk management when using artificial intelligence. https://www.finma.ch/en/~/media/finma/dokumente/dokumentencenter/myfinma/4dokumentation/finma-aufsichtsmitteilungen/20241218-finma-aufsichtsmitteilung-08-2024.pdf

FINMA. (2025). Press Release: Artificial Intelligence Gaining Traction at Swiss Financial Institutions [Survey results]. Swiss Financial Market Supervisory Authority. https://www.finma.ch/en/~/media/finma/dokumente/dokumentencenter/8news/medienmitteilungen/2025/04/20250424-mm-umfrage-ki.pdf

Gillieron, P. (2025). AI governance in Swiss banks: An operational guide to implementing FINMA Communication 08/2024. Published online. https://www.gillieronavocat.ch/en/2025/05/17/ai-governance-in-swiss-banks-an-operational-guide-to-implementing-finma-communication-08-2024

Grant Thornton Switzerland (2025). The Federal Council wants to ratify the Council of Europe's AI Convention: Finance challenges and implications for Swiss financial service providers. https://www.grantthornton.ch/en/insights/federal-council-ai-convention-finance-challenges/

Härting Rechtsanwälte (2025). AI in the financial sector: FINMA requirements for governance and risk management. Lexology. https://www.lexology.com/library/detail.aspx?g=45840365-3219-40d3-a19a-5e9ae5c0c794

Kuiper et. al. (2021). Kuiper, O., van den Berg, M., van der Burgt, J., & Leijnen, S. Exploring explainable AI in the financial sector: Perspectives of banks and supervisory authorities. arXiv preprint arXiv:2111.02244. https://doi.org/10.48550/arXiv.2111.02244

Kurshan, E., Shen, H., & Chen, J. (2020). Towards self-regulating AI: Challenges and opportunities of AI model governance in financial services. arXiv preprint arXiv:2010.04827. https://doi.org/10.48550/arXiv.2010.04827

Lenz & Staehelin. (2024). FINMA issues guidance on AI use in financial institutions. Lenz & Staehelin Insights. https://www.lenzstaehelin.com/news-and-insights/browse-thought-leadership-insights/insights-detail/finma-issues-guidance-on-ai-use-in-financial-institutions/

Lexology (2025). FINMA guidance on governance and risk management when using AI. Lexology. (2025). https://www.lexology.com/library/detail.aspx?g=b532f639-75b5-4893-918a-72bb6df535c0

Mirishli, S. (2025). Regulating AI in Financial Services: Legal Frameworks and Compliance Challenges. arXiv. https://arxiv.org/abs/2503.14541

MIT NANDA (2025). The GenAI Divide STATE OF AI IN BUSINESS 2025. Preliminary Findings from AI Implementation Research from Project NANDA. https://www.artificialintelligence-news.com/wp-content/uploads/2025/08/ai_report_2025.pdf

Neumann, S., & Wickihalder, S. (2025). A Scalable Framework for Implementing Artificial Intelligence in Swiss Financial Institutions. https://swissfintechinnovations.ch/wp-content/uploads/2024/10/White-Paper_AI-Scalable-Framework-for-Swiss-FI_V1.0-FINAL.pdf

PwC Switzerland (2025). What FINMA's Guidance 08/2024 means for your institution. https://www.pwc.ch/en/insights/fs/ai-in-the-financial-industry.html

Saha, B., Rani, N., & Shukla, S. K. (2025). Generative AI in financial institutions: A global survey of opportunities, threats, and regulation. arXiv preprint arXiv:2504.21574. https://doi.org/10.48550/arXiv.2504.21574

Swiss Federal Council (2023). Mandate to DETEC to develop options for AI regulation in Switzerland (press release/mandate note).

Unit8. (2025). Navigating AI regulation – A guide to AI governance for financial institutions. https://unit8.com/resources/navigating-ai-regulation-a-guide-to-ai-governance-for-financial-institutions/

8.4 Suggested Further Reading

Accenture (2024).
Scaling AI with Trust: How Financial Institutions Can Accelerate Responsible Deployment.
An industry study detailing governance, explainability, and data quality measures in AI scaling programs. Available at: https://www.accenture.com

Bank for International Settlements (2024).
Artificial Intelligence and Machine Learning in Financial Supervision (SupTech and RegTech Trends).
Provides supervisors' perspectives on using AI responsibly and how regulatory expectations are evolving globally.
Available at: https://www.bis.org/topic/fintech.htm?m=3096

Basel Committee on Banking Supervision (2023).
Principles for the Effective Management of Model Risk.
Although not AI-specific, this document is foundational for integrating AI models into regulated banking risk frameworks.
Available at: https://www.bis.org/bcbs

Deloitte (2025).
AI in Financial Services: The Path to Production.
Industry analysis of key drivers, obstacles, and risk controls observed in European financial institutions.
Available at: https://www.deloitte.com

European Commission (2024).
EU Artificial Intelligence Act – Official Text and Explanatory Memorandum.
Provides the final legislative version of the EU AI Act, including risk classifications and compliance obligations for financial services.
Available at: https://artificialintelligenceact.eu

Harvard Business Review (2024).
Turning AI Pilots into Profits: Why Scaling Requires Organizational Redesign.
Examines leadership and operating model changes necessary to move from AI experiments to enterprise impact.
Available at: https://hbr.org

Institute of Financial Services Zug (IFZ), Hochschule Luzern (2024).
Digital Banking and Data Readiness in Switzerland.
Continuation of Ankenbrand et al.'s research, exploring how Swiss banks are adopting data governance frameworks as AI enablers.
Available at: https://www.hslu.ch/en/ifz

MIT Sloan Management Review & Boston Consulting Group (2023).
Building Robust Responsible AI (RAI) Programs as Third-Party AI Tools Proliferate - Findings from the 2023 Responsible AI Global Executive Study and Research Project
A global benchmark report analyzing how leading organizations move from pilots to production. Useful context for comparing Swiss institutions to global peers.
Available at: https://sloanreview.mit.edu/projects/building-robust-rai-programs-as-third-party-ai-tools-proliferate/

OECD (2023).
OECD Framework for the Classification of AI Systems.
Introduces an international taxonomy for AI risk and governance applicable to financial institutions.
Available at: https://oecd.ai

World Economic Forum (2024).
AI Governance Alliance White Paper Series – Responsible AI for Financial Services.
Presents global best practices for embedding ethical and transparent AI in the financial sector.
Available at: https://www.weforum.org/publications