COLUMBIA | SIPA
School of International and Public Affairs

Artificial Intelligence in Banking:

A Strategy for Implementation amidst Legal and Policy Considerations

Columbia University
School of International and Public Affairs
Capstone
April 2025

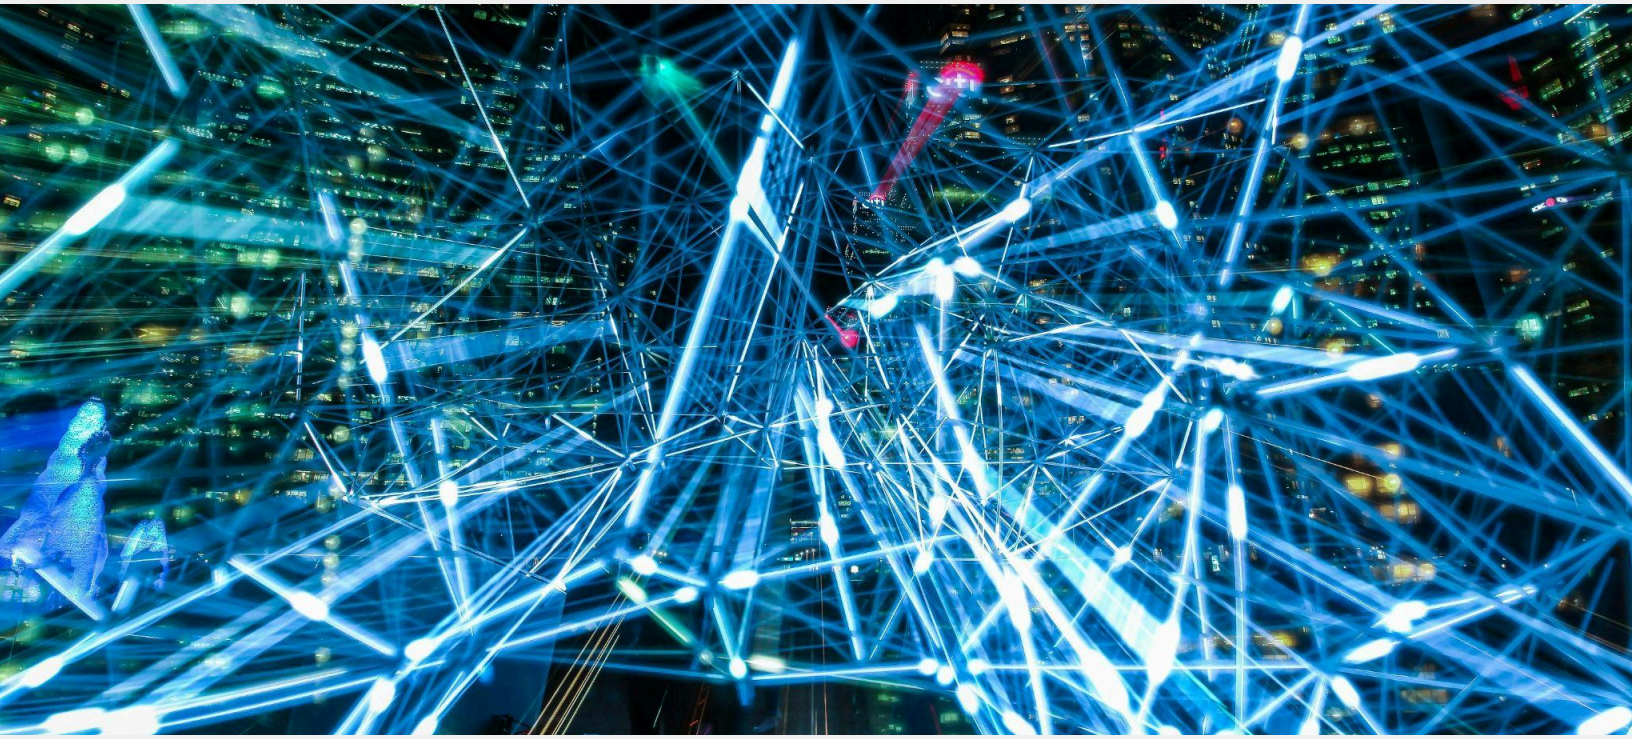

Prepared by

Andrea Vega
Anna Nasibyan
Carina Kaplan
Kennedi Carter
Lia Margarita Te
Louis Sellier
Lu Yu
Zahra Nasidi

Acknowledgements

The Columbia University SIPA Capstone team would like to extend our deepest gratitude to our mentors, interviewees, and advisors who made this project possible.

We would like to thank our wonderful capstone advisor, **Melanie Petsch,** for her advice that guided us throughout this research project. Her feedback challenged us and allowed us to think critically of the issues at hand.

We would also like to thank our client for the opportunity to pursue this project. Our client was engaging, responsive, and provided us invaluable exposure to key executives to give us a multidisciplinary perspective to our topic. Their participation in our project has furthered our understanding of the AI regulatory landscape, the challenges the industry is facing, and directly informed our future outlook on AI policy.

Table of Contents

Executive Summary

As the adoption of artificial intelligence (AI) accelerates, banks must strategically navigate regulatory uncertainty and technological advancement. This capstone project, conducted for a global financial institution, explores the intersection of artificial intelligence, regulation, public policy, and banking strategy to develop recommendations for responsible AI adoption in the financial sector. As AI-driven innovations reshape banking operations, regulatory oversight remains fragmented across federal, state, and international jurisdictions. This capstone project helps our clients navigate these complexities by evaluating multiple avenues the bank can take to leverage the integration of AI.

Through expert interviews, policy analysis, and industry case studies, this paper outlines four key recommendations of pathways for the use of AI within a financial institution for AI-driven innovation:

1. Build consensus to connect with policymakers
2. Align with new regulatory reality: refine internal governance
3. Navigate the AI Trilemma: strategy, development, integration
4. Win with the 3i Strategy: The Innovation Trajectory

This policy document project aims to inform BofA's position in AI-driven financial services, ensuring compliance while fostering trust and innovation.

Introduction

It is crucial for financial institutions to incorporate AI-driven solutions in order to remain competitive in a rapidly evolving artificial intelligence (AI) landscape. As peers and new entrants leverage AI to offer faster, more personalized, and more efficient services, banks that fail to adapt risk losing market share and client trust. AI also enables critical improvements in risk management, customer service, fraud detection, and operational efficiency. From a governance perspective, adopting AI has become a matter of sound business judgment, however this adoption must be done through thoughtful evaluation and adopted only where the technology is appropriate and regulatory risks are low. Failure to do so could weaken a bank's strategic position and expose it to scrutiny from shareholders, regulators, and customers in an industry increasingly shaped by innovation.

However, there are still many risks associated with the uncontrolled development and deployment of AI systems. One major concern is accuracy and reliability. AI systems can fabricate information and confuse correlation with causation, leading to potentially serious errors in financial forecasting, trading, and risk modeling. Without regulatory standards for model validation and testing, institutions may inadvertently introduce systemic risks into the financial system that could also have potential future compliance consequences. Similarly, there are significant risks related to bias and discrimination. AI models trained on skewed datasets can produce unfair or illegal outcomes, such as discriminatory lending decisions or biased investment recommendations. Encouraging thoughtful regulation can help ensure fairness by requiring bias testing, transparency, and responsible data use practices (in addition to internal governance at individual companies). Legal liability is another pressing concern, as institutions could face lawsuits if clients or the public are harmed by AI-driven decisions. Regulatory clarity on liability standards would provide financial institutions with a safer operating framework and help manage reputational risks.

Privacy protection is also a major issue, especially with the use of generative AI. Employees might unknowingly expose sensitive client information by inputting it into AI systems, leading to severe breaches. Regulation can impose clear limits on how financial institutions collect, use, and input client data into AI tools, reducing the likelihood of privacy violations. Additionally, the growing reliance on third-party AI

vendors creates buy-versus-build dilemmas and principal-agent risks. If a vendor's AI system fails, the institution still bears reputational and legal consequences. Regulation can require due diligence, vendor audits, and contractual safeguards to better manage third-party risks.

However, it is important to recognize that despite the risks mentioned, AI adoption offers a multitude of benefits to the financial sector. AI has the potential to revolutionize many aspects of service provision, making banking services faster, more accurate, and far more personalized. Properly deployed, AI can greatly enhance customer experiences, improve financial inclusion, and streamline complex processes such as onboarding, underwriting, and fraud detection. Banks that integrate AI thoughtfully stand to deliver real value to their clients and capture significant efficiency gains, thereby strengthening their competitive position. In an environment where technological innovation is reshaping consumer expectations, the failure to adopt AI could not only diminish a bank's relevance but also cause it to forfeit the opportunity to redefine how financial services are delivered for the better.

Recommendations

The Columbia SIPA capstone group offers four main recommendations for financial services to adopt to leverage AI without taking on compliance or regulatory burdens. These recommendations balance and sustain innovation with responsibility and follow deep research into the federal AI policy environment; state-level AI legislation (enacted, proposed, and failed); financial institution peer analysis; and strategy around integrating AI into existing financial services. Banks should:

1	**Build Consensus to Connect with Policymakers**
2	**Align with new Regulatory Reality: Refine Internal Governance**
3	**Navigate the AI Trilemma: Strategy, Development, Integration**
4	**Win with the 3i Strategy: The Innovation Trajectory**

By leveraging these recommendations, banks can solidify their position as integrators of AI-driven financial services, ensuring compliance while fostering trust and innovation.

01 Build Consensus to Connect with Policymakers

Companies who are interested in adopting AI to enhance their consumer products have been met with the challenging reality of navigating the regulatory and public policy landscape of emerging technology. The past few years have witnessed an advent of regulation and policies attempting to address the risks and harms associated with AI. This moment of time is a crucial moment for banks to seize an opportunity to shape legislative narratives, advocate for sector-specific needs, and ultimately connect with policymakers to help them understand the industry's key AI policy priorities.

Efforts in the Executive Branch:

The United States faces a new approach to its relationship with developers of AI under the Trump administration. President Trump has in part defined the federal government's approach to AI public policy with narratives focused on deregulation. Vice President JD Vance's speech at the AI Action Summit in February set the tone for a pro-innovation public policy approach that limits safety and transparency. Similarly, President Trump's appointed AI czar, David Sacks, has remained silent on the topic, which doubles down on the deregulation and pro-growth policy approach.

President Trump's January 2025 Executive Order (EO) on AI[1] has been the guiding document for AI policy from the executive branch. As the administration's primary overarching AI directive to date, it marks a sharp departure from President Biden's 2023 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.[2] Unlike its predecessor, the Trump 2025 EO positions AI as central to American leadership, economic competitiveness, and national security.[3]

---

[1] "REMOVING BARRIERS TO AMERICAN LEADERSHIP IN ARTIFICIAL INTELLIGENCE," *The White House,* January 23, 2025.
[2] "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," *Federal Register,* A Presidential Document by the Executive Office of the President on 11/01/2023.
[3] "Remarks by the Vice President at the Artificial Intelligence Action Summit in Paris, France," *The American Presidency Project,* February 11, 2025.

> "We believe that excessive regulation of the AI sector could kill a transformative industry... we'll make every effort to encourage pro-growth AI policies. And I'd like to see that deregulatory flavor making its way into a lot of the conversations this conference."
>
> **Vice President Vance,**
> **AI Action Summit, February 2025**

The 2025 EO falls short of proposing any regulation; instead of introducing guardrails or frameworks, it promotes deregulation and urges agencies to rescind or revise existing policies that might impede private sector innovation. While it highlights AI as a central driver of American innovation and leadership on the global stage, it offers no corresponding enforcement mechanisms or substantive regulatory guidance.

There has also been a noticeable shift in AI-related policy since President Trump's first directive on the subject during his initial term. The 2019 EO, Maintaining American Leadership in Artificial Intelligence, took a more measured, government-coordinated approach that prioritized research, security, and responsible governance. The emphasis on safety, security, interoperability, and protection of confidentiality is absent from current policy or public discourse from the executive branch.

In line with the new agenda to achieve AI dominance, the 2025 EO mandates the development of an Artificial Intelligence Action Plan within the first 180 days of the new administration. To support this effort, the administration issued a Request for Information (RFI) through the Office of Science and Technology Policy to seek input from interested public parties.[4] The RFI indicates that the Trump administration is willing to engage with relevant stakeholders, especially in the private sector. Several AI

[4] "Public Comment Invited on Artificial Intelligence Action Plan," *The White House,* February 25, 2025.

developers have already submitted public comments. Banks considering the deployment of third-party AI technologies should closely monitor these policy priorities, as they may influence future regulation and public policy narratives. There is also an opportunity to work with trade associations and ensure the banking sector's needs are included in future RFI's stemming from the executive branch.

Efforts in the Legislative Branch:

This year, the legislative branch, through its 119th Congress, proposed several bills on the use of AI in healthcare, safety, and workforce development. Some of these bills aim to mitigate harms, such as those stemming from digital forgery and price discrimination. While they have garnered bipartisan support, competing priorities, such as budget reconciliation, tax cuts, and debt limits, have stalled progress toward passing a comprehensive federal law to regulate AI in the United States.

Meanwhile, there are key members of Congress that are committed to advancing AI policy through the legislative branch. Congresswoman Maxine Waters' Responsible AI Disclosure Act of 2024[5] is the first and only piece of federal legislation to directly address the intersection of AI and financial services at present. While the Bill does not impose direct regulatory requirements, it directs banking, housing, and mortgage regulators to investigate AI-related risks and opportunities in the sector and report their findings. Although the bill has not been reintroduced, it illustrates Congress's interest in grounding future AI regulation in sector-specific research and stakeholder consultation.

Banks need to build new and existing consensus on sector-priorities with key staffers and members of Congress.

What role do banks play?

After an analysis of current priorities from regulators, this paper has identified five key themes in federal AI policy discussions: 1) data privacy and security, 2) ethical, fair and transparent development, 3) AI innovation and economic competitiveness, 4) national security and strategic AI leadership, and 5) interoperability and global cooperation.

[5] Congress.gov. "H.R.10263 - 118th Congress (2023-2024): Responsible AI Disclosure Act of 2024." November 26, 2024.

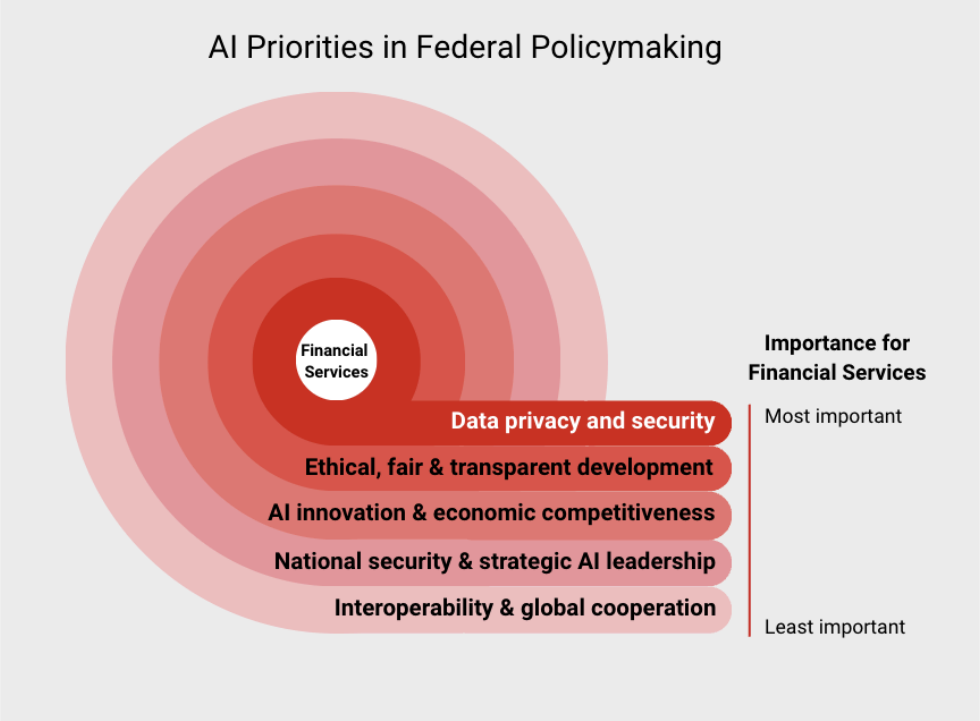

Figure 1: AI Priorities in Federal Policymaking: Key Themes

For the financial services sector, these priorities do not weigh equally. Data privacy and security stands out as the most pertinent issue for financial institutions due to existing compliance obligations, long-term reputational risk, and the responsibility of handling large volumes of sensitive consumer information. At the same time, security, privacy controls, and data governance need to be balanced with usability and innovation. Therefore, financial institutions must ensure any federal regulation is cognizant of usability and applicability.

However, efforts to regulate have been slow despite the public attention drawn to the topic. Beyond the administrations' executive order on AI in January, the executive branch has only released two EOs that focus specifically on the use of AI in classrooms and federal procurement. On April 28, the House passed The Take It Down Act, a bipartisan bill to criminalize deepfake pornography, which has the support of First Lady Melania Trump. That said, AI for use in the financial services industry is not a priority.

This indicates that Congress and the Executive branch are presently in the early stages of developing a united AI policy framework. Policymakers are first working to understand the risks and benefits of AI and identify key stakeholders. This is reflected in

the administration's RFI and limited movement on AI policy on the Hill. As policy takes shape, banks have an opportunity to influence outcomes by engaging strategically with regulators and policymakers. Moving forward, banks need to double down on their public policy approach, and take advantage that policymakers are still in this exploratory stage.

Given their existing relationships with financial regulators, large financial institutions already take several steps to work with policymakers on the use of technology in their services. In light of the emergence of AI, major banks should continue to:

1. Consult think tanks and NGOs to gain insights on risks and opportunities related to pending AI legislation.
2. Invest in dedicated government affairs and public policy teams that advocate independently and through trade associations.
3. Participate in government-led initiatives, such as the National Institute of Standards and Technology (NIST) Risk Management Framework and the EU AI Act. As early contributors to some of these frameworks, banks can position themselves as engaged members of the industry that are committed to finding the right balance between AI regulation and stifling innovation, especially one that recognizes existing financial regulations.
4. Tailor strategies for the U.S. and global engagement accordingly based on varying international standards.

Proactive participation in technical and policy discussions will help ensure that emerging regulations reflect both the complexity of financial services and the existing governance and risk mitigation frameworks of financial institutions. The financial services sector is stringently supervised, and companies in the industry are reckoning with how to define and participate in this new regulatory landscape. The Office of the Comptroller of the Currency (OCC) and other financial regulators have guidelines and expectations for how banks should manage risks associated with AI decision making and third-party technology usage.[6]

Therefore, the sector has an opportunity to collaborate with congressional actors, particularly the bipartisan House Financial Services Committee Working Group on AI.

[6]Matthew Bisanz, Tori K. Shinohara, "Supervisory Expectations for Artificial Intelligence Outlined by US OCC," *Mayer Brown,* May 20, 2022.

The financial services industry has had a collaborative relationship with the House Committee on Financial Service Working Group on AI, helping inform their research and white papers on understanding AI use in financial services applications.

Given existing regulation, financial institutions must foster relationships and build consensus with policymakers to help shape a regulatory environment for banks that supports organizational growth, protects consumers, and avoids one-size-fits-all mandates that overlook the financial sector's unique risks and responsibilities at the intersection of data, technology, and trust. Building multi-stakeholder consensus presents an opportunity to share best practices, identify sector-based risks, and contribute to the design of a regulatory approach that builds on what already works and avoids regulatory duplication. Advocacy should emphasize the importance of clear and consistent definitions as well as the value of leveraging current industry practices.

Continued engagement and participation in the broader regulatory conversation is paramount. While some banks were early participants in the development of the NIST AI Risk Management Framework[7], other avenues for discussion and shaping narratives such as and the Organisation for Economic Co-operation and Development's (OECD) AI principles[8] need to be considered. While neither is binding, both inform how industries define responsible development and deployment and offer structure for financial institutions looking to align governance with emerging expectations. As such, it is also crucial for banks to engage early and often in multi-stakeholder policy discussions and strengthen their participation in forums and platforms such as NIST stakeholder calls, the Cyber Risk Institute, and the Bank Policy Institute.

Working with trade associations in the United States is beneficial for devising a joint narrative, but it oftentimes may be convenient to engage independently abroad. Therefore, participation in these groups allows banks to further think about how to balance its strategies for engagement across jurisdictions.

[7] "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," *National Institute of Standards and Technology,* January 2023.
[8] "OECD AI Principles overview," *The OECD Artificial Intelligence Policy Observatory,* updated May 2024.

02 Align with new Regulatory Reality: Refine Internal Governance

There is currently an uncoordinated approach to federal AI regulation and a pronounced emphasis on competitive measures like domestic innovation, leadership, and economic dominance. While the federal government continues to say AI is a priority, the rhetoric has focused on further partnerships with the private sector rather than federal regulation. On the other hand, the lack of clear regulation also diminishes attention to safety, ethics, responsibility, and transparency. **Therefore, banks should anticipate strict regulation and get ahead in the development or deployment of AI through strong internal AI governance. Without a coordinated government approach, banks need to return to being proactive in their conversations with government, regulators, trade associations, and their peers to ensure they see a coordinated industry response.**

In the absence of federal legislation and leadership, states are left to lead when it comes to passing actual laws on AI in the United States. States are at the forefront of AI policy innovation with strong focuses on data privacy, biases & profiling, transparency, and the applicability of current common law to AI standards. Colorado and California have enacted AI laws emphasizing transparency and accountability in generative AI systems: Colorado's SB24-205 requires developers and deployers of high-risk AI to prevent algorithmic discrimination and provide free AI detection tools with provenance data, while California's AI Transparency Act goes further by mandating comprehensive watermarking requirements, setting a new technical standard for AI-generated content disclosures.[9] These ideas can be leveraged by banks to push forward policy both at the federal and state levels. The passed state legislation should be the benchmark for aligning internal governance policy of AI.

Understanding the following four takeaways from United States state legislation can lead to internal governance success around AI.

Data Privacy

[9] Bryan Cave Leighton Paisner LLP, "AI Litigation Tracker Map," accessed April 28, 2025, https://infographics.bclplaw.marketing/ai-litigation-map/.

AI poses a significant risk to data privacy in financial institutions due to its heavy reliance on vast amounts of personal and sensitive information. To train and operate effectively, AI systems often ingest data such as spending habits, credit histories, location information, and even biometric identifiers. While this enables powerful insights and personalized services, it also increases the potential for misuse, overcollection, and unauthorized access. Financial institutions may unintentionally violate consumer privacy by using data in ways that were not disclosed or consented to, especially when third-party vendors or opaque data pipelines are involved. Additionally, AI systems may be able to infer sensitive characteristics that individuals never voluntarily provided, raising ethical and legal concerns. A data breach or misuse not only puts individuals at risk but can also result in severe reputational damage and regulatory penalties for the financial institution. As financial firms adopt AI at scale, robust governance, transparency, and privacy-by-design principles become critical to managing these risks. States like California are mandating that developers disclose what data has been used to their AI system, while states like Colorado are requiring that the consumer is given the option to opt-out of personal data processing. In practice, this might not be that onerous of a requirement for companies to meet, in fact, it is likely that this will become standard practice similar to making consumers waive their data protection rights in end-license user agreements. It is imperative that financial institutions keep track of these changes.

Common Law

Attorney Generals (AGs) from states like Oregon and California have issued advisory opinions on the applicability of existing common law and statutory remedies to AI. In particular, they have highlighted the Unfair and Deceptive Practices Act, False Advertising Laws, Product Liability, and Negligence as potential existing solutions. While there have not been any cases testing how wide the applicability is going to be, there still is the potential for claimants who have been harmed by AI to go after various financial institutions by making use of these already well litigated remedies, instead of relying on the newer and more unpredictable AI legislation coming from the different states. Additionally, the enforcement of the new AI legislation coming from the states can only be done by AGs, while common law allows for private right of action.

Biases & Profiling

The most extensive laws that already mention AI focus on profiling, biases, and transparency. These considerations need to be at the center of the conversation as banks continue to think about their AI development and use, especially with automated decision-making spurring fears of profiling and discrimination. When AI is employed for internal functions like hiring and technical banking purposes like home loans, biases should be more comprehensively explored.

Further, the lack of mandatory bias auditing requirements for AI systems used in financial services will continue to pose a problem in state-level legislation. While federal laws like the Equal Credit Opportunity Act prohibit discrimination generally, they do not provide a detailed framework for evaluating whether an AI model is discriminatory in practice. At the state level, most jurisdictions do not require financial institutions to conduct or disclose audits that measure bias in their algorithms. As a result, models that systematically under-serve certain populations, such as applicants from marginalized communities, can continue to be used unchecked. This lack of oversight is particularly concerning because biases may not be obvious to the developers or users of the model, and the financial consequences for affected individuals can be long-lasting. Without standardized, periodic auditing, there's a significant risk that AI systems will reinforce existing inequalities in access to credit and financial services and in turn, create reputational risks for a bank.

Transparency

Legal transparency requirements around underlying training and mechanisms for AI models are unclear both at the federal and state levels. For example, one of the most pressing gaps in state-level AI regulation is the absence of clear, consistent standards for how AI can be used in credit decisioning. Financial institutions are increasingly relying on machine learning models to determine loan eligibility, set interest rates, and assess credit risk. However, these models often operate as "black boxes," making decisions based on complex patterns from data that are difficult to interpret, audit, or even understand. State laws are not fully addressing how these algorithms are developed, tested, or deployed, especially when it comes to ensuring fairness. However, a few failed bills demonstrate that this could become a problem of the future. For example, New Jersey S1588 "ensures AI transparency and accountability in recruitment

processes" and addresses potential biases.[10] This kind of thought process could be mirrored in financial services. Without specific guidance, lenders may inadvertently use features that serve as proxies for protected characteristics like race, gender, or age, leading to discriminatory outcomes. The lack of regulation around training data, model features, and performance metrics used in lending decisions leaves room for significant bias and makes it hard for regulators or consumers to challenge unfair outcomes.

Another major regulatory shortcoming is the limited requirement for explainability when AI is used to make financial decisions. In many states, there is no obligation for financial institutions to clearly explain how or why an algorithm denied someone a loan or flagged a transaction as fraudulent. This leaves consumers in the dark about decisions that deeply impact their financial well-being. While the federal Fair Credit Reporting Act and other consumer protection laws require some level of disclosure when adverse actions are taken, these rules were not written with complex AI models in mind. As a result, the explanations provided—if any—are often too vague or technical to be useful. This lack of transparency undermines consumer trust and makes it difficult to appeal incorrect or unfair decisions. A more robust regulatory framework would require institutions to provide clear, understandable explanations of automated decisions and allow consumers to request human review. This would protect financial institutions in the long-run from ambiguous government investigations and future brand risks. While additional transparency requirements could lead to liability for those decisions, it can also give financial institutions more clarity on compliance and ability to protect from these liabilities.

Given the strict state regulations and continued ambiguity at the federal level, banks are having to take into account and anticipate third-party oversight, biases and profiling, and transparency, and determine whether they are a developer or a deployer. For those considering third-party technologies, there is a lower regulatory risk of using large language models, such as Open AI GPT 3.0 given its widespread adoption in financial services. Banks need to double down on their internal AI governance strategy in order to have a compliance approach that anticipates strict regulation but does not jeopardize innovation.

[10] Bryan Cave Leighton Paisner LLP, "AI Litigation Tracker Map," accessed April 28, 2025, https://infographics.bclplaw.marketing/ai-litigation-map/.

Figure 2: Your Internal Governance Strategy

03 Navigate the AI Trilemma: Strategy, Development, Integration

Banks operate in an environment where technology is evolving rapidly, yet they remain subject to strict regulation, which poses unique challenges. The "AI Trilemma" in banking is the challenge of balancing the following three critical decisions. First, banks must choose a strategy based on their risk appetite and ambition for AI adoption. Second, they must decide whether to deploy external AI technologies or develop solutions internally—weighing speed, control, and cost. Finally, banks must integrate these technologies successfully by adapting internal processes to ensure operational effectiveness, regulatory compliance, and long-term scalability. Navigating the trilemma requires careful prioritization, as emphasizing one dimension often comes at the expense of the others. To guide a granular level, system by system decision-making process for AI adoption, banks should use the "AI Trilemma" as a framework:

Choose a Strategy
Based on how aggressive a bank wants to be with the adoption of AI technologies, they should choose an AI identity: **Conservative, Follower or Forefront**.

Develop or deploy
After choosing a strategy, banks need to balance **developing or deploying** AI technologies.

Integrate
Banks need to determine specific internal processes to maximize adoption and integration success.

Choose a Strategy

Banks must first choose a strategy based on their risk appetite and competitive goals.

Conservative: A conservative bank resists adopting AI, relies on traditional methods and uses human decision-making to prioritize stability, trust, and regulatory certainty. Choosing this strategy focuses on prioritizing stability and regulatory certainty and leans into only sustaining technologies. For example, the bank may try one or two tools, with low risk and mundane activities such as sorting documents or detecting

duplicates. Although this approach keeps costs low and minimizes risk, it presents the risk of falling behind the industry and missing out on tools that could make operations more efficient and better at serving customers.

Follower: A follower bank cautiously adopts AI by strategically implementing proven technologies and maintaining competitiveness while minimizing risk. The bank chooses selective adoption, which is to test AI on a small scale and expand successful tools over time. For example, a bank could deploy chatbots to have basic conversations with clients or internal AI tools for compliance reviews. This strategy avoids potential risks but still stays competitive. This approach can avoid high-level risks and build internal experience gradually. However, it takes a longer time to see the result and still requires investments in AI talent. There is also the risk of falling behind more aggressive competitors if pilots do not scale quickly.

Forefront: A forefront bank aggressively invests in, develops, and deploys AI technologies to lead industry innovation and transform operations. This approach positions the bank as a tech-driven financial powerhouse. A forefront bank innovates based on applying emerging technologies. The bank builds an ecosystem to adopt AI for complex decision-making, such as credit underwriting or customer behavior prediction. Although it demands significant resources for hiring AI talent and strong governance, it also improves long-term efficiency gains and presents an opportunity to be ahead of the industry.

Develop or Deploy

Banks must decide on how to source the technologies. The choice is whether to develop AI tools internally or deploy from third party vendors.

This framework highlights the distinct nature in which the bank **innovates**, hires **talent**, designs **transparency** activities, and structures **governance,** based on which sourcing strategy is chosen.

Developing: For banks who build specialized capacity to develop AI capabilities internally, they invest heavily in research, hiring AI talents, and internal R&D patents. This approach allows banks to design technology that meets their specific business and regulatory needs. These development activities are guided primarily by transparent and robust internal frameworks as developed in ***Figure 2***. In the context of AI developments, it would look like explainable AI, which is a broad field focused on making AI predictions

understandable to humans. In-house development supports stronger governance, as AI is embedded into internal compliance processes such as model risk validation and ethics reviews. The bank would need to invest heavily in talent, hiring AI engineers, data scientists, and governance experts, and often provide internal AI training programs, like those offered by BBVA or DBS (see ***Figure 3)***. While this approach gives banks more control and can lead to better long-term results, it comes with high costs, requires significant time and expertise, and demands strong internal oversight.

Deploying: Focusing solely on deployment is reserved for banks who buy from external vendors, like Palantir or SAS, and have a smaller internal team. This approach strongly relies on vendors' transparency disclosure and provides less visibility to how models were built and how they would perform under different conditions. A bank would need to set up a review process and compliance check when using the vendor's tool since it is a challenge to validate fairness or bias in vendor models. Talent needs under this approach are different from developing internally and require less internal hiring. A bank's internal team would focus on building around the vendor technologies and integrating into the bank's system. This approach is suitable especially when it comes to more complex technologies or sensitive use cases.

Integration

Once a bank has the technology, they must determine specific and internal processes needed to maximize adoption and integration success.

The next step for a bank is to determine how to integrate AI tools into daily operations including either by testing and gradually refining them over time or by prioritizing the deployment of incremental features before seeking full board approval.

One way to support integration of AI is by setting up an **"AI Lab for Innovation",** which allows banks to test new and fast moving technologies in a safe environment using synthetic data to reduce risks before scaling AI across the organizations. Another way is to build a tailored internal governance, where banks deploy AI securely on internal servers for sensitive use cases. Strong governance is key to guiding decisions in building and buying future tools. Finally, banks focus on embedding AI into existing AI tools. Banks are integrating generative AI into customer platforms and internal dashboards to improve personalization without significantly altering the user

experience. This step-by-step integration helps banks unlock value while staying cautious and in control.

Deciding on a strategy also requires consideration of the industry landscape to inform a relative position. See ***Figure 3*** below which outlines AI trends in peers in both the U.S. and in Europe.

UNITED STATES

JPMorgan Chase has chosen an aggressive approach to AI adoption. It has the largest AI workforce among its peers, which is larger than the next seven banks combined. The annual report indicates that the bank has made substantial investments in infrastructure, including migrating 75% of its data and 70% of its applications to the cloud.[11] When asked about whether there is a specific area in the business yielding good returns on AI investment, a Chief Data & Analytics Officer role at JPMorgan said "I would turn it around. I can't think of a business where that's not true." [12] The growth rate of employment of AI talent in JP Morgan is roughly 17% every year, which supports challenges in internal AI application.[13]

> "JP Morgan is set to leverage ChatGPT or its proprietary AI tools to enhance operational efficiency and client services."
>
> **Executive Director, JPMorgan AI Research**
> **February 2025**

The bank has also built internal models to support employee productivity, improve decision-making, and enhance customer service.[14] It emphasizes AI integration across all business lines, not only within the tech area, but has reported strong, tangible returns on its AI investments.

Capital One is a top-tier bank recognized for its advanced AI capabilities, especially in compliance and customer service functions. The company has demonstrated a strong commitment to AI research and development. Also, it is actively involved in Fintech Open

[11] JPMorgan Chase & Co., *2023 Annual Report* (New York: JPMorgan Chase, 2024), 18–21. https://www.jpmorganchase.com/content/dam/jpmc/jpmorgan-chase-and-co/investor-relations/documents/annualreport-2023.pdf
[12] Matt Ashare, "JPMorgan Chase Leads Banking Sector in AI Adoption: Report," *CIO Dive*, October 17, 2024, https://www.ciodive.com/news/jpmorgan-chase-capital-one-ai-adoption-leaders-evident/730208/.
[13] ibid.
[14] JPM AI research executive, New York (February 28 2025)

Source Foundation, where it helps shape industry standards and regulatory frameworks. With a focus on solving compliance challenges, such as transaction monitoring and fraud detection, Capital One has become a front-runner in using AI in managing risks.

Goldman Sachs has taken a careful and principle-driven approach to AI by emphasizing responsible AI, explainable AI, and strong governance from model creation to implementation.[15] AI is built into different workflows, such as helping developers work more efficiently, assisting analysts in pulling insights from documents, and generating personalized presentations for clients. Bing Xiang, the Head of AI Research, has highlighted the importance of embedding transparency from the outset to make sure that their AI systems are not only powerful but also trustworthy. This is consistent with responsible AI principles as an internally developed mechanism to govern AI integration in banking.

Citizens Bank highlights the adoption of AI more from the perspective of financial management among mid-sized firms and private equity partners. It actively promotes understanding AI concepts by producing educational resources tailored for finance professionals. The bank also recognizes and has designed mechanisms to address the regulatory and legal challenges, especially in data security and compliance.[16]

BBVA is one of the leading banks generating top assessment criteria for employee AI training. It has launched a Data University in collaboration with OpenAI, aimed at enhancing internal AI education and literacy.[17] This initiative reflects BBVA's strategy of long-term transformation through upskilling and preparing its workforce for deeper AI integration across banking.

EUROPE and SINGAPORE

Barclays and **Deutsche Bank** represent institutions that have engaged with AI primarily for signaling purposes. They have mentioned AI in investor communications demonstrating it is consistent with industry trends. However, compared to more advanced adopters like Capital One or JPMorgan Chase, there is less public evidence of a structured AI strategy, significant investment in infrastructure, or large-scale talent development efforts.

[15] " *The innovation advantage:* how can banks benefit from the most promising AI research?", Gbedemah, Luke (New York: Evident AI Index, 2024)
https://www.youtube.com/watch?v=QLq3-xsdI_Q
[16] Citizens Bank, *2025 AI Trends in Financial Management*
https://www.citizensbank.com/corporate-finance/insights/artificial-intelligence-trends-report-2025.aspx.
[17] https://evidentinsights.com/reports/key-findings-report-2024-public?id=ac66e0dc7e

DBS Bank based in Singapore has successfully transformed from a traditional bank into a tech-led organization. Its AI strategy centers on solving real customer pain points, such as reducing ATM wait times and streamlining account opening processes, through initiatives like the "Managing through Journeys" program.[18] DBS promotes collaboration across teams and has created an internal AI training program to support ongoing learning. The bank's strong customer-centric philosophy is deeply integrated into its AI deployment.

Figure 3: United States and Europe AI Industry and Peer Analysis

[18] McKinsey & Company, *DBS: Transforming a Banking Leader into a Technology Leader*, https://www.mckinsey.com/~/media/mckinsey/business%20functions/mckinsey%20digital/how%20we%20help%20clients/impact%20stories/dbs/dbs_casepdf.pdf.

04 Win with the 3i Strategy: The Innovation Trajectory

The **3i Strategy**, which stands for **"Incremental Innovation Integrator,"** is a proprietary concept crafted specifically for our client, inspired by Clayton Christensen's *The Innovator's Dilemma*. The central idea behind the 3i framework is that success in the AI era does not solely come from chasing disruptive technology. Instead, it is about strategically integrating incremental innovations that continually enhance existing offerings, especially for mainstream consumers. Rather than attempting to immediately leap to cutting-edge transformations, which could alienate customers and overreach capabilities, the 3i approach emphasizes steady progress, reinforcing customer trust, loyalty, and long-term competitive advantage. This strategy combines sustaining innovation with improving current services through experimental, controlled disruption, utilizing internal labs and rapid prototyping to stay ahead of potential market shifts.

Sustaining Innovation vs. Disruptive Technology

Sustaining Innovation focuses on refining and enhancing the bank's existing services and offerings. This includes embedding AI in customer-facing systems, improving operational efficiency, and augmenting personalization in a way that feels seamless and familiar to users. Banks should also recognize that disruptive technology and innovation can change market dynamics. With this in mind, the bank should maintain strategic scouting that observes emerging technologies using an "Internal AI Lab" rather than risking destabilizing the core business. Firms that master sustaining innovation climb the performance curve thoughtfully, defend their position against disruption, and maintain customer satisfaction without unnecessary risk.

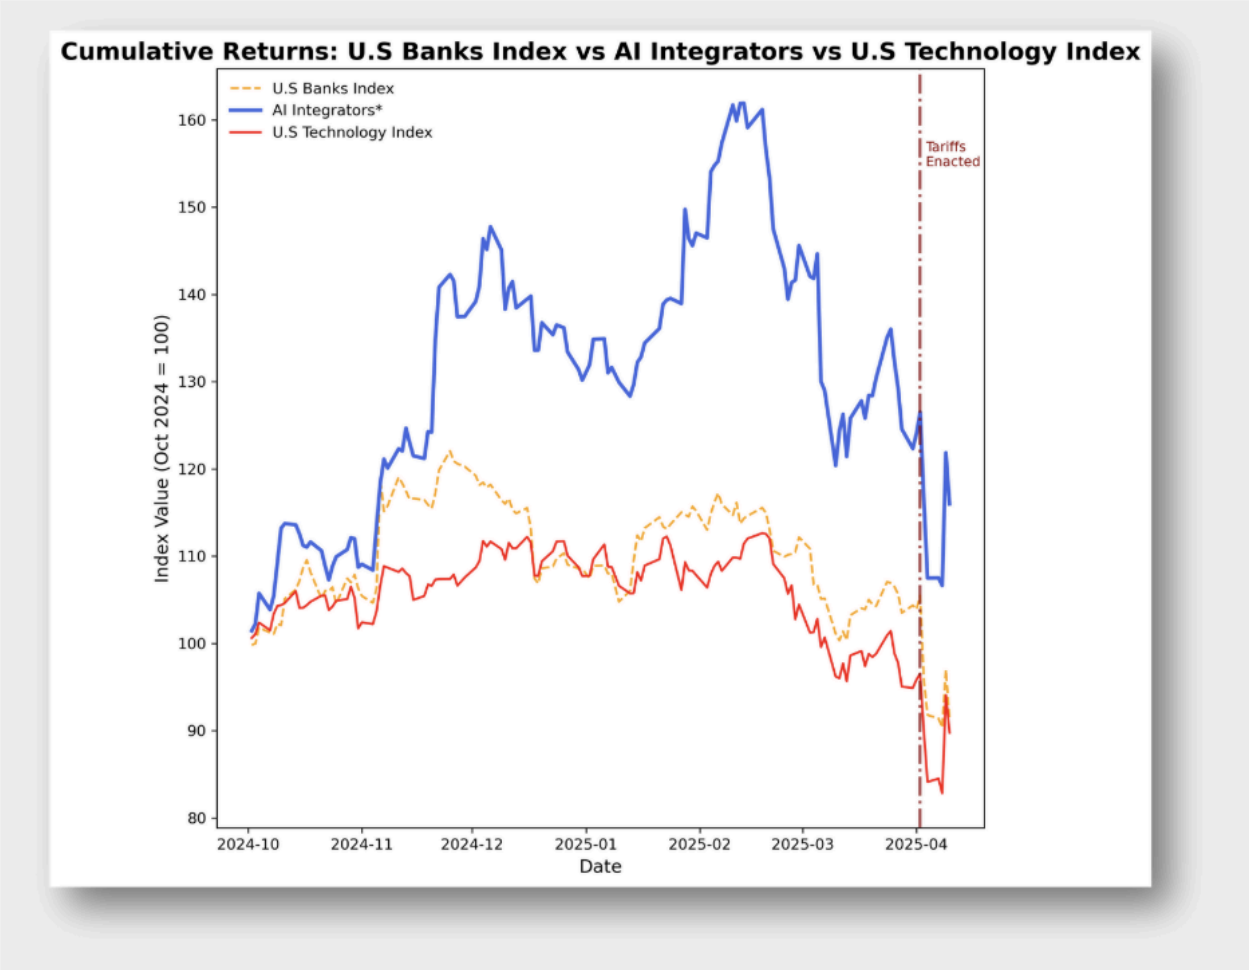

Figure 4: Cumulative Returns: U.S. Banks Index vs AI Integrators vs U.S. Technology Index (Source: SIPA Capstone Team, cf. Technical Appendix)

The AI Integrator Advantage

Building on Goldman Sachs's concept of "AI Integrators," Figure 4 shows that financial institutions that adopt AI actively and strategically are rewarded with a market premium. The benchmarking methodology compares market performance of the U.S. Banks Index, an equally weighted portfolio of AI Integrators, and the U.S. Tech Index. We use the U.S. Banks Index as a baseline, comparing its cumulative return against both AI Integrators and the U.S. Tech Index. This framework helps isolate the incremental "AI Integrator premium" that a bank could command by embracing AI technologies, beyond what's priced into the broader banking or tech sectors. The data shows that AI integrators consistently outperformed both traditional banks and even broader tech

companies. This shows that embracing AI systematically leads to stronger stock performance and investor confidence. If the bank positions itself as an AI-forward institution through strategic and visible AI adoption, it has the ability to unlock an AI premium, boosting its value and competitiveness.

The 3i Model in Action

The 3i Model comes to life through a dual-track approach: steadily integrating incremental, value-added AI into existing offerings while maintaining a proactive, experimental pipeline for future disruptive technologies. Inspired by Apple's seamless integration of generative AI into its ecosystem through Apple Intelligence, the bank should aim to integrate smart capabilities, such as predictive prompts, proactive alerts, and personalized insights, directly into its mobile apps and online platforms.[19] These enhancements should be introduced without disrupting the familiar user experience, ensuring that customers feel empowered rather than overwhelmed by new technology.

While incremental innovation sustains current strengths, the bank must also maintain a future-facing innovation pipeline to stay ahead of industry shifts. This requires establishing a dedicated "Internal AI Lab," serving as an "Innovation Watchtower." Within this lab, a specialized intrapreneurial team would utilize rapid build–measure–learn cycles to develop prototypes and conduct controlled experiments, often employing GAN-driven synthetic data to safeguard real customer information during testing. This setup enables the early identification of promising technologies and weeds out ineffective solutions without exposing the broader organization to unnecessary risk. Continuous, proactive tech scouting should be a priority, focusing on emerging advances such as foundation models, federated learning, and agentic AI. Strategic evaluation mechanisms will empower leadership to decide whether to develop solutions internally or deploy external products based on real-world testing data, ensuring that only the most effective and responsible innovations are scaled organization-wide. Through the 3i model, the bank will strike a balance between safe, continuous improvement and strategic readiness for disruption.

Internal AI Lab: The Disruptive Risk Hedge

[19] "Apple Intelligence." *Apple*, www.apple.com/apple-intelligence/. Accessed 29 Apr. 2025. https://www.apple.com/apple-intelligence/

The bank should establish an “Internal AI Lab” that monitors emerging technologies, analyzes market threats, and conducts small-scale experiments. This would help hedge against the risk of disruption by identifying disruptive competition sooner, rather than waiting until disruptive technologies are fully mainstream, thereby avoiding the challenge of catching up too late. The structure of this Internal AI Lab could have a lightweight, fast, and experimental identity, allowing it to pivot quickly based on its findings.

The AI revolution in banking is not about chasing the flashiest technology. It is about building a resilient, flexible, and customer-centered innovation pipeline that integrates improvements customers actually need and want today, while preparing stealthily for the disruptions of tomorrow. By combining incremental integration and strategic disruption readiness through the 3i Strategy, the bank can lead the next era of financial services with trust, stability, and intelligent growth.

The 3i Strategy enables the bank to lead in the AI era by striking a balance between innovation and intention. Rather than choosing between stability and disruption, this approach integrates both, enhancing today’s services with trusted AI capabilities while preparing for tomorrow’s breakthroughs through structured experimentation. By embedding AI incrementally and building a dedicated pipeline for disruptive technologies, the bank positions itself to earn customer trust, regulatory confidence, and long-term market rewards. In an environment defined by rapid change and rising expectations, the 3i Strategy ensures the bank remains not only compliant and competitive but also continuously evolving and future-ready.

Conclusion

In a rapidly evolving landscape of AI and regulations, ambiguity is not just a challenge, but an opportunity to meet the moment and integrate exciting uses of technology. Rather than waiting for top-down mandates that may overlook the sector's nuances, banks can proactively shape AI policy and internal governance frameworks that balance innovation with integrity.

This document presented a framework that acknowledges regulatory ambiguity not as a risk to avoid, but as a frontier to align with the bank's technology priorities. By building strategic relationships with policymakers, navigating the AI Trilemma with clarity, and embracing the 3i Strategy for innovation, the bank can secure a competitive edge while upholding trust and compliance.

Central to this strategy is the Internal AI Lab, a vital mechanism to convert uncertainty into insight and tailored to this global financial institution. As both a testbed and a monitoring mechanism, the lab will enable the bank to safely experiment with AI technologies, validate use cases using synthetic data, and stay ahead of disruptive trends before they impact the broader industry. It ensures that the client doesn't just respond to change, but actively drives it. Ultimately, the future of AI in banking will be dictated not solely by regulation, but by the foresight and discipline of institutions willing to innovate responsibly. With the tools, principles, and structure outlined in this paper, this global financial institution is poised to integrate AI with clarity, purpose, and resilience.

Technical Appendix

Selection of AI Integrators

We identified eight firms—Cloudflare, MongoDB, Nutanix, Autodesk, Datadog, Elastic, Snowflake and Palantir—whose **bottom-line margins** are predominantly driven by the deployment of AI technologies. These companies were chosen because a majority of their incremental gross profit can be directly traced to AI-enabled products or services, making them a reliable proxy for how effectively management teams convert AI investment into revenue and profitability. To qualify, each firm must publicly disclose its AI initiatives in quarterly filings and demonstrate consistent year-over-year margin expansion attributable to AI rollout.

Portfolio Construction

The AI Integrators portfolio is built as an equally-weighted basket of the eight names—Cloudflare, MongoDB, Nutanix, Autodesk, Datadog, Elastic, Snowflake and Palantir—with **daily rebalancing** to ensure each stock contributes exactly 1/8 of the portfolio return every trading day. For benchmarks, we use **KBE** (SPDR S&P Bank ETF) as a proxy for the U.S. Banks Index and **IYW** (iShares U.S. Technology ETF) as a proxy for the U.S. Technology Index, each held at a constant 100% weight without any internal drift.

Data & Normalization

All price data were pulled programmatically via the **Yahoo Finance API** (using the *yfinance* Python package). We focus exclusively on **price returns** (dividends excluded) and use daily closing prices to compute returns. To facilitate comparison, each series is **normalized to 100** on **October 1, 2024**, and cumulative returns are plotted from that date through the present.

Benchmarking Methodology

At its core, our analysis treats traditional banking performance as the baseline against which we gauge the true impact of AI adoption. By tracking how an equally-weighted

basket of AI-focused firms diverges from the U.S. Banks Index—and even the broader technology sector—we reveal the extra value that comes purely from executing on AI strategies, rather than from general market or tech-driven tailwinds. This "**AI Integrator premium**" shines a light on the benefit of leaning into AI as a transformative growth lever, isolating the gains that stem directly from intelligent automation and data-driven innovation.

Tools & Code Snippet

Analysis was performed in Python 3.10 with the *pandas*, *numpy*, *yfinance*, and *matplotlib* libraries. A representative snippet of our workflow is shown below:

```
import yfinance as yf

ai_integrator_list = ['NET', 'MDB', 'NTNX', 'ADSK', 'DDOG', 'ESTC', 'SNOW', 'PLTR']
tickers = ['KBE', 'IYW'] + ai_integrator_list

# Download daily closing prices
data = yf.download(tickers, start='2024-10-01', end=today, adjusted=False)['Close']

# Compute daily returns
returns = data.pct_change().fillna(0)

# Build equally-weighted AI Integrators portfolio
ai_portfolio = (returns[ai_integrator_list].mean(axis=1).add(1).cumprod()) * 100

# Normalize benchmarks
kbe_index = returns['KBE'].add(1).cumprod() * 100
iyw_index = returns['IYW'].add(1).cumprod() * 100
```