



FINANCIAL
STABILITY
BOARD

Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Consultation report

Response to Consultation

Mastercard

1. Do you agree with the benefits and risks of AI adoption by financial institutions described in this report? Are there any substantive benefits or risks not covered?

At Mastercard, we harness artificial intelligence (AI) every day across our ecosystem, using it to help verify identity, prevent fraud, and provide actionable cyber threat intelligence to our customers. To ensure our use of this technology is grounded in trust, we are committed to its responsible development and use. As such, Mastercard welcomes the Financial Stability Board's (FSB) efforts to develop the Sound Practices for Responsible Adoption of Artificial Intelligence ("Sound Practices") publication. The development of widely recognized standards and guidance for the adoption of AI by financial institutions will help ensure the benefits of AI adoption are experienced widely by consumers throughout the financial ecosystem while limiting exposure to the risks presented by the technology.

We have focused our comments on the Sound Practices publication in response to the first two questions posed by the FSB: (1) whether there should be any changes to the identified benefits and risks of AI adoption for financial institutions and (2) whether there should be any modifications to the recommended best practices that financial institutions can adopt to harness the benefits and protect against the risks of increased AI adoption.

Mastercard's key recommendations for the Sound Practices publication include the following:

(1) Provide additional description of key benefits of AI adoption for financial institutions focused on how AI can improve payments ecosystem integrity, network intelligence and operational resilience, as well as access to enterprise-grade security technology.

(2) Include more detail concerning important risks to financial institutions of increased AI adoption, including system concentration and agentic AI risks; and

(3) Finally, we believe the best practices identified in the publication would benefit from additional description concerning three key areas: (1) security best practices for AI agent identity; (2) oversight; and (3) the importance of leveraging threat intelligence to respond to AI security risks.

We believe the Sound Practices publication is a thoughtful resource for financial institutions looking to better understand the benefits and risks of AI adoption as well as key best practices to protect their institutions and customers from potential risks. Mastercard appreciates the opportunity to provide feedback on this publication.

Benefits and Risks of AI Adoption for Financial Institutions

In the first section of the Sound Practices publication, “[a]n overview of the use of AI in the financial system”, the FSB describes examples of the benefits and risks to financial institutions looking to increasingly adopt and incorporate AI into their operations and systems. Mastercard agrees with the articulation of several of these examples, including how AI can be used to improve fraud detection (p. 8-9), helping to identify synthetic identities and money laundering as well as enhancing operational resilience through improvements in cyber vulnerability detection (p. 12). However, we do believe that the Sound Practices publication could be improved by including more details concerning these issues.

Benefits

With respect to the benefits of AI adoption, we believe that the Sound Practices publication could be refined and improved by including further details, focused not only on benefits to individual financial institutions, but also system-wide benefits of AI adoption, such as improvements to the operation, security, and resilience of the financial payments ecosystem. Specifically, we have identified three areas of potential benefits of AI adoption that could be included in the Sound Practices publication: (1) payments ecosystem integrity; (2) network intelligence and operational resilience; and (3) democratizing access to enterprise-grade security technology.

Payments ecosystem integrity

The ability of AI to process large volumes of data while increasing the accuracy and speed of the interpretation and decision-making applied to that data will improve payments ecosystem integrity by: (1) increasing the consistent application of payment network rules, such as chargeback deadlines and compliance requirements; (2) enhancing the accuracy and timeliness of the dispute resolution process, including engaging with cardholders and resolving disputes; and (3) more quickly and accurately identifying emerging ecosystem risks, such as network-level money laundering and counterfeit payment card threats. This also implies that the integrity of ecosystem data that is being generated and ingested by AI is essential to the quality of outcomes.

Network intelligence and operational resilience

In addition to increasing ecosystem integrity, increased AI adoption by financial institutions can help enhance the security and operational resilience capabilities of payments networks, including: (1) enhancing cross-network pattern detection, such as improving the speed and accuracy of cross-border signals identifying potential large-scale global fraud events; (2) earlier detection of systemic risks, including empowering threat intelligence and AI-enabled authorization system testing to more quickly and accurately identify weaknesses in payment authorization systems; and (3) dramatically improving operational continuity and transaction decisioning by, among other things, leveraging AI models to better predict financial

institution transaction decisioning patterns in the event of potential network outages or downtime.

Democratizing access to enterprise-grade security technology

Finally, we believe increased AI adoption will enable small and medium-sized financial institutions to better use and scale sophisticated cybersecurity, anti-money laundering, and fraud prevention technology. Smaller financial institutions often confront significant resource and technical limitations in accessing, adopting, and deploying cutting-edge security technology. Increased AI adoption for smaller financial institutions can help them leverage tools like anomaly detection and behavioral analysis to assist in identifying, in real time and at significant scale, suspicious transactions that fall outside of usual and expected payment patterns. Leveling the playing field for access to state-of-the-art security technology will also help improve the overall security landscape for all participants in the broader payments ecosystem.

Risks

The Sound Practices publication identifies several important risks that financial institutions should consider when planning their adoption of AI, including “ineffective data governance and poor data quality,” “inadequate or inappropriate human oversight,” “cyber and ICT risks,” and “agentic AI risks.” While a significant number of the key risks of AI adoption have been identified in this publication, Mastercard believes additional risk categories, such as “AI homogeneity and system concentration risk” as well as supplemental detail relating to “agentic AI risks” would yield a more comprehensive picture of the risks of AI adoption by financial institutions.

AI homogeneity and system concentration risk

As noted above, the Sound Practices publication indicates how financial institutions can leverage AI to improve fraud detection and prevention. But, depending on how we build this technology, we could undermine its effectiveness and the security of the financial ecosystem. For example, if participants in the financial ecosystem build or use similar fraud detection AI models – or do not consistently improve such models, it will become easier for bad actors to devise solutions that take advantage of predictable computational blind spots, which could, in turn, leave a broad group of financial institutions susceptible to the same systemic fraud detection weaknesses. Network resilience will depend upon diversity of controls, decision-making, and technology.

Agentic AI risks

As AI agents begin to work across different business and organizational boundaries, such as shopping agents, business supply chain agents, as well as payment and contract execution, financial institutions will need to implement robust controls to manage security, resilience, and other operational risks.

AI agent systems used in agentic payments, for example, exhibit security risks that differ substantially from traditional software due to their autonomy, delegated authority, and dynamic prompt driven behavior. These systems are uniquely susceptible to prompt

injection, misconfiguration, and unauthorized tool execution, where adversarial instructions or improperly scoped permissions can cause agents to initiate unintended financial actions, modify trust related data, or circumvent intended constraints. Internal assessments of agent misconfiguration and prompt-based exploits highlight how hidden instructions or excessive permissions can lead agents to take alternative or unauthorized actions, with potential for material negative outcomes for individuals or organizations, including large-scale fraud and the theft or unintended disclosure of commercially sensitive or personal data.

Moreover, the concentration of sensitive identity, agentic trust (i.e., the confidence that an autonomous AI agent is authorized to act, is acting with legitimate intent, and can be held accountable on behalf of a real human within defined guardrails) and payment workflows within autonomous agents creates elevated exposure to agent hijacking and privilege escalation. As a result, securing agent payment systems requires dedicated controls including strict permissions, continuous monitoring, deterministic fallback mechanisms, and robust guardrails to ensure agent behavior remains aligned. The incredible pace of innovation in agentic AI today means that many of the challenges identified above are actively being addressed and mitigated. As these types of systems mature, it is likely that the industry's capacity to manage these risks will also improve substantially.

2. Are the sound practices sufficiently comprehensive and clear to enable financial institutions' responsible AI adoption?

Sound Practices for AI Adoption for Financial Institutions

The Sound Practices publication provides useful detail concerning a series of best practices that financial institutions should consider when adopting AI focused on "organisation-wide AI governance" and "AI lifecycle management." While we agree with the identified best practices, including the criticality of balancing "human oversight" with automated technology, we believe the Sound Practices publication should include additional detail focused on three categories: (1) AI agent identity and authentication; (2) oversight; and (3) the importance of threat intelligence adoption to protect financial institutions against AI security threats.

AI agent identity and authentication

The Sound Practices publication includes recommendations concerning the importance of implementing controls and security measures surrounding the adoption and use of AI agents. However, we believe there should be a clear emphasis on the importance of implementing best practices relating to AI agent identity and authentication as verifiable identity and authentication technology will be essential in creating the secure, reliable, resilience "on behalf of" authorization architecture that is inherent to a financial ecosystem in which AI agents are significant participants.

While Mastercard supports the Sound Practices publication's recommendation that financial institutions consider "assigning and documenting individual identifiers to AI agents to facilitate monitoring and identity management," we believe the FSB should go further and include more detail concerning AI agent identity and authentication, particularly where agents are expected to operate across or outside of company boundaries. We would suggest including language encouraging a "know your agent" approach to AI agent security, including references to agent registration, measures for cryptographically verifiable identity,

provenance, and traceability. We would also suggest including language strengthening guidance on agent credentials, noting that such credentials should be: (1) cryptographically verifiable; (2) tied to a human or enterprise principal; (3) fully lifecycle managed (from issuance to revocation); and (4) policy-constrained based on the scope of authority for that unique agent.

By including this additional detail, the FSB can ensure that the Sound Practices publication has more information for financial institutions concerning best practices for AI agent identity, authentication, and security.

Oversight

The ability to monitor, understand, and – where necessary – intervene, in the actions of AI is essential to building trust in a financial ecosystem in which AI is key to driving value and protecting consumers against evolving security threats. Meaningful human oversight of AI is vital for organizations maintaining the necessary visibility and control over AI – including AI agents – ensuring the proper foundation for accountability.

Even so, Mastercard believes it is critical to conceive of oversight as a multi-layered tapestry of techniques, some necessitating direct human involvement and others focused on leveraging automated monitoring mechanisms. As such, we support the manner in which the Sound Practices publication articulates “appropriate forms of human oversight” as including not only “human-in-the-loop” concepts but also alternative forms of oversight that depend on incorporating technology, such as “AI-in-the-loop” (AI as a supportive layer for decision-making generating alerts for potential human review) and “human-on-the-loop” (human intervention periodically depending on the triggering of alerts).

For financial institutions, a flexible approach to AI oversight is especially important with the expected growth of agentic commerce – where AI agents will be purchasing goods and services on behalf of users – as it may not be scalable and logistically feasible for organizations to implement human review of each individual purchasing decision. We believe the FSB could improve the Sound Practices publication by including more detail concerning the importance of balancing human oversight with automated techniques in the context of agentic commerce.

Threat intelligence adoption to protect financial institutions against AI security threats

The development and deployment of increasingly advanced AI models and technology is yielding several different AI-related security challenges for financial institutions, including: (1) the rapid discovery of a vast and growing number of security vulnerabilities in commonly used enterprise software and technology; (2) security risks unique to AI-powered technology, such as “data poisoning” and “prompt injection”; and (3) AI-powered cybercrime, financial crime, and fraud.

The key for financial institutions bolstering their security programs in the face of these challenges is adopting an intelligence-driven strategy that uses cutting-edge threat intelligence capable of quickly identifying, prioritizing, and remediating AI-related security threats. We support the Sound Practices publication’s acknowledgement of the importance of AI-powered threat intelligence and vulnerability discovery (pages 45-46) but believe this

security measure should be a more prominent focus, standing alone as a key security best practice for financial institutions looking to protect their systems, assets, and customers from a wide range of AI-related security threats.

The Sound Practices publication should include more detail concerning the different use cases where threat intelligence can help protect financial institutions and their customers from AI-related security threats. The importance of threat intelligence goes beyond vulnerability discovery. The Sound Practices publication should also reference the importance of organizations leveraging attack surface intelligence to better understand not only potential vulnerabilities but also the systems and assets that may be vulnerable, helping prioritize addressing those weaknesses. In addition, financial institutions should be looking to use intelligence to address financial crime threats, including the identification of synthetic identities and identity theft, fraudulent merchants and businesses, large scale phishing campaigns, money laundering networks, and large-scale scam operations.

- 3. Do the sound practices strike an appropriate balance between managing risks relating to all forms of AI, and addressing some of the risks relating to emerging and new complex forms of AI, such as GenAI and agentic AI?**
- 4. Are the sound practices sufficiently flexible to accommodate and address newer types of AI and responsible adoption over time?**
- 5. Do the case studies in this report sufficiently highlight how different types of financial institutions can benefit from responsible AI adoption? Are there additional case studies for inclusion in the report? If so, please provide such case studies, particularly for nonbanks.**
- 6. Do the case studies in this report provide actionable insights for financial institutions in their responsible AI adoption?**
- 7. Are the definitions in the glossary clear and aligned with industry sound practices, including recent developments in AI?**
- 8. Are there any comments you would like to make on this report? Please fill in.**