FSB FINANCIAL STABILITY BOARD

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

Response to Consultation

Finspector

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?

We agree with the benefits and risks as described. The analysis is comprehensive and the case studies are credible. We would add a few observations from the conduct and consumer protection side of the picture, which the final report could draw out more fully.

First, the report treats AI as both a threat and a defence in the cyber domain, and rightly so. The same duality applies to market conduct and consumer protection, but here the report leans towards treating AI mainly as a source of risk. In practice AI is now central to the defence as well. The volume and sophistication of misleading and fraudulent communications aimed at consumers has risen sharply, and much of it is AI assisted. UK Finance reported that criminals stole £1.28 billion through payment fraud in 2025, with investment fraud, which often begins with promotions on social media, the single largest category of authorised push payment losses at £221.5 million, up 40 per cent in a year. Around two thirds of authorised push payment fraud cases now start online. The FCA has also led repeated international actions against illegal ‘finfluencer’ promotions through 2025 and 2026. At the same time, financial promotions compliance has become a significant workload. Following FCA intervention, authorised firms amended or withdrew 19,766 promotions in 2024, almost double the figure for the year before. Defending against activity at this scale cannot be purely manual. We would encourage the FSB to recognise, symmetrically with its treatment of cyber, that responsible AI adoption is part of how firms meet their conduct and consumer protection obligations, not only a source of risk to them.

Second, and related, there is a distinct risk worth naming directly: the widening gap between the speed at which AI can generate communications and the largely manual processes many firms still use to review them. When promotional content can be produced and distributed in seconds but is checked by people working through queues, controls fall behind, non-compliant material reaches consumers, and firms and supervisors lose visibility. That is a consumer protection risk in its own right, and one that responsible AI adoption is well suited to address. The report’s risk analysis would be stronger for acknowledging it.

Third, the report’s treatment of shadow AI is good. We would simply note that marketing and front-office functions are a particular exposure. Staff under pressure to produce content quickly are among the most likely to reach for general-purpose generative tools without compliance-grade controls, which can introduce inaccurate, misleading or non-compliant claims into customer-facing material. Naming this pattern would help firms direct their controls where the risk is highest.

Sources:

UK Finance, Annual Fraud Report 2026 (2025 data), 15 June 2026. ukfinance.org.uk/news-and-insight/press-release/fraud-report-2026-press-release

FCA press releases on illegal ‘finfluencer’ promotions, including the international week of action of June 2025 and further enforcement in 2026. fca.org.uk/news

FCA, Financial promotions data 2024, published 7 February 2025: a 97.5 per cent increase on the 10,008 promotions amended or withdrawn in 2023. fca.org.uk/data/financial-promotions-data-2024

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

Broadly, yes. The twelve practices are coherent, the lifecycle structure is logical, and proportionality is applied consistently. Two observations from the position of a third-party provider.

The interaction between Sound Practice 8 (explainability and transparency) and Sound Practice 12 (third-party risk) is where much of the practical difficulty sits, because a large and growing share of AI is adopted through vendors rather than built in-house. The Bank of England and the FCA found that around a third of AI use cases in UK financial services are now third-party implementations, up from roughly a sixth two years earlier. A deploying firm stays accountable for outcomes it cannot fully see inside. The report makes this point well, including the observation that current transparency tools are generally not tailored to the financial sector. We would go a step further and suggest the final report state plainly that deployers should expect, and contract for, the specific evidence they need to discharge their own obligations: validation against the relevant regulatory requirements, audit and logging access, documented limitations, and clearly defined human-oversight hooks. Information asymmetry can be most acute for smaller firms, which lack the contractual leverage of larger banks to negotiate specific AI due diligence rights. Generic model or system cards are a starting point, not a sufficient one. We support the report’s encouragement of sector-specific assurance artefacts and shared-responsibility models, and as a provider we would welcome clearer expectations here, since they raise the floor for everyone.

On Sound Practice 10, the distinction between meaningful oversight and tick-box oversight is the most important idea in the report for our part of the market, and it is well drawn. In financial promotions work the pattern that works is simple: AI triages and drafts, a qualified person decides, and every decision leaves an audit trail. The report’s framing supports exactly this, and we would not change it.

Sources:

Bank of England and FCA, Artificial intelligence in UK financial services 2024 (third survey), November 2024: 75 per cent of responding firms reported using AI. bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024

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?

The balance is about right. The report avoids the common mistake of over-indexing on speculative agentic scenarios while still flagging them honestly, and it keeps proportionality at the centre. We would offer one caution. For conduct and customer-facing use cases, the most pressing near-term risk is not autonomous agents but generative AI producing content at scale, including promotions, summaries and advice-like outputs. That risk is present now, not on the horizon. As attention shifts towards agentic AI, we would encourage the FSB to keep appropriate weight on generative content risks, which are where most consumer harm is likely to arise over the next few years.

4. Are the sound practices sufficiently flexible to accommodate and address newer types of AI and responsible adoption over time?

Yes. The technology-neutral, outcomes-based, menu structure is the right design for durability, and it is the reason the report should stay useful as the technology changes. The decision to reference evolving external standards and knowledge bases, rather than freeze technical detail into the text, is sensible. We would encourage the FSB to commit to a regular refresh cycle and to keep those external references as living links, so the report keeps pace without needing to be rewritten. The public-private mechanisms the report cites, including the Bank of England and FCA AI Consortium and the FCA's AI Lab, are a good source of the signals that should drive those updates.

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.

The case studies are one of the strongest features of the report. They are specific and credible. As the report acknowledges, they lean towards large banks and insurers, and the request for nonbank examples is well judged. We would propose further case studies to illustrate proportionate approaches in fintechs, credit brokers and consumer finance firms, particularly illustrating where SMF holders may hold accountability for a range of functions recognising their less complex governance structure. We offer one below from our own domain, which shows Sound Practices 8, 9, 10 and 12 working together in a conduct setting.

Case study: AI-assisted review and monitoring of financial promotions (RegTech provider)

A specialist technology provider supplies regulated firms with an AI system that reviews financial promotions and customer communications against the applicable rules and firm policy before publication, and monitors live communications afterwards. The system uses generative AI to read each item, identify potential breaches such as an unbalanced presentation of risk and benefit, missing or inadequate risk warnings, or unsubstantiated claims, and to draft a plain-language explanation and suggested fix for each issue it raises.

The design mirrors the report's sound practices. Explainability is built in at the point of use: every flag is tied to the specific rule and the specific words that triggered it, so a reviewer can see why the system reached its view (Sound Practice 8). Performance is assessed against reviewer decisions over time, with precision and recall tracked by communication type and thresholds set before deployment (Sound Practice 9). Human oversight is meaningful rather than nominal: the system triages and recommends, but a qualified compliance reviewer makes and records the approval decision, and the audit trail captures both the system's reasoning and the human's (Sound Practice 10). Because the capability is delivered by a third party, the firm receives the validation evidence, audit access and documented limitations it needs to stay accountable for outcomes (Sound Practice 12).

Reported outcomes for adopting firms include a reduction of approximately 50-75 per cent in review turnaround time, the ability to review up to five times more communications without adding headcount, and more consistent application of the rules across high volumes of material. The human approval step and the audit trail are preserved, so faster review does not come at the cost of accountability.

6. Do the case studies in this report provide actionable insights for financial institutions in their responsible AI adoption?

Yes. They are concrete enough to be useful, which is not always the case with material of this kind. Two suggestions to make them more so. It would help readers if each case study stated explicitly which sound practices it illustrates, as we have done above, so the link between the example and the guidance is immediate. The report would also benefit from at least one example of a smaller, less complex firm adopting AI through a specialist provider rather than building it. Most firms are not global systemically important banks, and proportionality is easier to grasp when it is shown in the long tail as well as at the top. The report already notes that smaller firms in some markets are partnering with AI providers for tasks such as credit scoring, and a similar example in a conduct or operations setting would round out the picture. The case studies could benefit from focus on key controls, monitoring and assurance approach, and lessons learned to help translate into more actionable insights.

7. Are the definitions in the glossary clear and aligned with industry sound practices, including recent developments in AI?

The glossary is clear and well aligned with current usage. A few additions would make it more complete, given concepts the report already relies on in the body text:

*Retrieval-augmented generation (RAG), which appears in the text and case studies but is not defined.

*Prompt injection and jailbreaking, explained in line but worth naming as attack types in their own right.

*Model and system cards, given the weight the report places on them in Sound Practice 12.

*Red-teaming, referred to in Sound Practices 9 and 10.

*The human-oversight family (human-in-the-loop, on-the-loop, in-command), which the report defines in Sound Practice 10 and could usefully consolidate.

*Shadow AI, which is a term used but not defined and could be open to differing interpretations.

Given the report's attention to consumer protection and market conduct, the FSB might also add a short, neutral definition of financial promotion or financial communication, since the meaning varies across jurisdictions and the term carries specific regulatory weight in several of them. We would also suggest the definition of agentic AI note that autonomy is a spectrum rather than a binary, which is consistent with how the report treats the subject elsewhere.

8. Are there any comments you would like to make on this report? Please fill in.

Finspector builds AI to help regulated firms review, approve and monitor their financial promotions and customer advertisements. Our users are the compliance, risk and marketing teams across financial services, and our work sits squarely in the conduct and consumer protection territory this consultation touches on. We see the report from two angles at once. We are a user of AI, subject to the same governance and lifecycle expectations the report sets out, and we are a third-party AI provider to regulated firms, which gives us a practical view of the issues raised by Sound Practices 8, 10 and 12 in particular.

We welcome the FSB's work. The report is measured, proportionate and grounded in how AI is actually being adopted. Framing this as a menu of sound practices rather than a prescriptive standard is the right choice for a technology moving this quickly, and the emphasis on proportionality, technology neutrality and outcomes should age well. Our comments are offered in that spirit. They are meant to sharpen a strong report, not to redirect it. We have concentrated on the areas where we can add something from direct experience, and where we have no material concern we have said so briefly.

We support this report and the direction it sets. Responsible adoption is the right organising idea, and the proportionate, outcomes-based approach gives firms a workable path to capturing the benefits of AI while keeping risk within appetite. We would be happy to contribute further, including a metrics-backed case study or input on sector-specific transparency and assurance tooling. We are content for this response to be published. Thank you for the opportunity to respond.